Red Hat Enterprise Linux 5

Deployment Guide

Deployment, configuration and administration
of Red Hat Enterprise Linux 5

Deployment_Guide

Red Hat Enterprise Linux 5 Deployment Guide
Deployment, configuration and administration of Red Hat Enterprise Linux 5
엮음 8

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701

The Deployment Guide documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 5.

머리글

Red Hat Enterprise Linux 활용 가이드에 오신 것을 환영합니다!

Red Hat Enterprise Linux 활용 가이드에는 Red Hat Enterprise Linux 시스템을 사용자의 요구에 맞게 사용자 설정하는 방법에 관한 정보가 포함되어 있습니다. 시스템 설정 및 사용자 설정에 대한 포괄적이고, 실전 중심의 가이드를 찾고 계신 경우, 이 메뉴얼을 사용하시기 바랍니다.

이 메뉴얼에서는 다음과 같은 주제에 관하여 논의합니다:

- NIC (network interface card) 설정하기
- VPN (Virtual Private Network) 설정하기
- Samba 공유 설정하기
- RPM을 사용하여 소프트웨어를 관리하기
- 시스템에 관한 정보 선택하기
- 커널 업그레이드하기

이 메뉴얼은 다음과 같은 주요 범주로 나뉘어져 있습니다:

- 파일 시스템
- 패키지 관리
- 네트워크 관련 설정
- 시스템 설정
- 시스템 감시
- 커널 및 드라이버 설정
- 보안 및 인증
- Red Hat 교육 및 자격증

이 가이드는 사용자가 Red Hat Enterprise Linux 시스템에 관해 기본적인 내용을 이해하고 있다고 간주합니다. Red Hat Enterprise Linux 설치를 위한 도움이 필요하실 경우, Red Hat Enterprise Linux 설치 가이드를 참조하시기 바랍니다.

1. Document Conventions

In this manual, certain words are represented in different fonts, typefaces, sizes, and weights. This highlighting is systematic; different words are represented in the same style to indicate their inclusion in a specific category. The types of words that are represented this way include the following:

command

Linux commands (and other operating system commands, when used) are represented this way. This style should indicate to you that you can type the word or phrase on the command line and press Enter to invoke a command. Sometimes a command contains words that would be displayed in a different style on their own (such as file names). In these cases, they are considered to be part of the command, so the entire phrase is displayed as a command. For example:

Use the cat testfile command to view the contents of a file, named testfile, in the current working directory.

file name
: File names, directory names, paths, and RPM package names are represented this way. This style indicates that a particular file or directory exists with that name on your system. Examples:

  The .bashrc file in your home directory contains bash shell definitions and aliases for your own use.

  The /etc/fstab file contains information about different system devices and file systems.

  Install the webalizer RPM if you want to use a Web server log file analysis program.

application
: This style indicates that the program is an end-user application (as opposed to system software). For example:

  Use Mozilla to browse the Web.

key
: A key on the keyboard is shown in this style. For example:

  To use Tab completion to list particular files in a directory, type ls, then a character, and finally the Tab key. Your terminal displays the list of files in the working directory that begin with that character.

key+combination
: A combination of keystrokes is represented in this way. For example:

  The Ctrl+Alt+Backspace key combination exits your graphical session and returns you to the graphical login screen or the console.

text found on a GUI interface
: A title, word, or phrase found on a GUI interface screen or window is shown in this style. Text shown in this style indicates a particular GUI screen or an element on a GUI screen (such as text associated with a checkbox or field). Example:

  Select the Require Password checkbox if you would like your screensaver to require a password before stopping.

top level of a menu on a GUI screen or window
: A word in this style indicates that the word is the top level of a pulldown menu. If you click on the word on the GUI screen, the rest of the menu should appear. For example:

  Under File on a GNOME terminal, the New Tab option allows you to open multiple shell prompts in the same window.

  Instructions to type in a sequence of commands from a GUI menu look like the following example:

  Go to Applications (the main menu on the panel) > Programming > Emacs Text Editor to start the Emacs text editor.

button on a GUI screen or window
: This style indicates that the text can be found on a clickable button on a GUI screen. For example:

Click on the Back button to return to the webpage you last viewed.

computer output

Text in this style indicates text displayed to a shell prompt such as error messages and responses to commands. For example:

The ls command displays the contents of a directory. For example:

```
Desktop      about.html      logs       paulwesterberg.png
Mail      backupfiles      mail       reports
```

The output returned in response to the command (in this case, the contents of the directory) is shown in this style.

prompt

A prompt, which is a computer's way of signifying that it is ready for you to input something, is shown in this style. Examples:

$

#

[stephen@maturin stephen]$

leopard login:

user input

Text that the user types, either on the command line or into a text box on a GUI screen, is displayed in this style. In the following example, text is displayed in this style:

To boot your system into the text based installation program, you must type in the text command at the boot: prompt.

<replaceable>

Text used in examples that is meant to be replaced with data provided by the user is displayed in this style. In the following example, <version-number> is displayed in this style:

The directory for the kernel source is /usr/src/kernels/<version-number>/, where <version-number> is the version and type of kernel installed on this system.

Additionally, we use several different strategies to draw your attention to certain pieces of information. In order of urgency, these items are marked as a note, tip, important, caution, or warning. For example:

Note

Remember that Linux is case sensitive. In other words, a rose is not a ROSE is not a rOsE.

Tip

The directory /usr/share/doc/ contains additional documentation for packages installed on your system.

Important

If you modify the DHCP configuration file, the changes do not take effect until you restart the DHCP daemon.

Caution

Do not perform routine tasks as root — use a regular user account unless you need to use the root account for system administration tasks.

Warning

Be careful to remove only the necessary partitions. Removing other partitions could result in data loss or a corrupted system environment.

2. 여러분의 의견을 기다리고 있습니다

If you find an error in the Red Hat Enterprise Linux Deployment Guide, or if you have thought of a way to make this manual better, we would like to hear from you! Submit a report in Bugzilla (http://bugzilla.redhat.com/bugzilla/) against the component Deployment_Guide.

자료 문서 개선을 위한 제안이 있으시면, 최대한 명확히 설명해 주시기 바랍니다. 오류를 발견하셨다면, 저희가 쉽게 식별할 수 있도록 섹션 번호와 주위의 문장들을 함께 보내주시기 바랍니다.

부 I. 파일 시스템

File system refers to the files and directories stored on a computer. A file system can have different formats called file system types. These formats determine how the information is stored as files and directories. Some file system types store redundant copies of the data, while some file system types make hard drive access faster. This part discusses the ext3, swap, RAID, and LVM file system types. It also discusses the parted utility to manage partitions and access control lists (ACLs) to customize file permissions.

File System Structure

1.1. Why Share a Common Structure?

The file system structure is the most basic level of organization in an operating system. Almost all of the ways an operating system interacts with its users, applications, and security model are dependent upon the way it organizes files on storage devices. Providing a common file system structure ensures users and programs are able to access and write files.

File systems break files down into two logical categories:

- Shareable vs. unshareable files

- Variable vs. static files

Shareable files are those that can be accessed locally and by remote hosts; unshareable files are only available locally. Variable files, such as documents, can be changed at any time; static files, such as binaries, do not change without an action from the system administrator.

The reason for looking at files in this manner is to help correlate the function of the file with the permissions assigned to the directories which hold them. The way in which the operating system and its users interact with a given file determines the directory in which it is placed, whether that directory is mounted with read-only or read/write permissions, and the level of access each user has to that file. The top level of this organization is crucial. Access to the underlying directories can be restricted or security problems could manifest themselves if, from the top level down, it does not adhere to a rigid structure.

1.2. Overview of File System Hierarchy Standard (FHS)

Red Hat Enterprise Linux uses the Filesystem Hierarchy Standard (FHS) file system structure, which defines the names, locations, and permissions for many file types and directories.

The FHS document is the authoritative reference to any FHS-compliant file system, but the standard leaves many areas undefined or extensible. This section is an overview of the standard and a description of the parts of the file system not covered by the standard.

Compliance with the standard means many things, but the two most important are compatibility with other compliant systems and the ability to mount a /usr/ partition as read-only. This second point is important because the directory contains common executables and should not be changed by users. Also, since the /usr/ directory is mounted as read-only, it can be mounted from the CD-ROM or from another machine via a read-only NFS mount.

1.2.1. FHS Organization

The directories and files noted here are a small subset of those specified by the FHS document. Refer to the latest FHS document for the most complete information.

The complete standard is available online at http://www.pathname.com/fhs/[1].

[1] http://www.pathname.com/fhs

1.2.1.1. The /boot/ Directory

The /boot/ directory contains static files required to boot the system, such as the Linux kernel. These files are essential for the system to boot properly.

Warning

Do not remove the /boot/ directory. Doing so renders the system unbootable.

1.2.1.2. The /dev/ Directory

The /dev/ directory contains device nodes that either represent devices that are attached to the system or virtual devices that are provided by the kernel. These device nodes are essential for the system to function properly. The udev daemon takes care of creating and removing all these device nodes in /dev/.

Devices in the /dev directory and subdirectories are either character (providing only a serial stream of input/output) or block (accessible randomly). Character devices include mouse, keyboard, modem while block devices include hard disk, floppy drive etc. If you have GNOME or KDE installed in your system, devices such as external drives or cds are automatically detected when connected (e.g via usb) or inserted (e.g via CD or DVD drive) and a popup window displaying the contents is automatically displayed. Files in the /dev directory are essential for the system to function properly.

표 1.1. Examples of common files in the /dev

File	Description
/dev/hda	The master device on primary IDE channel.
/dev/hdb	The slave device on primary IDE channel.
/dev/tty0	The first virtual console.
/dev/tty1	The second virtual console.
/dev/sda	The first device on primary SCSI or SATA channel.
/dev/lp0	The first parallel port.

1.2.1.3. The /etc/ Directory

The /etc/ directory is reserved for configuration files that are local to the machine. No binaries are to be placed in /etc/. Any binaries that were once located in /etc/ should be placed into /sbin/ or /bin/.

Examples of directories in /etc are the X11/ and skel/:

```
/etc
  |- X11/
  |- skel/
```

The /etc/X11/ directory is for X Window System configuration files, such as xorg.conf. The /etc/skel/ directory is for "skeleton" user files, which are used to populate a home directory when a user is first created. Applications also store their configuration files in this directory and may reference them when they are executed.

1.2.1.4. The /lib/ Directory

The /lib/ directory should contain only those libraries needed to execute the binaries in /bin/ and /sbin/. These shared library images are particularly important for booting the system and executing commands within the root file system.

1.2.1.5. The /media/ Directory

The /media/ directory contains subdirectories used as mount points for removable media such as usb storage media, DVDs, CD-ROMs, and Zip disks.

1.2.1.6. The /mnt/ Directory

The /mnt/ directory is reserved for temporarily mounted file systems, such as NFS file system mounts. For all removable media, please use the /media/ directory. Automatically detected removable media will be mounted in the /media directory.

> **Note**
>
> The /mnt directory must not be used by installation programs.

1.2.1.7. The /opt/ Directory

The /opt/ directory provides storage for most application software packages.

A package placing files in the /opt/ directory creates a directory bearing the same name as the package. This directory, in turn, holds files that otherwise would be scattered throughout the file system, giving the system administrator an easy way to determine the role of each file within a particular package.

For example, if sample is the name of a particular software package located within the /opt/ directory, then all of its files are placed in directories inside the /opt/sample/ directory, such as /opt/sample/bin/ for binaries and /opt/sample/man/ for manual pages.

Packages that encompass many different sub-packages, data files, extra fonts, clipart etc are also located in the /opt/ directory, giving that large package a way to organize itself. In this way, our sample package may have different tools that each go in their own sub-directories, such as /opt/sample/tool1/ and /opt/sample/tool2/, each of which can have their own bin/, man/, and other similar directories.

1.2.1.8. The /proc/ Directory

The /proc/ directory contains special files that either extract information from or send information to the kernel. Examples include system memory, cpu information, hardware configuration etc.

Due to the great variety of data available within /proc/ and the many ways this directory can be used to communicate with the kernel, an entire chapter has been devoted to the subject. For more information, refer to 4장. The proc File System.

1.2.1.9. The /sbin/ Directory

The /sbin/ directory stores executables used by the root user. The executables in /sbin/ are used at boot time, for system administration and to perform system recovery operations. Of this directory, the FHS says:

> /sbin contains binaries essential for booting, restoring, recovering, and/or repairing the system in addition to the binaries in /bin. Programs executed after /usr/ is known to be mounted (when there are no problems) are generally placed into /usr/sbin. Locally-installed system administration programs should be placed into /usr/local/sbin.

At a minimum, the following programs should be in /sbin/:

```
arp, clock,
halt, init,
fsck.*, grub,
ifconfig, mingetty,
mkfs.*, mkswap,
reboot, route,
shutdown, swapoff,
swapon
```

1.2.1.10. The /srv/ Directory

The /srv/ directory contains site-specific data served by your system running Red Hat Enterprise Linux. This directory gives users the location of data files for a particular service, such as FTP, WWW, or CVS. Data that only pertains to a specific user should go in the /home/ directory.

1.2.1.11. The /sys/ Directory

The /sys/ directory utilizes the new sysfs virtual file system specific to the 2.6 kernel. With the increased support for hot plug hardware devices in the 2.6 kernel, the /sys/ directory contains information similarly held in /proc/, but displays a hierarchical view of specific device information in regards to hot plug devices.

1.2.1.12. The /usr/ Directory

The /usr/ directory is for files that can be shared across multiple machines. The /usr/ directory is often on its own partition and is mounted read-only. At a minimum, the following directories should be subdirectories of /usr/:

```
/usr
   |- bin/
   |- etc/
   |- games/
   |- include/
   |- kerberos/
   |- lib/
   |- libexec/
   |- local/
   |- sbin/
   |- share/
   |- src/
   |- tmp -> ../var/tmp/
```

Under the /usr/ directory, the bin/ subdirectory contains executables, etc/ contains system-wide configuration files, games is for games, include/ contains C header files, kerberos/ contains binaries and other Kerberos-related files, and lib/ contains object files and libraries that are not designed to be directly utilized by users or shell scripts. The libexec/ directory contains small helper programs called by other programs, sbin/ is for system administration binaries (those that do not belong in the /sbin/ directory), share/ contains files that are not architecture-specific, src/ is for source code.

1.2.1.13. The /usr/local/ Directory

The FHS says:

> The /usr/local hierarchy is for use by the system administrator when installing software locally. It needs to be safe from being overwritten when the system software is updated. It may be used for programs and data that are shareable among a group of hosts, but not found in /usr.

The /usr/local/ directory is similar in structure to the /usr/ directory. It has the following subdirectories, which are similar in purpose to those in the /usr/ directory:

```
/usr/local
 |- bin/
 |- etc/
 |- games/
 |- include/
 |- lib/
 |- libexec/
 |- sbin/
 |- share/
 |- src/
```

In Red Hat Enterprise Linux, the intended use for the /usr/local/ directory is slightly different from that specified by the FHS. The FHS says that /usr/local/ should be where software that is to remain safe from system software upgrades is stored. Since software upgrades can be performed safely with RPM Package Manager (RPM), it is not necessary to protect files by putting them in /usr/local/. Instead, the /usr/local/ directory is used for software that is local to the machine.

For instance, if the /usr/ directory is mounted as a read-only NFS share from a remote host, it is still possible to install a package or program under the /usr/local/ directory.

1.2.1.14. The /var/ Directory

Since the FHS requires Linux to mount /usr/ as read-only, any programs that write log files or need spool/ or lock/ directories should write them to the /var/ directory. The FHS states /var/ is for:

> ...variable data files. This includes spool directories and files, administrative and logging data, and transient and temporary files.

Below are some of the directories found within the /var/ directory:

```
/var
    |- account/
    |- arpwatch/
    |- cache/
    |- crash/
    |- db/
    |- empty/
    |- ftp/
    |- gdm/
    |- kerberos/
    |- lib/
    |- local/
    |- lock/
    |- log/
    |- mail -> spool/mail/
    |- mailman/
    |- named/
    |- nis/
    |- opt/
    |- preserve/
    |- run/
    +- spool/
        |- at/
        |- clientmqueue/
        |- cron/
        |- cups/
        |- exim/
        |- lpd/
        |- mail/
        |- mailman/
        |- mqueue/
        |- news/
        |- postfix/
        |- repackage/
        |- rwho/
        |- samba/
        |- squid/
        |- squirrelmail/
        |- up2date/
        |- uucp
        |- uucppublic/
        |- vbox/
|- tmp/
|- tux/
|- www/
|- yp/
```

System log files, such as messages and lastlog, go in the /var/log/ directory. The /var/lib/rpm/ directory contains RPM system databases. Lock files go in the /var/lock/ directory, usually in directories for the program using the file. The /var/spool/ directory has subdirectories for programs in which data files are stored.

1.3. Special File Locations Under Red Hat Enterprise Linux

Red Hat Enterprise Linux extends the FHS structure slightly to accommodate special files.

Most files pertaining to RPM are kept in the /var/lib/rpm/ directory. For more information on RPM, refer to the chapter 11장. RPM을 사용한 패키지 관리.

The /var/cache/yum/ directory contains files used by the Package Updater, including RPM header information for the system. This location may also be used to temporarily store RPMs downloaded while updating the system. For more information about Red Hat Network, refer to 14장. Product Subscriptions and Entitlements.

Another location specific to Red Hat Enterprise Linux is the /etc/sysconfig/ directory. This directory stores a variety of configuration information. Many scripts that run at boot time use the files in this directory. Refer to 30장. The sysconfig Directory for more information about what is within this directory and the role these files play in the boot process.

Using the mount Command

On Linux, UNIX, and similar operating systems, file systems on different partitions and removable devices like CDs, DVDs, or USB flash drives can be attached to a certain point (that is, the mount point) in the directory tree, and detached again. To attach or detach a file system, you can use the mount or umount command respectively. This chapter describes the basic usage of these commands, and covers some advanced topics such as moving a mount point or creating shared subtrees.

2.1. Listing Currently Mounted File Systems

To display all currently attached file systems, run the mount command with no additional arguments:

```
mount
```

This command displays the list of known mount points. Each line provides important information about the device name, the file system type, the directory in which it is mounted, and relevant mount options in the following form:

device on directory type type (options)

By default, the output includes various virtual file systems such as sysfs, tmpfs, and others. To display only the devices with a certain file system type, supply the -t option on the command line:

```
mount -t type
```

For a list of common file system types, refer to 표 2.1. "Common File System Types" . For an example on how to use the mount command to list the mounted file systems, see 예 2.1. "Listing Currently Mounted ext3 File Systems" .

> 예 2.1. Listing Currently Mounted ext3 File Systems
>
> Usually, both / and /boot partitions are formatted to use ext3. To display only the mount points that use this file system, type the following at a shell prompt:
>
> ```
> ~]$ mount -t ext3
> /dev/mapper/VolGroup00-LogVol00 on / type ext3 (rw)
> /dev/vda1 on /boot type ext3 (rw)
> ```

2.2. Mounting a File System

To attach a certain file system, use the mount command in the following form:

```
mount [option…] device directory
```

When the mount command is run, it reads the content of the /etc/fstab configuration file to see if the given file system is listed. This file contains a list of device names and the directory in which the selected file systems should be mounted, as well as the file system type and mount options. Because of this, when you are mounting a file system that is specified in this file, you can use one of the following variants of the command:

```
mount [option…] directory
mount [option…] device
```

Note that unless you are logged in as root, you must have permissions to mount the file system (see 2.2.2절. "Specifying the Mount Options").

2.2.1. Specifying the File System Type

In most cases, mount detects the file system automatically. However, there are certain file systems, such as NFS (Network File System) or CIFS (Common Internet File System), that are not recognized, and need to be specified manually. To specify the file system type, use the mount command in the following form:

```
mount -t type device directory
```

표 2.1. "Common File System Types" provides a list of common file system types that can be used with the mount command. For a complete list of all available file system types, consult the relevant manual page as referred to in 2.4.1절. "Installed Documentation" .

표 2.1. Common File System Types

Type	Description
ext2	The ext2 file system.
ext3	The ext3 file system.
iso9660	The ISO 9660 file system. It is commonly used by optical media, typically CDs.
jfs	The JFS file system created by IBM.
nfs	The NFS file system. It is commonly used to access files over the network.
nfs4	The NFSv4 file system. It is commonly used to access files over the network.
ntfs	The NTFS file system. It is commonly used on machines that are running the Windows operating system.
udf	The UDF file system. It is commonly used by optical media, typically DVDs.
vfat	The FAT file system. It is commonly used on machines that are running the Windows operating system, and on certain digital media such as USB flash drives or floppy disks.

See 예 2.2. "Mounting a USB Flash Drive" for an example usage.

예 2.2. Mounting a USB Flash Drive

Older USB flash drives often use the FAT file system. Assuming that such drive uses the /dev/sdc1 device and that the /media/flashdisk/ directory exists, you can mount it to this directory by typing the following at a shell prompt as root:

```
~]# mount -t vfat /dev/sdc1 /media/flashdisk
```

2.2.2. Specifying the Mount Options

To specify additional mount options, use the command in the following form:

```
mount -o options
```

When supplying multiple options, do not insert a space after a comma, or mount will incorrectly interpret the values following spaces as additional parameters.

표 2.2. "Common Mount Options" provides a list of common mount options. For a complete list of all available options, consult the relevant manual page as referred to in 2.4.1절. "Installed Documentation" .

표 2.2. Common Mount Options

Option	Description
async	Allows the asynchronous input/output operations on the file system.
auto	Allows the file system to be mounted automatically using the mount -a command.
defaults	Provides an alias for async,auto,dev,exec,nouser,rw,suid.
exec	Allows the execution of binary files on the particular file system.
loop	Mounts an image as a loop device.
noauto	Disallows the automatic mount of the file system using the mount -a command.
noexec	Disallows the execution of binary files on the particular file system.
nouser	Disallows an ordinary user (that is, other than root) to mount and unmount the file system.
remount	Remounts the file system in case it is already mounted.
ro	Mounts the file system for reading only.
rw	Mounts the file system for both reading and writing.
user	Allows an ordinary user (that is, other than root) to mount and unmount the file system.

See 예 2.3. "Mounting an ISO Image" for an example usage.

> 예 2.3. Mounting an ISO Image
>
> An ISO image (or a disk image in general) can be mounted by using the loop device. Assuming that the ISO image of the Fedora 14 installation disc is present in the current working directory and that the /media/cdrom/ directory exists, you can mount the image to this directory by running the following command as root:
>
> ```
> ~]# mount -o ro,loop Fedora-14-x86_64-Live-Desktop.iso /media/cdrom
> ```
>
> Note that ISO 9660 is by design a read-only file system.

2.2.3. Sharing Mounts

Occasionally, certain system administration tasks require access to the same file system from more than one place in the directory tree (for example, when preparing a chroot environment). To address such requirements, the mount command implements the --bind option that provides a means for duplicating certain mounts. Its usage is as follows:

```
mount --bind old_directory new_directory
```

Although the above command allows a user to access the file system from both places, it does not apply on the file systems that are mounted within the original directory. To include these mounts as well, type:

```
mount --rbind old_directory new_directory
```

Additionally, to provide as much flexibility as possible, Red Hat Enterprise Linux 5.8 implements the functionality known as shared subtrees. This feature allows you to use the following four mount types:

Shared Mount

A shared mount allows you to create an exact replica of a given mount point. When a shared mount is created, any mount within the original mount point is reflected in it, and vice versa. To create a shared mount, type the following at a shell prompt:

```
mount --make-shared mount_point
```

Alternatively, you can change the mount type for the selected mount point and all mount points under it:

```
mount --make-rshared mount_point
```

See 예 2.4. “Creating a Shared Mount Point” for an example usage.

예 2.4. Creating a Shared Mount Point

There are two places where other file systems are commonly mounted: the /media directory for removable media, and the /mnt directory for temporarily mounted file systems. By using a shared mount, you can make these two directories share the same content. To do so, as root, mark the /media directory as “shared” :

```
~]# mount --bind /media /media
~]# mount --make-shared /media
```

Then create its duplicate in /mnt by using the following command:

```
~]# mount --bind /media /mnt
```

You can now verify that a mount within /media also appears in /mnt. For example, if you have non-empty media in your CD-ROM drive and the /media/cdrom/ directory exists, run the following commands:

```
~]# mount /dev/cdrom /media/cdrom
~]# ls /media/cdrom
EFI   GPL   isolinux   LiveOS
~]# ls /mnt/cdrom
EFI   GPL   isolinux   LiveOS
```

Similarly, you can verify that any file system mounted in the /mnt directory is reflected in /media. For instance, if you have a non-empty USB flash drive that uses the /dev/sdc1 device plugged in and the /mnt/flashdisk/ directory is present, type:

```
~]# mount /dev/sdc1 /mnt/flashdisk
~]# ls /media/flashdisk
en-US   publican.cfg
~]# ls /mnt/flashdisk
en-US   publican.cfg
```

Slave Mount

A slave mount allows you to create a limited duplicate of a given mount point. When a slave mount is created, any mount within the original mount point is reflected in it, but no mount within a slave mount is reflected in its original. To create a slave mount, type the following at a shell prompt:

```
mount --make-slave mount_point
```

Alternatively, you can change the mount type for the selected mount point and all mount points under it:

```
mount --make-rslave mount_point
```

See 예 2.5. "Creating a Slave Mount Point" for an example usage.

예 2.5. Creating a Slave Mount Point

Imagine you want the content of the /media directory to appear in /mnt as well, but you do not want any mounts in the /mnt directory to be reflected in /media. To do so, as root, first mark the /media directory as "shared" :

```
~]# mount --bind /media /media
~]# mount --make-shared /media
```

Then create its duplicate in /mnt, but mark it as "slave" :

```
~]# mount --bind /media /mnt
~]# mount --make-slave /mnt
```

You can now verify that a mount within /media also appears in /mnt. For example, if you have non-empty media in your CD-ROM drive and the /media/cdrom/ directory exists, run the following commands:

```
~]# mount /dev/cdrom /media/cdrom
~]# ls /media/cdrom
EFI  GPL  isolinux  LiveOS
~]# ls /mnt/cdrom
EFI  GPL  isolinux  LiveOS
```

You can also verify that file systems mounted in the /mnt directory are not reflected in /media. For instance, if you have a non-empty USB flash drive that uses the /dev/sdc1 device plugged in and the /mnt/flashdisk/ directory is present, type: :

```
~]# mount /dev/sdc1 /mnt/flashdisk
~]# ls /media/flashdisk
~]# ls /mnt/flashdisk
en-US  publican.cfg
```

Private Mount

A private mount allows you to create an ordinary mount. When a private mount is created, no subsequent mounts within the original mount point are reflected in it, and no mount within a private mount is reflected in its original. To create a private mount, type the following at a shell prompt:

```
mount --make-private mount_point
```

Alternatively, you can change the mount type for the selected mount point and all mount points under it:

```
mount --make-rprivate mount_point
```

See 예 2.6. “Creating a Private Mount Point” for an example usage.

예 2.6. Creating a Private Mount Point

Taking into account the scenario in 예 2.4. “Creating a Shared Mount Point”, assume that you have previously created a shared mount point by using the following commands as root:

```
~]# mount --bind /media /media
~]# mount --make-shared /media
~]# mount --bind /media /mnt
```

To mark the /mnt directory as “private”, type:

```
~]# mount --make-private /mnt
```

You can now verify that none of the mounts within /media appears in /mnt. For example, if you have non-empty media in your CD-ROM drive and the /media/cdrom/ directory exists, run the following commands:

```
~]# mount /dev/cdrom /media/cdrom
~]# ls /media/cdrom
EFI  GPL  isolinux  LiveOS
~]# ls /mnt/cdrom
~]#
```

You can also verify that file systems mounted in the /mnt directory are not reflected in /media. For instance, if you have a non-empty USB flash drive that uses the /dev/sdc1 device plugged in and the /mnt/flashdisk/ directory is present, type:

```
~]# mount /dev/sdc1 /mnt/flashdisk
~]# ls /media/flashdisk
~]# ls /mnt/flashdisk
en-US  publican.cfg
```

Unbindable Mount

An unbindable mount allows you to prevent a given mount point from being duplicated whatsoever. To create an unbindable mount, type the following at a shell prompt:

```
mount --make-unbindable mount_point
```

Alternatively, you can change the mount type for the selected mount point and all mount points under it:

```
mount --make-runbindable mount_point
```

See 예 2.7. “Creating an Unbindable Mount Point” for an example usage.

예 2.7. Creating an Unbindable Mount Point

To prevent the /media directory from being shared, as root, type the following at a shell prompt:

```
~]# mount --bind /media /media
~]# mount --make-unbindable /media
```

This way, any subsequent attempt to make a duplicate of this mount will fail with an error:

```
~]# mount --bind /media /mnt
mount: wrong fs type, bad option, bad superblock on /media/,
        missing code page or other error
        In some cases useful info is found in syslog - try
        dmesg | tail  or so
```

2.2.4. Moving a Mount Point

To change the directory in which a file system is mounted, use the following command:

```
mount --move old_directory new_directory
```

See 예 2.8. "Moving an Existing NFS Mount Point" for an example usage.

예 2.8. Moving an Existing NFS Mount Point

Imagine that you have an NFS storage that contains user directories. Assuming that this storage is already mounted in /mnt/userdirs/, as root, you can move this mount point to /home by using the following command:

```
~]# mount --move /mnt/userdirs /home
```

To verify the mount point has been moved, list the content of both directories:

```
~]# ls /mnt/userdirs
~]# ls /home
jill  joe
```

2.3. Unmounting a File System

To detach a previously mounted file system, use either of the following variants of the umount command:

```
umount directory
umount device
```

Note that unless you are logged in as root, you must have permissions to unmount the file system (see 2.2.2절. "Specifying the Mount Options"). See 예 2.9. "Unmounting a CD" for an example usage.

Important: Make Sure the File System Is Not in Use

When a file system is in use (for example, when a process is reading a file on this file system), running the umount command will fail with an error. To determine which processes are accessing the file system, use the fuser command in the following form:

```
fuser -m directory
```

For example, to list the processes that are accessing a file system mounted to the /media/cdrom/ directory, type:

```
~]$ fuser -m /media/cdrom
/media/cdrom:            1793   2013   2022   2435  10532c  10672c
```

예 2.9. Unmounting a CD

To unmount a CD that was previously mounted to the /media/cdrom/ directory, type the following at a shell prompt:

```
~]$ umount /media/cdrom
```

2.4. Additional Resources

The following resources provide an in-depth documentation on the subject.

2.4.1. Installed Documentation

- man 8 mount — The manual page for the mount command that provides a full documentation on its usage.

- man 8 umount — The manual page for the umount command that provides a full documentation on its usage.

- man 5 fstab — The manual page providing a thorough description of the /etc/fstab file format.

2.4.2. Useful Websites

- Shared subtrees[1] — An LWN article covering the concept of shared subtrees.

- sharedsubtree.txt[2] — Extensive documentation that is shipped with the shared subtrees patches.

[1] http://lwn.net/Articles/159077/

[2] http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=blob_plain;f=Documentation/sharedsubtree.txt;hb=ce9e3d9953c8cb67001719b5516da2928e956be4

ext3 파일 시스템

기본 파일 시스템은 저널링 ext3 파일 시스템입니다.

3.1. ext3의 기능

ext3 파일 시스템은 ext2 형식의 기능을 강화시킨 파일 시스템 버전으로서, ext3 파일 시스템의 장점은 다음과 같습니다:

가용성 (Availability)

After an unexpected power failure or system crash (also called an unclean system shutdown), each mounted ext2 file system on the machine must be checked for consistency by the e2fsck program. This is a time-consuming process that can delay system boot time significantly, especially with large volumes containing a large number of files. During this time, any data on the volumes is unreachable.

ext3 파일 시스템의 저널링 기능을 이용하면, 시스템이 비정상적으로 종료된 후에도 이러한 시간 소모적인 파일 시스템 검사 작업을 수행할 필요가 전혀 없습니다. ext3 파일 시스템에서는 하드 드라이브가 고장난 경우와 같이 특정 하드웨어에 문제가 있는 경우에만 일관성 검사를 수행합니다. 시스템이 비정상적으로 종료된 후 ext3 파일 시스템을 복구하는데 걸리는 시간은 파일 시스템의 크기나 파일의 숫자에 따라 결정되지 않고; 파일 시스템의 일관성을 유지하는데 사용되는 저널 (journal)의 크기에 따라 결정됩니다. 하드웨어의 속도에 따라서 기본 저널 크기의 경우, 일반적으로 파일 시스템을 복구하는데 1초가 걸립니다.

데이터 신뢰성 강화 (Data Integrity)

The ext3 file system prevents loss of data integrity in the event that an unclean system shutdown occurs. The ext3 file system allows you to choose the type and level of protection that your data receives. By default, the ext3 volumes are configured to keep a high level of data consistency with regard to the state of the file system.

보다 빠른 수행 속도

Despite writing some data more than once, ext3 has a higher throughput in most cases than ext2 because ext3's journaling optimizes hard drive head motion. You can choose from three journaling modes to optimize speed, but doing so means trade-offs in regards to data integrity if the system was to fail.

손쉬운 변환 과정

It is easy to migrate from ext2 to ext3 and gain the benefits of a robust journaling file system without reformatting. Refer to 3.3절. "ext3 파일 시스템으로 변환하기" for more on how to perform this task.

The following sections walk you through the steps for creating and tuning ext3 partitions. For ext2 partitions, skip the partitioning and formatting sections below and go directly to 3.3절. "ext3 파일 시스템으로 변환하기".

3.2. ext3 파일 시스템 생성하기

설치를 마치신 후, 가끔씩 새로운 ext3 파일 시스템을 생성해야할 경우가 있습니다. 예를 들어, 시스템에 새로운 디스크 드라이브를 추가하실 경우, 드라이브를 파티션하신 후 ext3 파일 시스템을 사용 가능합니다.

ext3 파일 시스템을 생성하는 방법은 다음과 같습니다:

1. Format the partition with the ext3 file system using mkfs.

2. Label the partition using e2label.

3.3. ext3 파일 시스템으로 변환하기

The tune2fs allows you to convert an ext2 filesystem to ext3.

> **Note**
>
> Always use the e2fsck utility to check your filesystem before and after using tune2fs. A default installation of Red Hat Enterprise Linux uses ext3 for all file systems.

To convert an ext2 filesystem to ext3, log in as root and type the following command in a terminal:

```
tune2fs -j <block_device>
```

where <block_device> contains the ext2 filesystem you wish to convert.

A valid block device could be one of two types of entries:

- A mapped device — A logical volume in a volume group, for example, /dev/mapper/VolGroup00-LogVol02.

- A static device — A traditional storage volume, for example, /dev/hdbX, where hdb is a storage device name and X is the partition number.

Issue the df command to display mounted file systems.

For the remainder of this section, the sample commands use the following value for the block device:

```
/dev/mapper/VolGroup00-LogVol02
```

You must recreate the initrd image so that it will contain the ext3 kernel module. To create this, run the mkinitrd program. For information on using the mkinitrd command, type man mkinitrd. Also, make sure your GRUB configuration loads the initrd.

If you fail to make this change, the system still boots, but the file system is mounted as ext2 instead of ext3.

3.4. ext2 파일 시스템으로 되돌리기

If you wish to revert a partition from ext3 to ext2 for any reason, you must first unmount the partition by logging in as root and typing,

```
umount /dev/mapper/VolGroup00-LogVol02
```

이제 루트로 다음 명령을 입력하여 파일 시스템 유형을 ext2로 변경합니다:

```
tune2fs -O ^has_journal /dev/mapper/VolGroup00-LogVol02
```

루트로 다음과 같은 명령을 입력하여 파티션에 오류가 있는지 확인해 보시기 바랍니다:

```
e2fsck -y /dev/mapper/VolGroup00-LogVol02
```

다음으로 ext2 파일 시스템으로 파티션을 마운트하기 위하여 다음 명령을 입력해 주십시오:

```
mount -t ext2 /dev/mapper/VolGroup00-LogVol02 /mount/point
```

위의 명령에서 /mount/point 부분에 파티션의 마운트 지점을 입력해 주십시오

Next, remove the .journal file at the root level of the partition by changing to the directory where it is mounted and typing:

```
rm -f .journal
```

이제 다시 ext2 파티션이 생성되었습니다.

If you want to permanently change the partition to ext2, remember to update the /etc/fstab file.

The proc File System

The Linux kernel has two primary functions: to control access to physical devices on the computer and to schedule when and how processes interact with these devices. The /proc/ directory — also called the proc file system — contains a hierarchy of special files which represent the current state of the kernel — allowing applications and users to peer into the kernel's view of the system.

Within the /proc/ directory, one can find a wealth of information detailing the system hardware and any processes currently running. In addition, some of the files within the /proc/ directory tree can be manipulated by users and applications to communicate configuration changes to the kernel.

4.1. A Virtual File System

Under Linux, all data are stored as files. Most users are familiar with the two primary types of files: text and binary. But the /proc/ directory contains another type of file called a virtual file. It is for this reason that /proc/ is often referred to as a virtual file system.

These virtual files have unique qualities. Most of them are listed as zero bytes in size and yet when one is viewed, it can contain a large amount of information. In addition, most of the time and date settings on virtual files reflect the current time and date, indicative of the fact they are constantly updated.

Virtual files such as /proc/interrupts, /proc/meminfo, /proc/mounts, and /proc/partitions provide an up-to-the-moment glimpse of the system's hardware. Others, like the /proc/filesystems file and the /proc/sys/ directory provide system configuration information and interfaces.

For organizational purposes, files containing information on a similar topic are grouped into virtual directories and sub-directories. For instance, /proc/ide/ contains information for all physical IDE devices. Likewise, process directories contain information about each running process on the system.

4.1.1. Viewing Virtual Files

By using the cat, more, or less commands on files within the /proc/ directory, users can immediately access enormous amounts of information about the system. For example, to display the type of CPU a computer has, type cat /proc/cpuinfo to receive output similar to the following:

```
processor : 0
vendor_id : AuthenticAMD
cpu family : 5
model   : 9
model name : AMD-K6(tm) 3D+
Processor stepping : 1 cpu
MHz   : 400.919
cache size : 256 KB
fdiv_bug : no
hlt_bug   : no
f00f_bug : no
coma_bug : no
fpu   : yes
fpu_exception : yes
cpuid level : 1
wp   : yes
flags   : fpu vme de pse tsc msr mce cx8 pge mmx syscall 3dnow k6_mtrr
bogomips : 799.53
```

When viewing different virtual files in the /proc/ file system, some of the information is easily understandable while some is not human-readable. This is in part why utilities exist to pull data from virtual files and display it in a useful way. Examples of these utilities include lspci, apm, free, and top.

Note

Some of the virtual files in the /proc/ directory are readable only by the root user.

4.1.2. Changing Virtual Files

As a general rule, most virtual files within the /proc/ directory are read-only. However, some can be used to adjust settings in the kernel. This is especially true for files in the /proc/sys/ subdirectory.

To change the value of a virtual file, use the echo command and a greater than symbol (>) to redirect the new value to the file. For example, to change the hostname on the fly, type:

```
echo www.example.com > /proc/sys/kernel/hostname
```

Other files act as binary or Boolean switches. Typing cat /proc/sys/net/ipv4/ip_forward returns either a 0 or a 1. A 0 indicates that the kernel is not forwarding network packets. Using the echo command to change the value of the ip_forward file to 1 immediately turns packet forwarding on.

Tip

Another command used to alter settings in the /proc/sys/ subdirectory is /sbin/sysctl. For more information on this command, refer to 4.4절. "Using the sysctl Command"

For a listing of some of the kernel configuration files available in the /proc/sys/ subdirectory, refer to 4.3.9절. " /proc/sys/ " .

4.2. Top-level Files within the proc File System

Below is a list of some of the more useful virtual files in the top-level of the /proc/ directory.

Note

In most cases, the content of the files listed in this section are not the same as those installed on your machine. This is because much of the information is specific to the hardware on which Red Hat Enterprise Linux is running for this documentation effort.

4.2.1. /proc/apm

This file provides information about the state of the Advanced Power Management (APM) system and is used by the apm command. If a system with no battery is connected to an AC power source, this virtual file would look similar to the following:

```
1.16 1.2 0x07 0x01 0xff 0x80 -1% -1 ?
```

Running the apm -v command on such a system results in output similar to the following:

```
APM BIOS 1.2 (kernel driver 1.16ac) AC on-line, no system battery
```

For systems which do not use a battery as a power source, apm is able do little more than put the machine in standby mode. The apm command is much more useful on laptops. For example, the following output is from the command cat /proc/apm on a laptop while plugged into a power outlet:

```
1.16 1.2 0x03 0x01 0x03 0x09 100% -1 ?
```

When the same laptop is unplugged from its power source for a few minutes, the content of the apm file changes to something like the following:

```
1.16 1.2 0x03 0x00 0x00 0x01 99% 1792 min
```

The apm -v command now yields more useful data, such as the following:

```
APM BIOS 1.2 (kernel driver 1.16) AC off-line, battery status high: 99% (1 day, 5:52)
```

4.2.2. /proc/buddyinfo

This file is used primarily for diagnosing memory fragmentation issues. Using the buddy algorithm, each column represents the number of pages of a certain order (a certain size) that are available at any given time. For example, for zone DMA (direct memory access), there are 90 of 2^(0*PAGE_SIZE) chunks of memory. Similarly, there are 6 of 2^(1*PAGE_SIZE) chunks, and 2 of 2^(2*PAGE_SIZE) chunks of memory available.

The DMA row references the first 16 MB on a system, the HighMem row references all memory greater than 4 GB on a system, and the Normal row references all memory in between.

The following is an example of the output typical of /proc/buddyinfo:

```
Node 0, zone      DMA     90      6      2      1      1      ...
Node 0, zone   Normal   1650    310      5      0      0      ...
Node 0, zone  HighMem      2      0      0      1      1      ...
```

4.2.3. /proc/cmdline

This file shows the parameters passed to the kernel at the time it is started. A sample /proc/cmdline file looks like the following:

```
ro root=/dev/VolGroup00/LogVol00 rhgb quiet 3
```

This output tells us the following:

ro
: The root device is mounted read-only at boot time. The presence of ro on the kernel boot line overrides any instances of rw.

root=/dev/VolGroup00/LogVol00
: This tells us on which disk device or, in this case, on which logical volume, the root filesystem image is located. With our sample /proc/cmdline output, the root filesystem image is located on the first logical volume (LogVol00) of the first LVM volume group (VolGroup00). On a system not using Logical Volume Management, the root file system might be located on /dev/sda1 or /dev/sda2, meaning on either the first or second partition of the first SCSI or SATA disk drive, depending on whether we have a separate (preceding) boot or swap partition on that drive.

For more information on LVM used in Red Hat Enterprise Linux, refer to http://www.tldp.org/HOWTO/LVM-HOWTO/index.html.

rhgb
: A short lowercase acronym that stands for Red Hat Graphical Boot, providing "rhgb" on the kernel command line signals that graphical booting is supported, assuming that /etc/inittab shows that the default runlevel is set to 5 with a line like this:

```
id:5:initdefault:
```

quiet
: Indicates that all verbose kernel messages except those which are extremely serious should be suppressed at boot time.

4.2.4. /proc/cpuinfo

This virtual file identifies the type of processor used by your system. The following is an example of the output typical of /proc/cpuinfo:

```
processor : 0
vendor_id : GenuineIntel
cpu family : 15
model  : 2
model name : Intel(R) Xeon(TM) CPU 2.40GHz
stepping : 7 cpu
MHz  : 2392.371
cache size : 512 KB
physical id : 0
siblings : 2
runqueue : 0
fdiv_bug : no
hlt_bug  : no
f00f_bug : no
coma_bug : no
fpu  : yes
fpu_exception : yes
cpuid level : 2
wp  : yes
flags  : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca  cmov pat pse36 clflush dts acpi mmx fxsr sse sse2 ss
 ht tm
bogomips : 4771.02
```

- processor — Provides each processor with an identifying number. On systems that have one processor, only a 0 is present.

- cpu family — Authoritatively identifies the type of processor in the system. For an Intel-based system, place the number in front of "86" to determine the value. This is particularly helpful for those attempting to identify the architecture of an older system such as a 586, 486, or 386. Because some RPM packages are compiled for each of these particular architectures, this value also helps users determine which packages to install.
- model name — Displays the common name of the processor, including its project name.
- cpu MHz — Shows the precise speed in megahertz for the processor to the thousandths decimal place.
- cache size — Displays the amount of level 2 memory cache available to the processor.
- siblings — Displays the number of sibling CPUs on the same physical CPU for architectures which use hyper-threading.
- flags — Defines a number of different qualities about the processor, such as the presence of a floating point unit (FPU) and the ability to process MMX instructions.

4.2.5. /proc/crypto

This file lists all installed cryptographic ciphers used by the Linux kernel, including additional details for each. A sample /proc/crypto file looks like the following:

```
name         : sha1
module       : kernel
type         : digest
blocksize    : 64
digestsize   : 20
name         : md5
module       : md5
type         : digest
blocksize    : 64
digestsize   : 16
```

4.2.6. /proc/devices

This file displays the various character and block devices currently configured (not including devices whose modules are not loaded). Below is a sample output from this file:

```
Character devices:
  1 mem
  4 /dev/vc/0
  4 tty
  4 ttyS
  5 /dev/tty
  5 /dev/console
  5 /dev/ptmx
  7 vcs
  10 misc
  13 input
  29 fb
  36 netlink
  128 ptm
  136 pts
  180 usb
```

```
Block devices:
  1 ramdisk
  3 ide0
  9 md
  22 ide1
  253 device-mapper
  254 mdp
```

The output from /proc/devices includes the major number and name of the device, and is broken into two major sections: Character devices and Block devices.

Character devices are similar to block devices, except for two basic differences:

1. Character devices do not require buffering. Block devices have a buffer available, allowing them to order requests before addressing them. This is important for devices designed to store information — such as hard drives — because the ability to order the information before writing it to the device allows it to be placed in a more efficient order.

2. Character devices send data with no preconfigured size. Block devices can send and receive information in blocks of a size configured per device.

For more information about devices refer to the following installed documentation:

```
/usr/share/doc/kernel-doc-<version>/Documentation/devices.txt
```

4.2.7. /proc/dma

This file contains a list of the registered ISA DMA channels in use. A sample /proc/dma files looks like the following:

```
4: cascade
```

4.2.8. /proc/execdomains

This file lists the execution domains currently supported by the Linux kernel, along with the range of personalities they support.

```
0-0    Linux           [kernel]
```

Think of execution domains as the "personality" for an operating system. Because other binary formats, such as Solaris, UnixWare, and FreeBSD, can be used with Linux, programmers can change the way the operating system treats system calls from these binaries by changing the personality of the task. Except for the PER_LINUX execution domain, different personalities can be implemented as dynamically loadable modules.

4.2.9. /proc/fb

This file contains a list of frame buffer devices, with the frame buffer device number and the driver that controls it. Typical output of /proc/fb for systems which contain frame buffer devices looks similar to the following:

```
0 VESA VGA
```

4.2.10. /proc/filesystems

This file displays a list of the file system types currently supported by the kernel. Sample output from a generic /proc/filesystems file looks similar to the following:

```
nodev   sysfs
nodev   rootfs
nodev   bdev
nodev   proc
nodev   sockfs
nodev   binfmt_misc
nodev   usbfs
nodev   usbdevfs
nodev   futexfs
nodev   tmpfs
nodev   pipefs
nodev   eventpollfs
nodev   devpts
 ext2
nodev   ramfs
nodev   hugetlbfs
 iso9660
nodev   mqueue
 ext3
nodev   rpc_pipefs
nodev   autofs
```

The first column signifies whether the file system is mounted on a block device. Those beginning with nodev are not mounted on a device. The second column lists the names of the file systems supported.

The mount command cycles through the file systems listed here when one is not specified as an argument.

4.2.11. /proc/interrupts

This file records the number of interrupts per IRQ on the x86 architecture. A standard /proc/interrupts looks similar to the following:

```
  CPU0
  0:   80448940          XT-PIC  timer
  1:     174412          XT-PIC  keyboard
  2:          0          XT-PIC  cascade
  8:          1          XT-PIC  rtc
 10:     410964          XT-PIC  eth0
 12:      60330          XT-PIC  PS/2 Mouse
 14:    1314121          XT-PIC  ide0
 15:    5195422          XT-PIC  ide1
NMI:          0
ERR:          0
```

For a multi-processor machine, this file may look slightly different:

```
    CPU0       CPU1
  0: 1366814704          0          XT-PIC  timer
  1:        128        340    IO-APIC-edge  keyboard
  2:          0          0          XT-PIC  cascade
  8:          0          1    IO-APIC-edge  rtc
 12:       5323       5793    IO-APIC-edge  PS/2 Mouse
 13:          1          0          XT-PIC  fpu
 16:   11184294   15940594   IO-APIC-level  Intel EtherExpress Pro 10/100 Ethernet
 20:    8450043   11120093   IO-APIC-level  megaraid
 30:      10432      10722   IO-APIC-level  aic7xxx
 31:         23         22   IO-APIC-level  aic7xxx
NMI:          0
ERR:          0
```

The first column refers to the IRQ number. Each CPU in the system has its own column and its own number of interrupts per IRQ. The next column reports the type of interrupt, and the last column contains the name of the device that is located at that IRQ.

Each of the types of interrupts seen in this file, which are architecture-specific, mean something different. For x86 machines, the following values are common:

- XT-PIC — This is the old AT computer interrupts.
- IO-APIC-edge — The voltage signal on this interrupt transitions from low to high, creating an edge, where the interrupt occurs and is only signaled once. This kind of interrupt, as well as the IO-APIC-level interrupt, are only seen on systems with processors from the 586 family and higher.
- IO-APIC-level — Generates interrupts when its voltage signal is high until the signal is low again.

4.2.12. /proc/iomem

This file shows you the current map of the system's memory for each physical device:

```
00000000-0009fbff : System RAM
0009fc00-0009ffff : reserved
000a0000-000bffff : Video RAM area
000c0000-000c7fff : Video ROM
000f0000-000fffff : System ROM
00100000-07ffffff : System RAM
00100000-00291ba8 : Kernel code
00291ba9-002e09cb : Kernel data
e0000000-e3ffffff : VIA Technologies, Inc. VT82C597 [Apollo VP3] e4000000-e7ffffff : PCI Bus #01
e4000000-e4003fff : Matrox Graphics, Inc. MGA G200 AGP
e5000000-e57fffff : Matrox Graphics, Inc. MGA G200 AGP
e8000000-e8ffffff : PCI Bus #01
e8000000-e8ffffff : Matrox Graphics, Inc. MGA G200 AGP
ea000000-ea00007f : Digital Equipment Corporation DECchip 21140 [FasterNet]
ea000000-ea00007f : tulip ffff0000-ffffffff : reserved
```

The first column displays the memory registers used by each of the different types of memory. The second column lists the kind of memory located within those registers and displays which memory registers are used by the kernel within the system RAM or, if the network interface card has multiple Ethernet ports, the memory registers assigned for each port.

4.2.13. /proc/ioports

The output of /proc/ioports provides a list of currently registered port regions used for input or output communication with a device. This file can be quite long. The following is a partial listing:

```
0000-001f : dma1
```

```
0020-003f : pic1
0040-005f : timer
0060-006f : keyboard
0070-007f : rtc
0080-008f : dma page reg
00a0-00bf : pic2
00c0-00df : dma2
00f0-00ff : fpu
0170-0177 : ide1
01f0-01f7 : ide0
02f8-02ff : serial(auto)
0376-0376 : ide1
03c0-03df : vga+
03f6-03f6 : ide0
03f8-03ff : serial(auto)
0cf8-0cff : PCI conf1
d000-dfff : PCI Bus #01
e000-e00f : VIA Technologies, Inc. Bus Master IDE
e000-e007 : ide0
e008-e00f : ide1
e800-e87f : Digital Equipment Corporation DECchip 21140 [FasterNet]
e800-e87f : tulip
```

The first column gives the I/O port address range reserved for the device listed in the second column.

4.2.14. /proc/kcore

This file represents the physical memory of the system and is stored in the core file format. Unlike most /proc/ files, kcore displays a size. This value is given in bytes and is equal to the size of the physical memory (RAM) used plus 4 KB.

The contents of this file are designed to be examined by a debugger, such as gdb, and is not human readable.

Caution

Do not view the /proc/kcore virtual file. The contents of the file scramble text output on the terminal. If this file is accidentally viewed, press Ctrl+C to stop the process and then type reset to bring back the command line prompt.

4.2.15. /proc/kmsg

This file is used to hold messages generated by the kernel. These messages are then picked up by other programs, such as /sbin/klogd or /bin/dmesg.

4.2.16. /proc/loadavg

This file provides a look at the load average in regard to both the CPU and IO over time, as well as additional data used by uptime and other commands. A sample /proc/loadavg file looks similar to the following:

```
0.20 0.18 0.12 1/80 11206
```

The first three columns measure CPU and IO utilization of the last one, five, and 15 minute periods. The fourth column shows the number of currently running processes and the total number of processes. The last column displays the last process ID used.

In addition, load average also refers to the number of processes ready to run (i.e. in the run queue, waiting for a CPU share.

4.2.17. /proc/locks

This file displays the files currently locked by the kernel. The contents of this file contain internal kernel debugging data and can vary tremendously, depending on the use of the system. A sample /proc/locks file for a lightly loaded system looks similar to the following:

```
1: POSIX  ADVISORY  WRITE 3568 fd:00:2531452 0 EOF
2: FLOCK  ADVISORY  WRITE 3517 fd:00:2531448 0 EOF
3: POSIX  ADVISORY  WRITE 3452 fd:00:2531442 0 EOF
4: POSIX  ADVISORY  WRITE 3443 fd:00:2531440 0 EOF
5: POSIX  ADVISORY  WRITE 3326 fd:00:2531430 0 EOF
6: POSIX  ADVISORY  WRITE 3175 fd:00:2531425 0 EOF
7: POSIX  ADVISORY  WRITE 3056 fd:00:2548663 0 EOF
```

Each lock has its own line which starts with a unique number. The second column refers to the class of lock used, with FLOCK signifying the older-style UNIX file locks from a flock system call and POSIX representing the newer POSIX locks from the lockf system call.

The third column can have two values: ADVISORY or MANDATORY. ADVISORY means that the lock does not prevent other people from accessing the data; it only prevents other attempts to lock it. MANDATORY means that no other access to the data is permitted while the lock is held. The fourth column reveals whether the lock is allowing the holder READ or WRITE access to the file. The fifth column shows the ID of the process holding the lock. The sixth column shows the ID of the file being locked, in the format of MAJOR-DEVICE:MINOR-DEVICE:INODE-NUMBER . The seventh and eighth column shows the start and end of the file's locked region.

4.2.18. /proc/mdstat

This file contains the current information for multiple-disk, RAID configurations. If the system does not contain such a configuration, then /proc/mdstat looks similar to the following:

```
Personalities :  read_ahead not set unused devices: <none>
```

This file remains in the same state as seen above unless a software RAID or md device is present. In that case, view /proc/mdstat to find the current status of mdX RAID devices.

The /proc/mdstat file below shows a system with its md0 configured as a RAID 1 device, while it is currently re-syncing the disks:

```
Personalities : [linear] [raid1] read_ahead 1024 sectors
md0: active raid1 sda2[1] sdb2[0] 9940 blocks [2/2] [UU] resync=1% finish=12.3min algorithm 2 [3/3] [UUU]
unused devices: <none>
```

4.2.19. /proc/meminfo

This is one of the more commonly used files in the /proc/ directory, as it reports a large amount of valuable information about the systems RAM usage.

The following sample /proc/meminfo virtual file is from a system with 256 MB of RAM and 512 MB of swap space:

```
MemTotal:         255908 kB
MemFree:           69936 kB
Buffers:           15812 kB
Cached:           115124 kB
SwapCached:            0 kB
Active:            92700 kB
Inactive:          63792 kB
HighTotal:             0 kB
HighFree:              0 kB
LowTotal:         255908 kB
LowFree:           69936 kB
SwapTotal:        524280 kB
SwapFree:         524280 kB
Dirty:                 4 kB
Writeback:             0 kB
Mapped:            42236 kB
Slab:              25912 kB
Committed_AS:     118680 kB
PageTables:         1236 kB
VmallocTotal:    3874808 kB
VmallocUsed:        1416 kB
VmallocChunk:    3872908 kB
HugePages_Total:       0
HugePages_Free:        0
Hugepagesize:       4096 kB
```

Much of the information here is used by the free, top, and ps commands. In fact, the output of the free command is similar in appearance to the contents and structure of /proc/meminfo. But by looking directly at /proc/meminfo, more details are revealed:

- MemTotal — Total amount of physical RAM, in kilobytes.
- MemFree — The amount of physical RAM, in kilobytes, left unused by the system.
- Buffers — The amount of physical RAM, in kilobytes, used for file buffers.
- Cached — The amount of physical RAM, in kilobytes, used as cache memory.
- SwapCached — The amount of swap, in kilobytes, used as cache memory.
- Active — The total amount of buffer or page cache memory, in kilobytes, that is in active use. This is memory that has been recently used and is usually not reclaimed for other purposes.
- Inactive — The total amount of buffer or page cache memory, in kilobytes, that are free and available. This is memory that has not been recently used and can be reclaimed for other purposes.
- HighTotal and HighFree — The total and free amount of memory, in kilobytes, that is not directly mapped into kernel space. The HighTotal value can vary based on the type of kernel used.
- LowTotal and LowFree — The total and free amount of memory, in kilobytes, that is directly mapped into kernel space. The LowTotal value can vary based on the type of kernel used.
- SwapTotal — The total amount of swap available, in kilobytes.
- SwapFree — The total amount of swap free, in kilobytes.

- Dirty — The total amount of memory, in kilobytes, waiting to be written back to the disk.
- Writeback — The total amount of memory, in kilobytes, actively being written back to the disk.
- Mapped — The total amount of memory, in kilobytes, which have been used to map devices, files, or libraries using the mmap command.
- Slab — The total amount of memory, in kilobytes, used by the kernel to cache data structures for its own use.
- Committed_AS — The total amount of memory, in kilobytes, estimated to complete the workload. This value represents the worst case scenario value, and also includes swap memory.
- PageTables — The total amount of memory, in kilobytes, dedicated to the lowest page table level.
- VMallocTotal — The total amount of memory, in kilobytes, of total allocated virtual address space.
- VMallocUsed — The total amount of memory, in kilobytes, of used virtual address space.
- VMallocChunk — The largest contiguous block of memory, in kilobytes, of available virtual address space.
- HugePages_Total — The total number of hugepages for the system. The number is derived by dividing Hugepagesize by the megabytes set aside for hugepages specified in /proc/sys/vm/hugetlb_pool. This statistic only appears on the x86, Itanium, and AMD64 architectures.
- HugePages_Free — The total number of hugepages available for the system. This statistic only appears on the x86, Itanium, and AMD64 architectures.
- Hugepagesize — The size for each hugepages unit in kilobytes. By default, the value is 4096 KB on uniprocessor kernels for 32 bit architectures. For SMP, hugemem kernels, and AMD64, the default is 2048 KB. For Itanium architectures, the default is 262144 KB. This statistic only appears on the x86, Itanium, and AMD64 architectures.

4.2.20. /proc/misc

This file lists miscellaneous drivers registered on the miscellaneous major device, which is device number 10:

```
63 device-mapper 175 agpgart 135 rtc 134 apm_bios
```

The first column is the minor number of each device, while the second column shows the driver in use.

4.2.21. /proc/modules

This file displays a list of all modules loaded into the kernel. Its contents vary based on the configuration and use of your system, but it should be organized in a similar manner to this sample /proc/modules file output:

Note

This example has been reformatted into a readable format. Most of this information can also be viewed via the /sbin/lsmod command.

```
nfs        170109   0 -             Live 0x129b0000
lockd      51593    1 nfs,          Live 0x128b0000
nls_utf8  1729      0 -             Live 0x12830000
vfat       12097    0 -             Live 0x12823000
fat        38881    1 vfat,         Live 0x1287b000
autofs4   20293     2 -             Live 0x1284f000
sunrpc     140453   3 nfs,lockd,  Live 0x12954000
3c59x       33257    0 -             Live 0x12871000
uhci_hcd 28377     0 -             Live 0x12869000
md5         3777      1 -             Live 0x1282c000
ipv6       211845 16 -             Live 0x128de000
ext3       92585     2 -             Live 0x12886000
jbd         65625     1 ext3,         Live 0x12857000
dm_mod     46677     3 -             Live 0x12833000
```

The first column contains the name of the module.

The second column refers to the memory size of the module, in bytes.

The third column lists how many instances of the module are currently loaded. A value of zero represents an unloaded module.

The fourth column states if the module depends upon another module to be present in order to function, and lists those other modules.

The fifth column lists what load state the module is in: Live, Loading, or Unloading are the only possible values.

The sixth column lists the current kernel memory offset for the loaded module. This information can be useful for debugging purposes, or for profiling tools such as oprofile.

4.2.22. /proc/mounts

This file provides a list of all mounts in use by the system:

```
rootfs / rootfs rw 0 0
/proc /proc proc rw,nodiratime 0 0 none
/dev ramfs rw 0 0
/dev/mapper/VolGroup00-LogVol00 / ext3 rw 0 0
none /dev ramfs rw 0 0
/proc /proc proc rw,nodiratime 0 0
/sys /sys sysfs rw 0 0
none /dev/pts devpts rw 0 0
usbdevfs /proc/bus/usb usbdevfs rw 0 0
/dev/hda1 /boot ext3 rw 0 0
none /dev/shm tmpfs rw 0 0
none /proc/sys/fs/binfmt_misc binfmt_misc rw 0 0
sunrpc /var/lib/nfs/rpc_pipefs rpc_pipefs rw 0 0
```

The output found here is similar to the contents of /etc/mtab, except that /proc/mount is more up-to-date.

The first column specifies the device that is mounted, the second column reveals the mount point, and the third column tells the file system type, and the fourth column tells you if it is mounted read-only (ro) or read-write (rw). The fifth and sixth columns are dummy values designed to match the format used in /etc/mtab.

4.2.23. /proc/mtrr

This file refers to the current Memory Type Range Registers (MTRRs) in use with the system. If the system architecture supports MTRRs, then the /proc/mtrr file may look similar to the following:

```
reg00: base=0x00000000 (   0MB), size= 256MB: write-back, count=1
reg01: base=0xe8000000 (3712MB), size=  32MB: write-combining, count=1
```

MTRRs are used with the Intel P6 family of processors (Pentium II and higher) and control processor access to memory ranges. When using a video card on a PCI or AGP bus, a properly configured /proc/mtrr file can increase performance more than 150%.

Most of the time, this value is properly configured by default. More information on manually configuring this file can be found locally at the following location:

```
/usr/share/doc/kernel-doc-<version>/Documentation/mtrr.txt
```

4.2.24. /proc/partitions

This file contains partition block allocation information. A sampling of this file from a basic system looks similar to the following:

```
major minor  #blocks  name
  3     0   19531250 hda
  3     1     104391 hda1
  3     2   19422585 hda2
253     0   22708224 dm-0
253     1     524288 dm-1
```

Most of the information here is of little importance to the user, except for the following columns:

- major — The major number of the device with this partition. The major number in the /proc/partitions, (3), corresponds with the block device ide0, in /proc/devices.
- minor — The minor number of the device with this partition. This serves to separate the partitions into different physical devices and relates to the number at the end of the name of the partition.
- #blocks — Lists the number of physical disk blocks contained in a particular partition.
- name — The name of the partition.

4.2.25. /proc/pci

This file contains a full listing of every PCI device on the system. Depending on the number of PCI devices, /proc/pci can be rather long. A sampling of this file from a basic system looks similar to the following:

```
Bus  0, device 0, function 0: Host bridge: Intel Corporation 440BX/ZX - 82443BX/ZX Host bridge (rev 3). Master Capable.
 Latency=64. Prefetchable 32 bit memory at 0xe4000000 [0xe7ffffff].
Bus  0, device 1, function 0: PCI bridge: Intel Corporation 440BX/ZX - 82443BX/ZX AGP bridge (rev 3).   Master
 Capable. Latency=64. Min Gnt=128.
Bus  0, device 4, function 0: ISA bridge: Intel Corporation 82371AB PIIX4 ISA (rev 2).
Bus  0, device 4, function 1: IDE interface: Intel Corporation 82371AB PIIX4 IDE (rev 1). Master Capable. Latency=32. I/O
 at 0xd800 [0xd80f].
Bus  0, device 4, function 2: USB Controller: Intel Corporation 82371AB PIIX4 USB (rev 1). IRQ 5. Master Capable.
 Latency=32. I/O at 0xd400 [0xd41f].
Bus  0, device 4, function 3: Bridge: Intel Corporation 82371AB PIIX4 ACPI (rev 2). IRQ 9.
Bus  0, device 9, function 0: Ethernet controller: Lite-On Communications Inc LNE100TX (rev 33). IRQ 5. Master Capable.
 Latency=32. I/O at 0xd000 [0xd0ff].
Bus  0, device 12, function  0: VGA compatible controller: S3 Inc. ViRGE/DX or /GX (rev 1). IRQ 11. Master Capable.
 Latency=32. Min Gnt=4.Max Lat=255.
```

This output shows a list of all PCI devices, sorted in the order of bus, device, and function. Beyond providing the name and version of the device, this list also gives detailed IRQ information so an administrator can quickly look for conflicts.

Tip

To get a more readable version of this information, type:

```
lspci -vb
```

4.2.26. /proc/slabinfo

This file gives full information about memory usage on the slab level. Linux kernels greater than version 2.2 use slab pools to manage memory above the page level. Commonly used objects have their own slab pools.

Instead of parsing the highly verbose /proc/slabinfo file manually, the /usr/bin/slabtop program displays kernel slab cache information in real time. This program allows for custom configurations, including column sorting and screen refreshing.

A sample screen shot of /usr/bin/slabtop usually looks like the following example:

```
Active / Total Objects (% used)    : 133629 / 147300 (90.7%)
Active / Total Slabs (% used)      : 11492 / 11493 (100.0%)
Active / Total Caches (% used)      : 77 / 121 (63.6%)
Active / Total Size (% used)       : 41739.83K / 44081.89K (94.7%)
Minimum / Average / Maximum Object : 0.01K / 0.30K / 128.00K
OBJS   ACTIVE  USE      OBJ    SIZE     SLABS OBJ/SLAB CACHE SIZE NAME
44814  43159   96%     0.62K   7469        6      29876K ext3_inode_cache
36900  34614   93%     0.05K    492       75       1968K buffer_head
35213  33124   94%     0.16K   1531       23       6124K dentry_cache
7364    6463   87%    0.27K    526       14       2104K radix_tree_node
2585    1781   68%    0.08K     55       47        220K vm_area_struct
2263    2116   93%    0.12K     73       31        292K size-128
1904    1125   59%    0.03K     16      119         64K size-32
1666     768   46%    0.03K     14      119         56K anon_vma
1512    1482   98%    0.44K    168        9        672K inode_cache
```

```
1464  1040  71%   0.06K   24    61     96K size-64
1320   820  62%   0.19K   66    20    264K filp
678    587  86%   0.02K    3   226     12K dm_io
678    587  86%   0.02K    3   226     12K dm_tio
576    574  99%   0.47K   72     8    288K proc_inode_cache
528    514  97%   0.50K   66     8    264K size-512
492    372  75%   0.09K   12    41     48K bio
465    314  67%   0.25K   31    15    124K size-256
452    331  73%   0.02K    2   226      8K biovec-1
420    420 100%   0.19K   21    20     84K skbuff_head_cache
305    256  83%   0.06K    5    61     20K biovec-4
290      4   1%   0.01K    1   290      4K revoke_table
264    264 100%   4.00K  264     1   1056K size-4096
260    256  98%   0.19K   13    20     52K biovec-16
260    256  98%   0.75K   52     5    208K biovec-64
```

Some of the more commonly used statistics in /proc/slabinfo that are included into /usr/bin/slabtop include:

- OBJS — The total number of objects (memory blocks), including those in use (allocated), and some spares not in use.
- ACTIVE — The number of objects (memory blocks) that are in use (allocated).
- USE — Percentage of total objects that are active. ((ACTIVE/OBJS)(100))
- OBJ SIZE — The size of the objects.
- SLABS — The total number of slabs.
- OBJ/SLAB — The number of objects that fit into a slab.
- CACHE SIZE — The cache size of the slab.
- NAME — The name of the slab.

For more information on the /usr/bin/slabtop program, refer to the slabtop man page.

4.2.27. /proc/stat

This file keeps track of a variety of different statistics about the system since it was last restarted. The contents of /proc/stat, which can be quite long, usually begins like the following example:

```
cpu  259246 7001 60190 34250993 137517 772 0
cpu0 259246 7001 60190 34250993 137517 772 0
intr 354133732 347209999 2272 0 4 4 0 0 3 1 1249247 0 0 80143 0 422626 5169433
ctxt 12547729
btime 1093631447
processes 130523
procs_running 1
procs_blocked 0
preempt 5651840
cpu  209841 1554 21720 118519346 72939 154 27168
cpu0 42536 798 4841 14790880 14778 124 3117
cpu1 24184 569 3875 14794524 30209 29 3130
cpu2 28616 11 2182 14818198 4020 1 3493
cpu3 35350 6 2942 14811519 3045 0 3659
cpu4 18209 135 2263 14820076 12465 0 3373
cpu5 20795 35 1866 14825701 4508 0 3615
cpu6 21607 0 2201 14827053 2325 0 3334
cpu7 18544 0 1550 14831395 1589 0 3447
```

```
intr 15239682 14857833 6 0 6 6 0 5 0 1 0 0 0 29 0 2 0 0 0 0 0 0 0 94982 0 286812
ctxt 4209609
btime 1078711415
processes 21905
procs_running 1
procs_blocked 0
```

Some of the more commonly used statistics include:

- cpu — Measures the number of jiffies (1/100 of a second for x86 systems) that the system has been in user mode, user mode with low priority (nice), system mode, idle task, I/O wait, IRQ (hardirq), and softirq respectively. The IRQ (hardirq) is the direct response to a hardware event. The IRQ takes minimal work for queuing the "heavy" work up for the softirq to execute. The softirq runs at a lower priority than the IRQ and therefore may be interrupted more frequently. The total for all CPUs is given at the top, while each individual CPU is listed below with its own statistics. The following example is a 4-way Intel Pentium Xeon configuration with multi-threading enabled, therefore showing four physical processors and four virtual processors totaling eight processors.

- page — The number of memory pages the system has written in and out to disk.

- swap — The number of swap pages the system has brought in and out.

- intr — The number of interrupts the system has experienced.

- btime — The boot time, measured in the number of seconds since January 1, 1970, otherwise known as the epoch.

4.2.28. /proc/swaps

This file measures swap space and its utilization. For a system with only one swap partition, the output of /proc/swaps may look similar to the following:

```
Filename                          Type        Size      Used    Priority
/dev/mapper/VolGroup00-LogVol01   partition   524280    0       -1
```

While some of this information can be found in other files in the /proc/ directory, /proc/swaps provides a snapshot of every swap file name, the type of swap space, the total size, and the amount of space in use (in kilobytes). The priority column is useful when multiple swap files are in use. The lower the priority, the more likely the swap file is to be used.

4.2.29. /proc/sysrq-trigger

Using the echo command to write to this file, a remote root user can execute most System Request Key commands remotely as if at the local terminal. To echo values to this file, the /proc/sys/kernel/sysrq must be set to a value other than 0. For more information about the System Request Key, refer to 4.3.9.3절. “ /proc/sys/kernel/ ” .

Although it is possible to write to this file, it cannot be read, even by the root user.

4.2.30. /proc/uptime

This file contains information detailing how long the system has been on since its last restart. The output of /proc/uptime is quite minimal:

```
350735.47 234388.90
```

The first number is the total number of seconds the system has been up. The second number is how much of that time the machine has spent idle, in seconds.

4.2.31. /proc/version

This file specifies the version of the Linux kernel and gcc in use, as well as the version of Red Hat Enterprise Linux installed on the system:

```
Linux version 2.6.8-1.523 (user@foo.redhat.com) (gcc version 3.4.1 20040714 \  (Red Hat Enterprise Linux 3.4.1-7)) #1 Mon
 Aug 16 13:27:03 EDT 2004
```

This information is used for a variety of purposes, including the version data presented when a user logs in.

4.3. Directories within /proc/

Common groups of information concerning the kernel are grouped into directories and subdirectories within the /proc/ directory.

4.3.1. Process Directories

Every /proc/ directory contains a number of directories with numerical names. A listing of them may be similar to the following:

```
dr-xr-xr-x    3 root      root            0 Feb 13 01:28 1
dr-xr-xr-x    3 root      root            0 Feb 13 01:28 1010
dr-xr-xr-x    3 xfs       xfs             0 Feb 13 01:28 1087
dr-xr-xr-x    3 daemon    daemon          0 Feb 13 01:28 1123
dr-xr-xr-x    3 root      root            0 Feb 13 01:28 11307
dr-xr-xr-x    3 apache    apache          0 Feb 13 01:28 13660
dr-xr-xr-x    3 rpc       rpc             0 Feb 13 01:28 637
dr-xr-xr-x    3 rpcuser   rpcuser         0 Feb 13 01:28 666
```

These directories are called process directories, as they are named after a program's process ID and contain information specific to that process. The owner and group of each process directory is set to the user running the process. When the process is terminated, its /proc/ process directory vanishes.

Each process directory contains the following files:

- cmdline — Contains the command issued when starting the process.
- cwd — A symbolic link to the current working directory for the process.
- environ — A list of the environment variables for the process. The environment variable is given in all upper-case characters, and the value is in lower-case characters.
- exe — A symbolic link to the executable of this process.
- fd — A directory containing all of the file descriptors for a particular process. These are given in numbered links:

```
total 0
lrwx------    1 root     root          64 May  8 11:31 0 -> /dev/null
lrwx------    1 root     root          64 May  8 11:31 1 -> /dev/null
lrwx------    1 root     root          64 May  8 11:31 2 -> /dev/null
lrwx------    1 root     root          64 May  8 11:31 3 -> /dev/ptmx
lrwx------    1 root     root          64 May  8 11:31 4 -> socket:[7774817]
lrwx------    1 root     root          64 May  8 11:31 5 -> /dev/ptmx
lrwx------    1 root     root          64 May  8 11:31 6 -> socket:[7774829]
lrwx------    1 root     root          64 May  8 11:31 7 -> /dev/ptmx
```

- maps — A list of memory maps to the various executables and library files associated with this process. This file can be rather long, depending upon the complexity of the process, but sample output from the sshd process begins like the following:

```
08048000-08086000 r-xp 00000000 03:03 391479     /usr/sbin/sshd
08086000-08088000 rw-p 0003e000 03:03 391479 /usr/sbin/sshd
08088000-08095000 rwxp 00000000 00:00 0
40000000-40013000 r-xp 0000000 03:03 293205 /lib/ld-2.2.5.so
40013000-40014000 rw-p 00013000 03:03 293205 /lib/ld-2.2.5.so
40031000-40038000 r-xp 00000000 03:03 293282 /lib/libpam.so.0.75
40038000-40039000 rw-p 00006000 03:03 293282 /lib/libpam.so.0.75
40039000-4003a000 rw-p 00000000 00:00 0
4003a000-4003c000 r-xp 00000000 03:03 293218 /lib/libdl-2.2.5.so
4003c000-4003d000 rw-p 00001000 03:03 293218 /lib/libdl-2.2.5.so
```

- mem — The memory held by the process. This file cannot be read by the user.

- root — A link to the root directory of the process.

- stat — The status of the process.

- statm — The status of the memory in use by the process. Below is a sample /proc/statm file:

```
263 210 210 5 0 205 0
```

The seven columns relate to different memory statistics for the process. From left to right, they report the following aspects of the memory used:

1. Total program size, in kilobytes.

2. Size of memory portions, in kilobytes.

3. Number of pages that are shared.

4. Number of pages that are code.

5. Number of pages of data/stack.

6. Number of library pages.

7. Number of dirty pages.

- status — The status of the process in a more readable form than stat or statm. Sample output for sshd looks similar to the following:

```
Name: sshd
State: S (sleeping)
Tgid: 797
```

```
Pid: 797
PPid: 1
TracerPid: 0
Uid: 0 0 0 0
Gid: 0 0 0 0
FDSize: 32
Groups:
VmSize:       3072 kB
VmLck:           0 kB
VmRSS:         840 kB
VmData:        104 kB
VmStk:          12 kB
VmExe:         300 kB
VmLib:        2528 kB
SigPnd: 0000000000000000
SigBlk: 0000000000000000
SigIgn: 8000000000001000
SigCgt: 0000000000014005
CapInh: 0000000000000000
CapPrm: 00000000fffffeff
CapEff: 00000000fffffeff
```

The information in this output includes the process name and ID, the state (such as S (sleeping) or R (running)), user/group ID running the process, and detailed data regarding memory usage.

4.3.1.1. /proc/self/

The /proc/self/ directory is a link to the currently running process. This allows a process to look at itself without having to know its process ID.

Within a shell environment, a listing of the /proc/self/ directory produces the same contents as listing the process directory for that process.

4.3.2. /proc/bus/

This directory contains information specific to the various buses available on the system. For example, on a standard system containing PCI and USB buses, current data on each of these buses is available within a subdirectory within /proc/bus/ by the same name, such as /proc/bus/pci/.

The subdirectories and files available within /proc/bus/ vary depending on the devices connected to the system. However, each bus type has at least one directory. Within these bus directories are normally at least one subdirectory with a numerical name, such as 001, which contain binary files.

For example, the /proc/bus/usb/ subdirectory contains files that track the various devices on any USB buses, as well as the drivers required for them. The following is a sample listing of a /proc/bus/usb/ directory:

```
total 0 dr-xr-xr-x    1 root     root            0 May  3 16:25 001
-r--r--r--     1 root     root            0 May  3 16:25 devices
-r--r--r--     1 root     root            0 May  3 16:25 drivers
```

The /proc/bus/usb/001/ directory contains all devices on the first USB bus and the devices file identifies the USB root hub on the motherboard.

The following is a example of a /proc/bus/usb/devices file:

```
T:  Bus=01 Lev=00 Prnt=00 Port=00 Cnt=00 Dev#=  1 Spd=12  MxCh= 2
```

```
B:  Alloc=  0/900 us ( 0%), #Int=  0, #Iso=  0
D:  Ver= 1.00 Cls=09(hub  ) Sub=00 Prot=00 MxPS= 8 #Cfgs=  1
P:  Vendor=0000 ProdID=0000 Rev= 0.00
S:  Product=USB UHCI Root Hub
S:  SerialNumber=d400
C:* #Ifs= 1 Cfg#= 1 Atr=40 MxPwr=  0mA
I:  If#= 0 Alt= 0 #EPs= 1 Cls=09(hub  ) Sub=00 Prot=00 Driver=hub
E:  Ad=81(I) Atr=03(Int.) MxPS=   8 Ivl=255ms
```

4.3.3. /proc/driver/

This directory contains information for specific drivers in use by the kernel.

A common file found here is rtc which provides output from the driver for the system's Real Time Clock (RTC), the device that keeps the time while the system is switched off. Sample output from /proc/driver/rtc looks like the following:

```
rtc_time        : 16:21:00
rtc_date        : 2004-08-31
rtc_epoch       : 1900
alarm           : 21:16:27
DST_enable      : no
BCD             : yes
24hr            : yes
square_wave     : no
alarm_IRQ       : no
update_IRQ      : no
periodic_IRQ    : no
periodic_freq   : 1024
batt_status     : okay
```

For more information about the RTC, refer to the following installed documentation:

/usr/share/doc/kernel-doc-<version>/Documentation/rtc.txt.

4.3.4. /proc/fs

This directory shows which file systems are exported. If running an NFS server, typing cat /proc/fs/nfsd/exports displays the file systems being shared and the permissions granted for those file systems. For more on file system sharing with NFS, refer to 20장. 네트워크 파일 시스템 (NFS).

4.3.5. /proc/ide/

This directory contains information about IDE devices on the system. Each IDE channel is represented as a separate directory, such as /proc/ide/ide0 and /proc/ide/ide1. In addition, a drivers file is available, providing the version number of the various drivers used on the IDE channels:

```
ide-floppy version 0.99.
newide ide-cdrom version 4.61
ide-disk version 1.18
```

Many chipsets also provide a file in this directory with additional data concerning the drives connected through the channels. For example, a generic Intel PIIX4 Ultra 33 chipset produces the /proc/ide/piix file which reveals whether DMA or UDMA is enabled for the devices on the IDE channels:

```
Intel PIIX4 Ultra 33 Chipset.
------------- Primary Channel ---------------- Secondary Channel -------------
  enabled                                  enabled

------------- drive0 --------- drive1 -------- drive0 ---------- drive1 ------
DMA enabled:      yes                 no                  yes                 no
UDMA enabled:     yes                 no                  no                  no
UDMA enabled:     2                   X                   X                   X
UDMA DMA PIO
```

Navigating into the directory for an IDE channel, such as ide0, provides additional information. The channel file provides the channel number, while the model identifies the bus type for the channel (such as pci).

4.3.5.1. Device Directories

Within each IDE channel directory is a device directory. The name of the device directory corresponds to the drive letter in the /dev/ directory. For instance, the first IDE drive on ide0 would be hda.

Note

There is a symbolic link to each of these device directories in the /proc/ide/ directory.

Each device directory contains a collection of information and statistics. The contents of these directories vary according to the type of device connected. Some of the more useful files common to many devices include:

- cache — The device cache.
- capacity — The capacity of the device, in 512 byte blocks.
- driver — The driver and version used to control the device.
- geometry — The physical and logical geometry of the device.
- media — The type of device, such as a disk.
- model — The model name or number of the device.
- settings — A collection of current device parameters. This file usually contains quite a bit of useful, technical information. A sample settings file for a standard IDE hard disk looks similar to the following:

```
name                 value           min            max            mode
----                 -----           ---            ---            ----
acoustic             0               0              254            rw
address              0               0              2              rw
bios_cyl             38752           0              65535          rw
bios_head            16              0              255            rw
bios_sect            63              0              63             rw
bswap                0               0              1              r
current_speed        68              0              70             rw
failures             0               0              65535          rw
```

```
init_speed       68            0          70          rw
io_32bit          0            0           3          rw
keepsettings      0            0           1          rw
lun               0            0           7          rw
max_failures      1            0       65535          rw
multcount        16            0          16          rw
nice1             1            0           1          rw
nowerr            0            0           1          rw
number            0            0           3          rw
pio_mode          write-only   0         255          w
unmaskirq         0            0           1          rw
using_dma         1            0           1          rw
wcache            1            0           1          rw
```

4.3.6. /proc/irq/

This directory is used to set IRQ to CPU affinity, which allows the system to connect a particular IRQ to only one CPU. Alternatively, it can exclude a CPU from handling any IRQs.

Each IRQ has its own directory, allowing for the individual configuration of each IRQ. The /proc/irq/prof_cpu_mask file is a bitmask that contains the default values for the smp_affinity file in the IRQ directory. The values in smp_affinity specify which CPUs handle that particular IRQ.

For more information about the /proc/irq/ directory, refer to the following installed documentation:

```
/usr/share/doc/kernel-doc-<version>/Documentation/filesystems/proc.txt
```

4.3.7. /proc/net/

This directory provides a comprehensive look at various networking parameters and statistics. Each directory and virtual file within this directory describes aspects of the system's network configuration. Below is a partial list of the /proc/net/ directory:

- arp — Lists the kernel's ARP table. This file is particularly useful for connecting a hardware address to an IP address on a system.
- atm/ directory — The files within this directory contain Asynchronous Transfer Mode (ATM) settings and statistics. This directory is primarily used with ATM networking and ADSL cards.
- dev — Lists the various network devices configured on the system, complete with transmit and receive statistics. This file displays the number of bytes each interface has sent and received, the number of packets inbound and outbound, the number of errors seen, the number of packets dropped, and more.
- dev_mcast — Lists Layer2 multicast groups on which each device is listening.
- igmp — Lists the IP multicast addresses which this system joined.
- ip_conntrack — Lists tracked network connections for machines that are forwarding IP connections.
- ip_tables_names — Lists the types of iptables in use. This file is only present if iptables is active on the system and contains one or more of the following values: filter, mangle, or nat.
- ip_mr_cache — Lists the multicast routing cache.
- ip_mr_vif — Lists multicast virtual interfaces.

- netstat — Contains a broad yet detailed collection of networking statistics, including TCP timeouts, SYN cookies sent and received, and much more.
- psched — Lists global packet scheduler parameters.
- raw — Lists raw device statistics.
- route — Lists the kernel's routing table.
- rt_cache — Contains the current routing cache.
- snmp — List of Simple Network Management Protocol (SNMP) data for various networking protocols in use.
- sockstat — Provides socket statistics.
- tcp — Contains detailed TCP socket information.
- tr_rif — Lists the token ring RIF routing table.
- udp — Contains detailed UDP socket information.
- unix — Lists UNIX domain sockets currently in use.
- wireless — Lists wireless interface data.

4.3.8. /proc/scsi/

This directory is analogous to the /proc/ide/ directory, but it is for connected SCSI devices.

The primary file in this directory is /proc/scsi/scsi, which contains a list of every recognized SCSI device. From this listing, the type of device, as well as the model name, vendor, SCSI channel and ID data is available.

For example, if a system contains a SCSI CD-ROM, a tape drive, a hard drive, and a RAID controller, this file looks similar to the following:

```
Attached devices:
Host: scsi1
Channel: 00
Id: 05
Lun: 00
Vendor: NEC
Model: CD-ROM DRIVE:466
Rev: 1.06
Type:    CD-ROM
ANSI SCSI revision: 02
Host: scsi1
Channel: 00
Id: 06
Lun: 00
Vendor: ARCHIVE
Model: Python 04106-XXX
Rev: 7350
Type:    Sequential-Access
ANSI SCSI revision: 02
Host: scsi2
Channel: 00
Id: 06
Lun: 00
Vendor: DELL
```

```
Model: 1x6 U2W SCSI BP
Rev: 5.35
Type:    Processor
ANSI SCSI revision: 02
Host: scsi2
Channel: 02
Id: 00
Lun: 00
Vendor: MegaRAID
Model: LD0 RAID5 34556R
Rev: 1.01
Type:    Direct-Access
ANSI SCSI revision: 02
```

Each SCSI driver used by the system has its own directory within /proc/scsi/, which contains files specific to each SCSI controller using that driver. From the previous example, aic7xxx/ and megaraid/ directories are present, since two drivers are in use. The files in each of the directories typically contain an I/O address range, IRQ information, and statistics for the SCSI controller using that driver. Each controller can report a different type and amount of information. The Adaptec AIC-7880 Ultra SCSI host adapter's file in this example system produces the following output:

```
Adaptec AIC7xxx driver version: 5.1.20/3.2.4
Compile Options:
TCQ Enabled By Default : Disabled
AIC7XXX_PROC_STATS      : Enabled
AIC7XXX_RESET_DELAY     : 5
Adapter Configuration:
SCSI Adapter: Adaptec AIC-7880 Ultra SCSI host adapter
Ultra Narrow Controller      PCI MMAPed
I/O Base: 0xfcffe000
Adapter SEEPROM Config: SEEPROM found and used.
Adaptec SCSI BIOS: Enabled
IRQ: 30
SCBs: Active 0, Max Active 1, Allocated 15, HW 16, Page 255
Interrupts: 33726
BIOS Control Word: 0x18a6
Adapter Control Word: 0x1c5f
Extended Translation: Enabled
Disconnect Enable Flags: 0x00ff
Ultra Enable Flags: 0x0020
Tag Queue Enable Flags: 0x0000
Ordered Queue Tag Flags: 0x0000
Default Tag Queue Depth: 8
Tagged Queue By Device array for aic7xxx
host instance 1:        {255,255,255,255,255,255,255,255,255,255,255,255,255,255,255,255}
Actual queue depth per device for aic7xxx host instance 1:        {1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1}
Statistics:

(scsi1:0:5:0) Device using Narrow/Sync transfers at 20.0 MByte/sec, offset 15
Transinfo settings: current(12/15/0/0), goal(12/15/0/0), user(12/15/0/0)
Total transfers 0 (0 reads and 0 writes)
  < 2K      2K+     4K+     8K+    16K+    32K+    64K+   128K+
Reads:        0       0       0       0       0       0       0       0
Writes:       0       0       0       0       0       0       0       0

(scsi1:0:6:0) Device using Narrow/Sync transfers at 10.0 MByte/sec, offset 15
Transinfo settings: current(25/15/0/0), goal(12/15/0/0), user(12/15/0/0)
Total transfers 132 (0 reads and 132 writes)
  < 2K      2K+     4K+     8K+    16K+    32K+    64K+   128K+
Reads:        0       0       0       0       0       0       0       0
Writes:       0       0       0       1     131       0       0       0
```

This output reveals the transfer speed to the SCSI devices connected to the controller based on channel ID, as well as detailed statistics concerning the amount and sizes of files read or written by that device. For example, this controller is communicating with the CD-ROM at 20 megabytes per second, while the tape drive is only communicating at 10 megabytes per second.

4.3.9. /proc/sys/

The /proc/sys/ directory is different from others in /proc/ because it not only provides information about the system but also allows the system administrator to immediately enable and disable kernel features.

Caution

Use caution when changing settings on a production system using the various files in the /proc/sys/ directory. Changing the wrong setting may render the kernel unstable, requiring a system reboot.

For this reason, be sure the options are valid for that file before attempting to change any value in /proc/sys/.

A good way to determine if a particular file can be configured, or if it is only designed to provide information, is to list it with the -l option at the shell prompt. If the file is writable, it may be used to configure the kernel. For example, a partial listing of /proc/sys/fs looks like the following:

```
-r--r--r--    1 root    root         0 May 10 16:14 dentry-state
-rw-r--r--    1 root    root         0 May 10 16:14 dir-notify-enable
-r--r--r--    1 root    root         0 May 10 16:14 dquot-nr
-rw-r--r--    1 root    root         0 May 10 16:14 file-max
-r--r--r--    1 root    root         0 May 10 16:14 file-nr
```

In this listing, the files dir-notify-enable and file-max can be written to and, therefore, can be used to configure the kernel. The other files only provide feedback on current settings.

Changing a value within a /proc/sys/ file is done by echoing the new value into the file. For example, to enable the System Request Key on a running kernel, type the command:

```
echo 1 > /proc/sys/kernel/sysrq
```

This changes the value for sysrq from 0 (off) to 1 (on).

A few /proc/sys/ configuration files contain more than one value. To correctly send new values to them, place a space character between each value passed with the echo command, such as is done in this example:

```
echo 4 2 45 > /proc/sys/kernel/acct
```

Note

Any configuration changes made using the echo command disappear when the system is restarted. To make configuration changes take effect after the system is rebooted, refer to 4.4절. “Using the sysctl Command” .

The /proc/sys/ directory contains several subdirectories controlling different aspects of a running kernel.

4.3.9.1. /proc/sys/dev/

This directory provides parameters for particular devices on the system. Most systems have at least two directories, cdrom/ and raid/. Customized kernels can have other directories, such as parport/, which provides the ability to share one parallel port between multiple device drivers.

The cdrom/ directory contains a file called info, which reveals a number of important CD-ROM parameters:

```
CD-ROM information, Id: cdrom.c 3.20 2003/12/17
drive name:              hdc
drive speed:             48
drive # of slots:        1
Can close tray:          1
Can open tray:           1
Can lock tray:           1
Can change speed:        1
Can select disk:         0
Can read multisession:   1
Can read MCN:            1
Reports media changed:   1
Can play audio:          1
Can write CD-R:          0
Can write CD-RW:         0
Can read DVD:            0
Can write DVD-R:         0
Can write DVD-RAM:       0
Can read MRW:            0
Can write MRW:           0
Can write RAM:           0
```

This file can be quickly scanned to discover the qualities of an unknown CD-ROM. If multiple CD-ROMs are available on a system, each device is given its own column of information.

Various files in /proc/sys/dev/cdrom, such as autoclose and checkmedia, can be used to control the system's CD-ROM. Use the echo command to enable or disable these features.

If RAID support is compiled into the kernel, a /proc/sys/dev/raid/ directory becomes available with at least two files in it: speed_limit_min and speed_limit_max. These settings determine the acceleration of RAID devices for I/O intensive tasks, such as resyncing the disks.

4.3.9.2. /proc/sys/fs/

This directory contains an array of options and information concerning various aspects of the file system, including quota, file handle, inode, and dentry information.

The binfmt_misc/ directory is used to provide kernel support for miscellaneous binary formats.

The important files in /proc/sys/fs/ include:

- dentry-state — Provides the status of the directory cache. The file looks similar to the following:

  ```
  57411 52939 45 0 0 0
  ```

  The first number reveals the total number of directory cache entries, while the second number displays the number of unused entries. The third number tells the number of seconds between when a directory has been freed and when it can be reclaimed, and the fourth measures the pages currently requested by the system. The last two numbers are not used and display only zeros.

- dquot-nr — Lists the maximum number of cached disk quota entries.

- file-max — Lists the maximum number of file handles that the kernel allocates. Raising the value in this file can resolve errors caused by a lack of available file handles.

- file-nr — Lists the number of allocated file handles, used file handles, and the maximum number of file handles.

- overflowgid and overflowuid — Defines the fixed group ID and user ID, respectively, for use with file systems that only support 16-bit group and user IDs.

- super-max — Controls the maximum number of superblocks available.

- super-nr — Displays the current number of superblocks in use.

4.3.9.3. /proc/sys/kernel/

This directory contains a variety of different configuration files that directly affect the operation of the kernel. Some of the most important files include:

- acct — Controls the suspension of process accounting based on the percentage of free space available on the file system containing the log. By default, the file looks like the following:

  ```
  4 2 30
  ```

  The first value dictates the percentage of free space required for logging to resume, while the second value sets the threshold percentage of free space when logging is suspended. The third value sets the interval, in seconds, that the kernel polls the file system to see if logging should be suspended or resumed.

- cap-bound — Controls the capability bounding settings, which provides a list of capabilities for any process on the system. If a capability is not listed here, then no process, no matter how privileged, can do it. The idea is to make the system more secure by ensuring that certain things cannot happen, at least beyond a certain point in the boot process.

  For a valid list of values for this virtual file, refer to the following installed documentation:

  /lib/modules/<kernel-version>/build/include/linux/capability.h.

- ctrl-alt-del — Controls whether Ctrl+Alt+Delete gracefully restarts the computer using init (0) or forces an immediate reboot without syncing the dirty buffers to disk (1).
- domainname — Configures the system domain name, such as example.com.
- exec-shield — Configures the Exec Shield feature of the kernel. Exec Shield provides protection against certain types of buffer overflow attacks.

  There are two possible values for this virtual file:

  - 0 — Disables Exec Shield.
  - 1 — Enables Exec Shield. This is the default value.

  Important

  If a system is running security-sensitive applications that were started while Exec Shield was disabled, these applications must be restarted when Exec Shield is enabled in order for Exec Shield to take effect.

- exec-shield-randomize — Enables location randomization of various items in memory. This helps deter potential attackers from locating programs and daemons in memory. Each time a program or daemon starts, it is put into a different memory location each time, never in a static or absolute memory address.

  There are two possible values for this virtual file:

  - 0 — Disables randomization of Exec Shield. This may be useful for application debugging purposes.
  - 1 — Enables randomization of Exec Shield. This is the default value. Note: The exec-shield file must also be set to 1 for exec-shield-randomize to be effective.

- hostname — Configures the system hostname, such as www.example.com.
- hotplug — Configures the utility to be used when a configuration change is detected by the system. This is primarily used with USB and Cardbus PCI. The default value of /sbin/hotplug should not be changed unless testing a new program to fulfill this role.
- modprobe — Sets the location of the program used to load kernel modules. The default value is /sbin/modprobe which means kmod calls it to load the module when a kernel thread calls kmod.
- msgmax — Sets the maximum size of any message sent from one process to another and is set to 8192 bytes by default. Be careful when raising this value, as queued messages between processes are stored in non-swappable kernel memory. Any increase in msgmax would increase RAM requirements for the system.
- msgmnb — Sets the maximum number of bytes in a single message queue. The default is 16384.
- msgmni — Sets the maximum number of message queue identifiers. The default is 16.
- osrelease — Lists the Linux kernel release number. This file can only be altered by changing the kernel source and recompiling.

- ostype — Displays the type of operating system. By default, this file is set to Linux, and this value can only be changed by changing the kernel source and recompiling.
- overflowgid and overflowuid — Defines the fixed group ID and user ID, respectively, for use with system calls on architectures that only support 16-bit group and user IDs.
- panic — Defines the number of seconds the kernel postpones rebooting when the system experiences a kernel panic. By default, the value is set to 0, which disables automatic rebooting after a panic.
- printk — This file controls a variety of settings related to printing or logging error messages. Each error message reported by the kernel has a loglevel associated with it that defines the importance of the message. The loglevel values break down in this order:

  - 0 — Kernel emergency. The system is unusable.
  - 1 — Kernel alert. Action must be taken immediately.
  - 2 — Condition of the kernel is considered critical.
  - 3 — General kernel error condition.
  - 4 — General kernel warning condition.
  - 5 — Kernel notice of a normal but significant condition.
  - 6 — Kernel informational message.
  - 7 — Kernel debug-level messages.

  Four values are found in the printk file:

  ```
  6       4       1       7
  ```

  Each of these values defines a different rule for dealing with error messages. The first value, called the console loglevel, defines the lowest priority of messages printed to the console. (Note that, the lower the priority, the higher the loglevel number.) The second value sets the default loglevel for messages without an explicit loglevel attached to them. The third value sets the lowest possible loglevel configuration for the console loglevel. The last value sets the default value for the console loglevel.

- random/ directory — Lists a number of values related to generating random numbers for the kernel.
- rtsig-max — Configures the maximum number of POSIX real-time signals that the system may have queued at any one time. The default value is 1024.
- rtsig-nr — Lists the current number of POSIX real-time signals queued by the kernel.
- sem — Configures semaphore settings within the kernel. A semaphore is a System V IPC object that is used to control utilization of a particular process.
- shmall— Sets the total amount of shared memory pages that can be used at one time, system-wide. By default, this value is 2097152.
- shmmax — Sets the largest shared memory segment size allowed by the kernel. By default, this value is 33554432. However, the kernel supports much larger values than this.

- shmmni — Sets the maximum number of shared memory segments for the whole system. By default, this value is 4096.
- sysrq — Activates the System Request Key, if this value is set to anything other than zero (0), the default.

  The System Request Key allows immediate input to the kernel through simple key combinations. For example, the System Request Key can be used to immediately shut down or restart a system, sync all mounted file systems, or dump important information to the console. To initiate a System Request Key, type Alt+SysRq+ <system request code> . Replace <system request code> with one of the following system request codes:

  - r — Disables raw mode for the keyboard and sets it to XLATE (a limited keyboard mode which does not recognize modifiers such as Alt, Ctrl, or Shift for all keys).
  - k — Kills all processes active in a virtual console. Also called Secure Access Key (SAK), it is often used to verify that the login prompt is spawned from init and not a Trojan copy designed to capture usernames and passwords.
  - b — Reboots the kernel without first unmounting file systems or syncing disks attached to the system.
  - c — Crashes the system without first unmounting file systems or syncing disks attached to the system.
  - o — Shuts off the system.
  - s — Attempts to sync disks attached to the system.
  - u — Attempts to unmount and remount all file systems as read-only.
  - p — Outputs all flags and registers to the console.
  - t — Outputs a list of processes to the console.
  - m — Outputs memory statistics to the console.
  - 0 through 9 — Sets the log level for the console.
  - e — Kills all processes except init using SIGTERM.
  - i — Kills all processes except init using SIGKILL.
  - l — Kills all processes using SIGKILL (including init). The system is unusable after issuing this System Request Key code.
  - h — Displays help text.

  This feature is most beneficial when using a development kernel or when experiencing system freezes.

Caution

The System Request Key feature is considered a security risk because an unattended console provides an attacker with access to the system. For this reason, it is turned off by default.

Refer to /usr/share/doc/kernel-doc-<version>/Documentation/sysrq.txt for more information about the System Request Key.

- sysrq-key — Defines the key code for the System Request Key (84 is the default).
- sysrq-sticky — Defines whether the System Request Key is a chorded key combination. The accepted values are as follows:
  - 0 — Alt+SysRq and the system request code must be pressed simultaneously. This is the default value.
  - 1 — Alt+SysRq must be pressed simultaneously, but the system request code can be pressed anytime before the number of seconds specified in /proc/sys/kernel/sysrq-timer elapses.
- sysrq-timer — Specifies the number of seconds allowed to pass before the system request code must be pressed. The default value is 10.
- tainted — Indicates whether a non-GPL module is loaded.
  - 0 — No non-GPL modules are loaded.
  - 1 — At least one module without a GPL license (including modules with no license) is loaded.
  - 2 — At least one module was force-loaded with the command insmod -f.
- threads-max — Sets the maximum number of threads to be used by the kernel, with a default value of 2048.
- version — Displays the date and time the kernel was last compiled. The first field in this file, such as #3, relates to the number of times a kernel was built from the source base.

4.3.9.4. /proc/sys/net/

This directory contains subdirectories concerning various networking topics. Various configurations at the time of kernel compilation make different directories available here, such as ethernet/, ipv4/, ipx/, and ipv6/. By altering the files within these directories, system administrators are able to adjust the network configuration on a running system.

Given the wide variety of possible networking options available with Linux, only the most common /proc/sys/net/ directories are discussed.

The /proc/sys/net/core/ directory contains a variety of settings that control the interaction between the kernel and networking layers. The most important of these files are:

- message_burst — Sets the amount of time in tenths of a second required to write a new warning message. This setting is used to mitigate Denial of Service (DoS) attacks. The default setting is 50.

- message_cost — Sets a cost on every warning message. The higher the value of this file (default of 5), the more likely the warning message is ignored. This setting is used to mitigate DoS attacks.

  The idea of a DoS attack is to bombard the targeted system with requests that generate errors and fill up disk partitions with log files or require all of the system's resources to handle the error logging. The settings in message_burst and message_cost are designed to be modified based on the system's acceptable risk versus the need for comprehensive logging.

- netdev_max_backlog — Sets the maximum number of packets allowed to queue when a particular interface receives packets faster than the kernel can process them. The default value for this file is 300.

- optmem_max — Configures the maximum ancillary buffer size allowed per socket.

- rmem_default — Sets the receive socket buffer default size in bytes.

- rmem_max — Sets the receive socket buffer maximum size in bytes.

- wmem_default — Sets the send socket buffer default size in bytes.

- wmem_max — Sets the send socket buffer maximum size in bytes.

The /proc/sys/net/ipv4/ directory contains additional networking settings. Many of these settings, used in conjunction with one another, are useful in preventing attacks on the system or when using the system to act as a router.

Caution

An erroneous change to these files may affect remote connectivity to the system.

The following is a list of some of the more important files within the /proc/sys/net/ipv4/ directory:

- icmp_destunreach_rate, icmp_echoreply_rate, icmp_paramprob_rate, and icmp_timeexeed_rate — Set the maximum ICMP send packet rate, in 1/100 of a second, to hosts under certain conditions. A setting of 0 removes any delay and is not a good idea.

- icmp_echo_ignore_all and icmp_echo_ignore_broadcasts — Allows the kernel to ignore ICMP ECHO packets from every host or only those originating from broadcast and multicast addresses, respectively. A value of 0 allows the kernel to respond, while a value of 1 ignores the packets.

- ip_default_ttl — Sets the default Time To Live (TTL), which limits the number of hops a packet may make before reaching its destination. Increasing this value can diminish system performance.

- ip_forward — Permits interfaces on the system to forward packets to one other. By default, this file is set to 0. Setting this file to 1 enables network packet forwarding.

- ip_local_port_range — Specifies the range of ports to be used by TCP or UDP when a local port is needed. The first number is the lowest port to be used and the second number specifies the highest port. Any systems that expect to require more ports than the default 1024 to 4999 should use a range from 32768 to 61000.

- tcp_syn_retries — Provides a limit on the number of times the system re-transmits a SYN packet when attempting to make a connection.

- tcp_retries1 — Sets the number of permitted re-transmissions attempting to answer an incoming connection. Default of 3.
- tcp_retries2 — Sets the number of permitted re-transmissions of TCP packets. Default of 15.

The file called

```
/usr/share/doc/kernel-doc-<version>/Documentation/networking/ ip-sysctl.txt
```

contains a complete list of files and options available in the /proc/sys/net/ipv4/ directory.

A number of other directories exist within the /proc/sys/net/ipv4/ directory and each covers a different aspect of the network stack. The /proc/sys/net/ipv4/conf/ directory allows each system interface to be configured in different ways, including the use of default settings for unconfigured devices (in the /proc/sys/net/ipv4/conf/default/ subdirectory) and settings that override all special configurations (in the /proc/sys/net/ipv4/conf/all/ subdirectory).

The /proc/sys/net/ipv4/neigh/ directory contains settings for communicating with a host directly connected to the system (called a network neighbor) and also contains different settings for systems more than one hop away.

Routing over IPV4 also has its own directory, /proc/sys/net/ipv4/route/. Unlike conf/ and neigh/, the /proc/sys/net/ipv4/route/ directory contains specifications that apply to routing with any interfaces on the system. Many of these settings, such as max_size, max_delay, and min_delay, relate to controlling the size of the routing cache. To clear the routing cache, write any value to the flush file.

Additional information about these directories and the possible values for their configuration files can be found in:

```
/usr/share/doc/kernel-doc-<version>/Documentation/filesystems/proc.txt
```

4.3.9.5. /proc/sys/vm/

This directory facilitates the configuration of the Linux kernel's virtual memory (VM) subsystem. The kernel makes extensive and intelligent use of virtual memory, which is commonly referred to as swap space.

The following files are commonly found in the /proc/sys/vm/ directory:

- block_dump — Configures block I/O debugging when enabled. All read/write and block dirtying operations done to files are logged accordingly. This can be useful if diagnosing disk spin up and spin downs for laptop battery conservation. All output when block_dump is enabled can be retrieved via dmesg. The default value is 0.

  > **Tip**
  >
  > If block_dump is enabled at the same time as kernel debugging, it is prudent to stop the klogd daemon, as it generates erroneous disk activity caused by block_dump.

- dirty_background_ratio — Starts background writeback of dirty data at this percentage of total memory, via a pdflush daemon. The default value is 10.

- dirty_expire_centisecs — Defines when dirty in-memory data is old enough to be eligible for writeout. Data which has been dirty in-memory for longer than this interval is written out next time a pdflush daemon wakes up. The default value is 3000, expressed in hundredths of a second.
- dirty_ratio — Starts active writeback of dirty data at this percentage of total memory for the generator of dirty data, via pdflush. The default value is 40.
- dirty_writeback_centisecs — Defines the interval between pdflush daemon wakeups, which periodically writes dirty in-memory data out to disk. The default value is 500, expressed in hundredths of a second.
- laptop_mode — Minimizes the number of times that a hard disk needs to spin up by keeping the disk spun down for as long as possible, therefore conserving battery power on laptops. This increases efficiency by combining all future I/O processes together, reducing the frequency of spin ups. The default value is 0, but is automatically enabled in case a battery on a laptop is used.

  This value is controlled automatically by the acpid daemon once a user is notified battery power is enabled. No user modifications or interactions are necessary if the laptop supports the ACPI (Advanced Configuration and Power Interface) specification.

  For more information, refer to the following installed documentation:

  /usr/share/doc/kernel-doc-<version>/Documentation/laptop-mode.txt
- lower_zone_protection — Determines how aggressive the kernel is in defending lower memory allocation zones. This is effective when utilized with machines configured with highmem memory space enabled. The default value is 0, no protection at all. All other integer values are in megabytes, and lowmem memory is therefore protected from being allocated by users.

  For more information, refer to the following installed documentation:

  /usr/share/doc/kernel-doc-<version>/Documentation/filesystems/proc.txt
- max_map_count — Configures the maximum number of memory map areas a process may have. In most cases, the default value of 65536 is appropriate.
- min_free_kbytes — Forces the Linux VM (virtual memory manager) to keep a minimum number of kilobytes free. The VM uses this number to compute a pages_min value for each lowmem zone in the system. The default value is in respect to the total memory on the machine.
- nr_hugepages — Indicates the current number of configured hugetlb pages in the kernel.

  For more information, refer to the following installed documentation:

  /usr/share/doc/kernel-doc-<version>/Documentation/vm/hugetlbpage.txt
- nr_pdflush_threads — Indicates the number of pdflush daemons that are currently running. This file is read-only, and should not be changed by the user. Under heavy I/O loads, the default value of two is increased by the kernel.
- overcommit_memory — Configures the conditions under which a large memory request is accepted or denied. The following three modes are available:
  - 0 — The kernel performs heuristic memory over commit handling by estimating the amount of memory available and failing requests that are blatantly invalid. Unfortunately, since memory is allocated using a heuristic rather than a precise algorithm, this setting can sometimes allow available memory on the system to be overloaded. This is the default setting.

  - 1 — The kernel performs no memory over commit handling. Under this setting, the potential for memory overload is increased, but so is performance for memory intensive tasks (such as those executed by some scientific software).
  - 2 — The kernel fails requests for memory that add up to all of swap plus the percent of physical RAM specified in /proc/sys/vm/overcommit_ratio. This setting is best for those who desire less risk of memory overcommitment.

  > Note
  >
  > This setting is only recommended for systems with swap areas larger than physical memory.

- overcommit_ratio — Specifies the percentage of physical RAM considered when /proc/sys/vm/overcommit_memory is set to 2. The default value is 50.
- page-cluster — Sets the number of pages read in a single attempt. The default value of 3, which actually relates to 16 pages, is appropriate for most systems.
- swappiness — Determines how much a machine should swap. The higher the value, the more swapping occurs. The default value, as a percentage, is set to 60.

All kernel-based documentation can be found in the following locally installed location:

/usr/share/doc/kernel-doc-<version>/Documentation/, which contains additional information.

4.3.10. /proc/sysvipc/

This directory contains information about System V IPC resources. The files in this directory relate to System V IPC calls for messages (msg), semaphores (sem), and shared memory (shm).

4.3.11. /proc/tty/

This directory contains information about the available and currently used tty devices on the system. Originally called teletype devices, any character-based data terminals are called tty devices.

In Linux, there are three different kinds of tty devices. Serial devices are used with serial connections, such as over a modem or using a serial cable. Virtual terminals create the common console connection, such as the virtual consoles available when pressing Alt+<F-key> at the system console. Pseudo terminals create a two-way communication that is used by some higher level applications, such as XFree86. The drivers file is a list of the current tty devices in use, as in the following example:

```
serial               /dev/cua       5  64-127 serial:callout
serial               /dev/ttyS      4  64-127 serial
pty_slave            /dev/pts     136   0-255 pty:slave
pty_master           /dev/ptm     128   0-255 pty:master
pty_slave            /dev/ttyp      3   0-255 pty:slave
pty_master           /dev/pty       2   0-255 pty:master
/dev/vc/0            /dev/vc/0      4       0 system:vtmaster
/dev/ptmx            /dev/ptmx      5       2 system
/dev/console         /dev/console   5       1 system:console
```

```
/dev/tty             /dev/tty        5       0 system:/dev/tty
unknown              /dev/vc/%d      4    1-63 console
```

The /proc/tty/driver/serial file lists the usage statistics and status of each of the serial tty lines.

In order for tty devices to be used as network devices, the Linux kernel enforces line discipline on the device. This allows the driver to place a specific type of header with every block of data transmitted over the device, making it possible for the remote end of the connection to a block of data as just one in a stream of data blocks. SLIP and PPP are common line disciplines, and each are commonly used to connect systems to one other over a serial link.

Registered line disciplines are stored in the ldiscs file, and more detailed information is available within the ldisc/ directory.

4.3.12. /proc/<PID>/

Out of Memory (OOM) refers to a computing state where all available memory, including swap space, has been allocated. When this situation occurs, it will cause the system to panic and stop functioning as expected. There is a switch that controls OOM behavior in /proc/sys/vm/panic_on_oom. When set to 1 the kernel will panic on OOM. A setting of 0 instructs the kernel to call a function named oom_killer on an OOM. Usually, oom_killer can kill rogue processes and the system will survive.

The easiest way to change this is to echo the new value to /proc/sys/vm/panic_on_oom.

```
~]# cat /proc/sys/vm/panic_on_oom
1
~]# echo 0 > /proc/sys/vm/panic_on_oom
~]# cat /proc/sys/vm/panic_on_oom
0
```

It is also possible to prioritize which processes get killed by adjusting the oom_killer score. In /proc/<PID>/ there are two tools labelled oom_adj and oom_score. Valid scores for oom_adj are in the range -16 to +15. To see the current oom_killer score, view the oom_score for the process. oom_killer will kill processes with the highest scores first.

This example adjusts the oom_score of a process with a PID of 12465 to make it less likely that oom_killer will kill it.

```
~]# cat /proc/12465/oom_score
79872
~]# echo -5 > /proc/12465/oom_adj
~]# cat /proc/12465/oom_score
78
```

There is also a special value of -17, which disables oom_killer for that process. In the example below, oom_score returns a value of 0, indicating that this process would not be killed.

```
~]# cat /proc/12465/oom_score
78
~]# echo -17 > /proc/12465/oom_adj
~]# cat /proc/12465/oom_score
0
```

A function called badness() is used to determine the actual score for each process. This is done by adding up 'points' for each examined process. The process scoring is done in the following way:

1. The basis of each process's score is its memory size.

2. The memory size of any of the process's children (not including a kernel thread) is also added to the score

3. The process's score is increased for 'niced' processes and decreased for long running processes.

4. Processes with the CAP_SYS_ADMIN and CAP_SYS_RAWIO capabilities have their scores reduced.

5. The final score is then bitshifted by the value saved in the oom_adj file.

Thus, a process with the highest oom_score value will most probably be a non-privileged, recently started process that, along with its children, uses a large amount of memory, has been 'niced', and handles no raw I/O.

4.4. Using the sysctl Command

The /sbin/sysctl command is used to view, set, and automate kernel settings in the /proc/sys/ directory.

For a quick overview of all settings configurable in the /proc/sys/ directory, type the /sbin/sysctl -a command as root. This creates a large, comprehensive list, a small portion of which looks something like the following:

```
net.ipv4.route.min_delay = 2 kernel.sysrq = 0 kernel.sem = 250     32000     32     128
```

This is the same information seen if each of the files were viewed individually. The only difference is the file location. For example, the /proc/sys/net/ipv4/route/min_delay file is listed as net.ipv4.route.min_delay, with the directory slashes replaced by dots and the proc.sys portion assumed.

The sysctl command can be used in place of echo to assign values to writable files in the /proc/sys/ directory. For example, instead of using the command

```
echo 1 > /proc/sys/kernel/sysrq
```

use the equivalent sysctl command as follows:

```
~]# sysctl -w kernel.sysrq="1"
kernel.sysrq = 1
```

While quickly setting single values like this in /proc/sys/ is helpful during testing, this method does not work as well on a production system as special settings within /proc/sys/ are lost when the machine is rebooted. To preserve custom settings, add them to the /etc/sysctl.conf file.

Each time the system boots, the init program runs the /etc/rc.d/rc.sysinit script. This script contains a command to execute sysctl using /etc/sysctl.conf to determine the values passed to the kernel. Any values added to /etc/sysctl.conf therefore take effect each time the system boots.

4.5. Additional Resources

Below are additional sources of information about proc file system.

4.5.1. Installed Documentation

Some of the best documentation about the proc file system is installed on the system by default.

- /usr/share/doc/kernel-doc-<version>/Documentation/filesystems/proc.txt — Contains assorted, but limited, information about all aspects of the /proc/ directory.
- /usr/share/doc/kernel-doc-<version>/Documentation/sysrq.txt — An overview of System Request Key options.
- /usr/share/doc/kernel-doc-<version>/Documentation/sysctl/ — A directory containing a variety of sysctl tips, including modifying values that concern the kernel (kernel.txt), accessing file systems (fs.txt), and virtual memory use (vm.txt).
- /usr/share/doc/kernel-doc-<version>/Documentation/networking/ip-sysctl.txt — A detailed overview of IP networking options.

4.5.2. Useful Websites

- http://www.linuxhq.com/ — This website maintains a complete database of source, patches, and documentation for various versions of the Linux kernel.

Redundant Array of Independent Disks (RAID)

The basic idea behind RAID is to combine multiple small, inexpensive disk drives into an array to accomplish performance or redundancy goals not attainable with one large and expensive drive. This array of drives appears to the computer as a single logical storage unit or drive.

5.1. What is RAID?

RAID allows information to access several disks. RAID uses techniques such as disk striping (RAID Level 0), disk mirroring (RAID Level 1), and disk striping with parity (RAID Level 5) to achieve redundancy, lower latency, increased bandwidth, and maximized ability to recover from hard disk crashes.

RAID consistently distributes data across each drive in the array. RAID then breaks down the data into consistently-sized chunks (commonly 32K or 64k, although other values are acceptable). Each chunk is then written to a hard drive in the RAID array according to the RAID level employed. When the data is read, the process is reversed, giving the illusion that the multiple drives in the array are actually one large drive.

5.1.1. Who Should Use RAID?

System Administrators and others who manage large amounts of data would benefit from using RAID technology. Primary reasons to deploy RAID include:

- Enhances speed
- Increases storage capacity using a single virtual disk
- Minimizes disk failure

5.1.2. Hardware RAID versus Software RAID

There are two possible RAID approaches: hardware RAID and software RAID.

Hardware RAID

The hardware-based array manages the RAID subsystem independently from the host. It presents a single disk per RAID array to the host.

A hardware RAID device connects to the SCSI controller and presents the RAID arrays as a single SCSI drive. An external RAID system moves all RAID handling "intelligence" into a controller located in the external disk subsystem. The whole subsystem is connected to the host via a normal SCSI controller and appears to the host as a single disk.

RAID controller cards function like a SCSI controller to the operating system, and handle all the actual drive communications. The user plugs the drives into the RAID controller (just like a normal SCSI controller) and then adds them to the RAID controllers configuration, and the operating system won't know the difference.

Software RAID

Software RAID implements the various RAID levels in the kernel disk (block device) code. It offers the cheapest possible solution, as expensive disk controller cards or hot-swap chassis[1] are not required. Software RAID also works with cheaper IDE disks as well as SCSI disks. With today's faster CPUs, software RAID outperforms hardware RAID.

The Linux kernel contains an MD driver that allows the RAID solution to be completely hardware independent. The performance of a software-based array depends on the server CPU performance and load.

To learn more about software RAID, here are the key features:

- Threaded rebuild process
- Kernel-based configuration
- Portability of arrays between Linux machines without reconstruction
- Backgrounded array reconstruction using idle system resources
- Hot-swappable drive support
- Automatic CPU detection to take advantage of certain CPU optimizations

5.1.3. RAID Levels and Linear Support

RAID supports various configurations, including levels 0, 1, 4, 5, and linear. These RAID types are defined as follows:

Level 0

RAID level 0, often called "striping", is a performance-oriented striped data mapping technique. This means the data being written to the array is broken down into strips and written across the member disks of the array, allowing high I/O performance at low inherent cost but provides no redundancy. The storage capacity of a level 0 array is equal to the total capacity of the member disks in a hardware RAID or the total capacity of member partitions in a software RAID.

Level 1

RAID level 1, or "mirroring", has been used longer than any other form of RAID. Level 1 provides redundancy by writing identical data to each member disk of the array, leaving a "mirrored" copy on each disk. Mirroring remains popular due to its simplicity and high level of data availability. Level 1 operates with two or more disks that may use parallel access for high data-transfer rates when reading but more commonly operate independently to provide high I/O transaction rates. Level 1 provides very good data reliability and improves performance for read-intensive applications but at a relatively high cost. The storage capacity of the level 1 array is equal to the capacity of one of the mirrored hard disks in a hardware RAID or one of the mirrored partitions in a software RAID.

[1] A hot-swap chassis allows you to remove a hard drive without having to power-down your system.

알림

RAID level 1 comes at a high cost because you write the same information to all of the disks in the array, which wastes drive space. For example, if you have RAID level 1 set up so that your root (/) partition exists on two 40G drives, you have 80G total but are only able to access 40G of that 80G. The other 40G acts like a mirror of the first 40G.

Level 4

RAID level 4 uses parity[2] concentrated on a single disk drive to protect data. It is better suited to transaction I/O rather than large file transfers. Because the dedicated parity disk represents an inherent bottleneck, level 4 is seldom used without accompanying technologies such as write-back caching. Although RAID level 4 is an option in some RAID partitioning schemes, it is not an option allowed in Red Hat Enterprise Linux RAID installations. The storage capacity of hardware RAID level 4 is equal to the capacity of member disks, minus the capacity of one member disk. The storage capacity of software RAID level 4 is equal to the capacity of the member partitions, minus the size of one of the partitions if they are of equal size.

알림

RAID level 4 takes up the same amount of space as RAID level 5, but level 5 has more advantages. For this reason, level 4 is not supported.

Level 5

RAID level 5 is the most common type of RAID. By distributing parity across some or all of an array's member disk drives, RAID level 5 eliminates the write bottleneck inherent in level 4. The only performance bottleneck is the parity calculation process. With modern CPUs and software RAID, that usually is not a very big problem. As with level 4, the result is asymmetrical performance, with reads substantially outperforming writes. Level 5 is often used with write-back caching to reduce the asymmetry. The storage capacity of hardware RAID level 5 is equal to the capacity of member disks, minus the capacity of one member disk. The storage capacity of software RAID level 5 is equal to the capacity of the member partitions, minus the size of one of the partitions if they are of equal size.

Linear RAID

Linear RAID is a simple grouping of drives to create a larger virtual drive. In linear RAID, the chunks are allocated sequentially from one member drive, going to the next drive only when the first is completely filled. This grouping provides no performance benefit, as it is unlikely that any I/O operations will be split between member drives. Linear RAID also offers no redundancy and, in fact, decreases reliability — if any one member drive fails, the entire array cannot be used. The capacity is the total of all member disks.

[2] Parity information is calculated based on the contents of the rest of the member disks in the array. This information can then be used to reconstruct data when one disk in the array fails. The reconstructed data can then be used to satisfy I/O requests to the failed disk before it is replaced and to repopulate the failed disk after it has been replaced.

5.2. Configuring Software RAID

Users can configure software RAID during the graphical installation process, the text-based installation process, or during a kickstart installation. This section discusses software RAID configuration during the installation process using the Disk Druid application, and covers the following steps:

1. Creating software RAID partitions on physical hard drives.

2. Creating RAID devices from the software RAID partitions.

3. (Optional) Configuring LVM from the RAID devices.

4. Creating file systems from the RAID devices.

To configure software RAID, select Create custom layout from the pulldown list on the Disk Partitioning Setup screen, click the Next button, and follow the instructions in the rest of this section. The example screenshots in this section use two 10 GB disk drives (/dev/hda and /dev/hdb) to illustrate the creation of simple RAID 1 and RAID 0 configurations, and detail how to create a simple RAID configuration by implementing multiple RAID devices.

5.2.1. Creating the RAID Partitions

In a typical situation, the disk drives are new or are formatted. Both drives are shown as raw devices with no partition configuration in 그림 5.1. "Two Blank Drives, Ready For Configuration" .

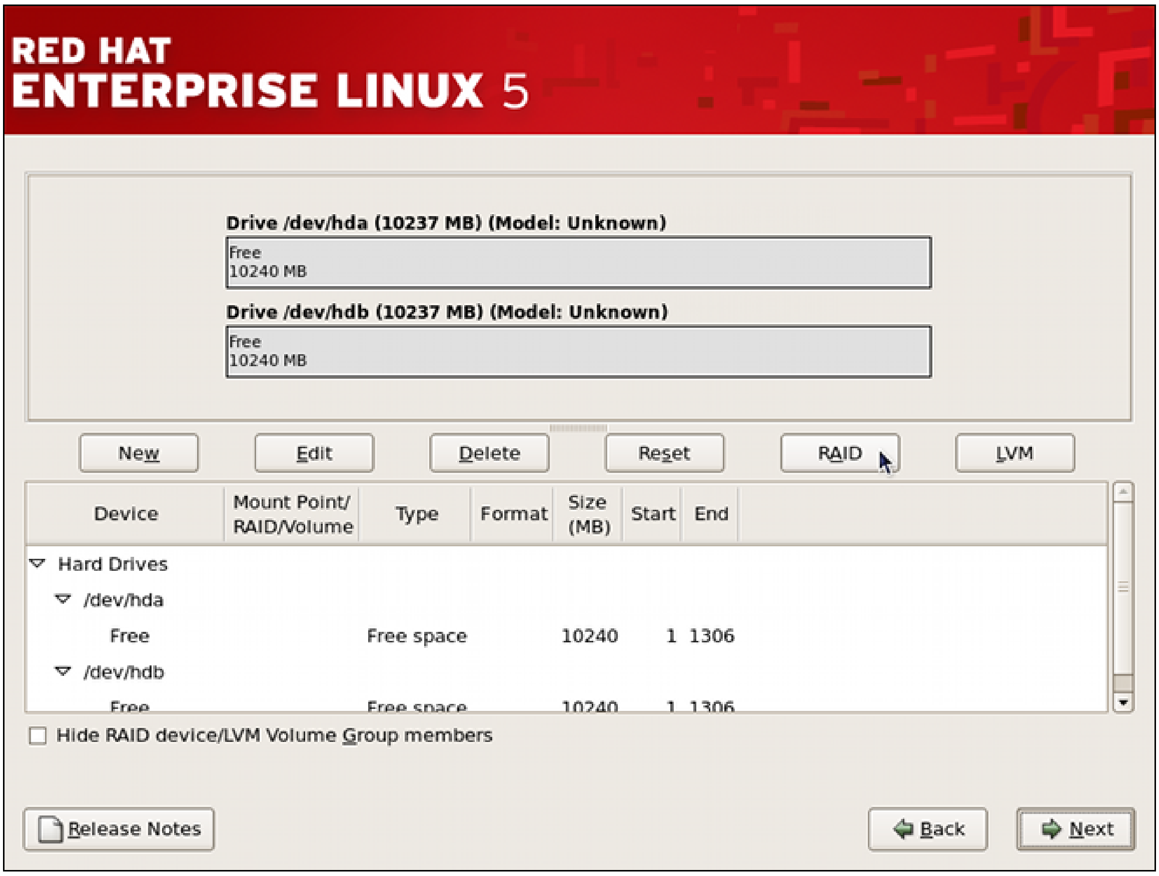

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
Hard Drives						
/dev/hda						
Free		Free space		10240	1	1306
/dev/hdb						
Free		Free space		10240	1	1306

□ Hide RAID device/LVM Volume Group members

그림 5.1. Two Blank Drives, Ready For Configuration

1. In Disk Druid, click the RAID button to enter the software RAID creation screen.

2. Choose Create a software RAID partition to create a RAID partition as shown in 그림 5.2. "RAID Partition Options" . Note that no other RAID options (such as entering a mount point) are available until RAID partitions, as well as RAID devices, are created. Click OK to confirm the choice.

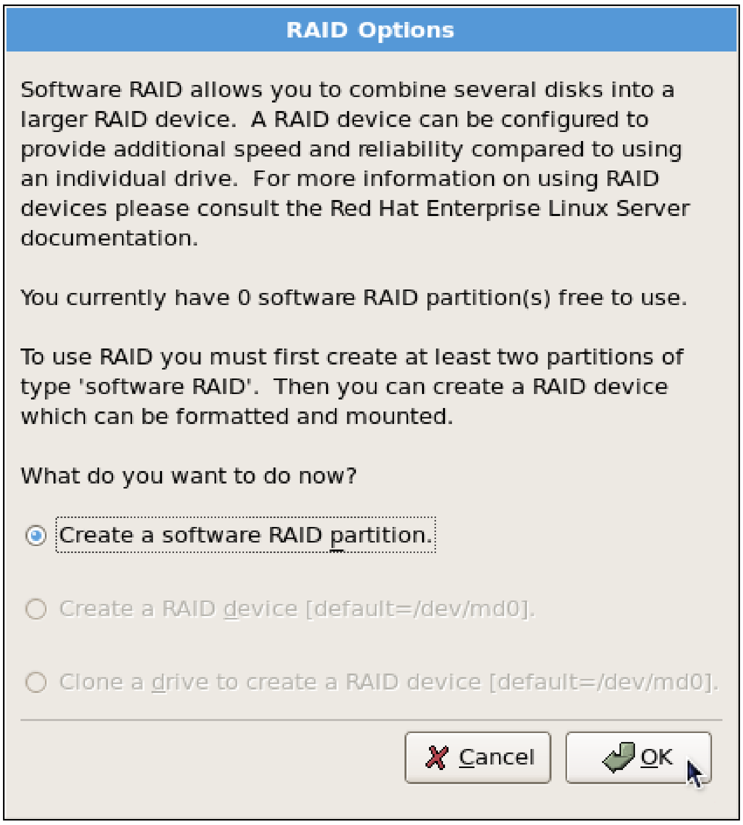

RAID Options

Software RAID allows you to combine several disks into a larger RAID device. A RAID device can be configured to provide additional speed and reliability compared to using an individual drive. For more information on using RAID devices please consult the Red Hat Enterprise Linux Server documentation.

You currently have 0 software RAID partition(s) free to use.

To use RAID you must first create at least two partitions of type 'software RAID'. Then you can create a RAID device which can be formatted and mounted.

What do you want to do now?

● Create a software RAID partition.

○ Create a RAID device [default=/dev/md0].

○ Clone a drive to create a RAID device [default=/dev/md0].

그림 5.2. RAID Partition Options

3. A software RAID partition must be constrained to one drive. For Allowable Drives, select the drive to use for RAID. If you have multiple drives, by default all drives are selected and you must deselect the drives you do not want.

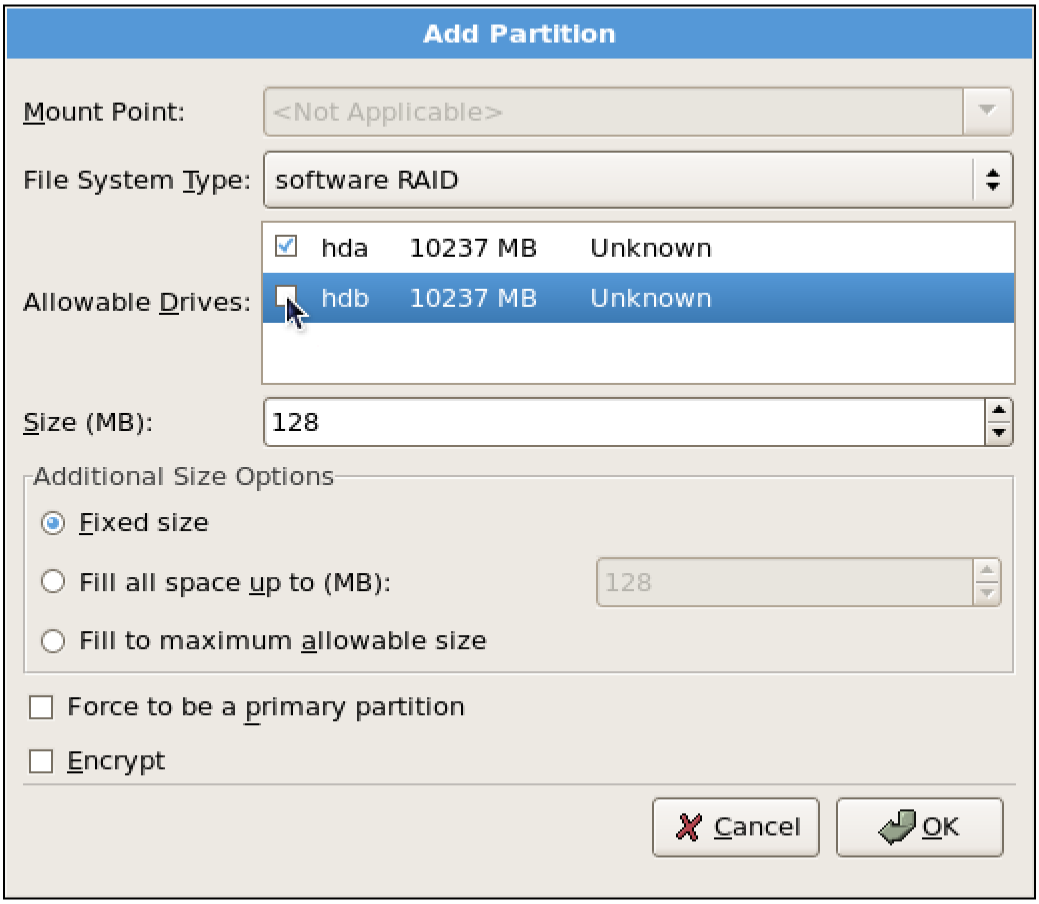

Add Partition

Mount Point: <Not Applicable>

File System Type: software RAID

Allowable Drives:

☑	hda	10237 MB	Unknown
□	hdb	10237 MB	Unknown

Size (MB): 128

Additional Size Options

● Fixed size

○ Fill all space up to (MB): 128

○ Fill to maximum allowable size

□ Force to be a primary partition

□ Encrypt

그림 5.3. Adding a RAID Partition

4. Edit the Size (MB) field, and enter the size that you want the partition to be (in MB).

5. Select Fixed Size to specify partition size. Select Fill all space up to (MB) and enter a value (in MB) to specify partition size range. Select Fill to maximum allowable size to allow maximum available space of the hard disk. Note that if you make more than one space growable, they share the available free space on the disk.

6. Select Force to be a primary partition if you want the partition to be a primary partition. A primary partition is one of the first four partitions on the hard drive. If unselected, the partition is created as a logical partition. If other operating systems are already on the system, unselecting this option should be considered. For more information on primary versus logical/extended partitions, refer to the appendix section of the Red Hat Enterprise Linux Installation Guide.

Repeat these steps to create as many partitions as needed for your RAID setup. Notice that all the partitions do not have to be RAID partitions. For example, you can configure only the /boot partition as a software RAID device, leaving the root partition (/), /home, and swap as regular file systems. 그림 5.4. "RAID 1 Partitions Ready, Pre-Device and Mount Point Creation" shows successfully allocated space for the RAID 1 configuration (for /boot), which is now ready for RAID device and mount point creation:

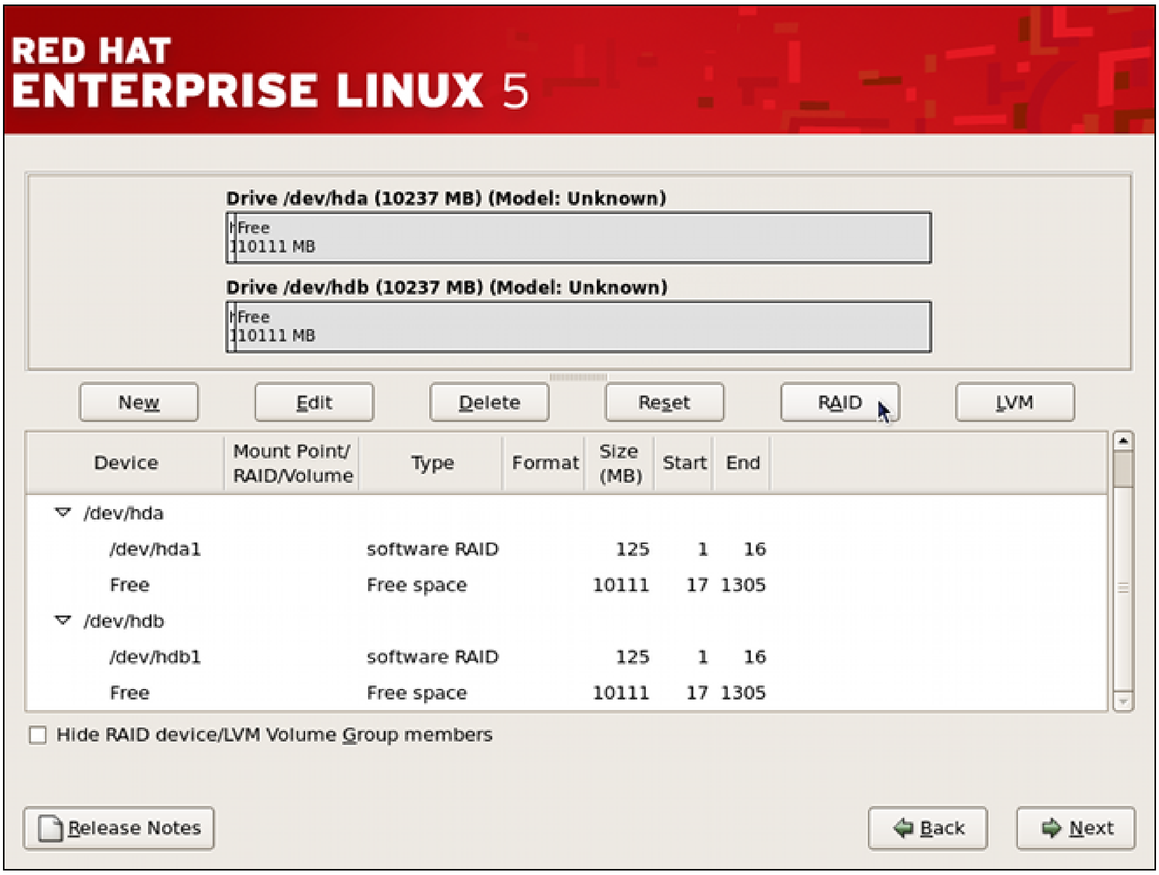

Drive /dev/hda (10237 MB) (Model: Unknown)

Free
10111 MB

Drive /dev/hdb (10237 MB) (Model: Unknown)

Free
10111 MB

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
/dev/hda						
/dev/hda1		software RAID		125	1	16
Free		Free space		10111	17	1305
/dev/hdb						
/dev/hdb1		software RAID		125	1	16
Free		Free space		10111	17	1305

☐ Hide RAID device/LVM Volume Group members

그림 5.4. RAID 1 Partitions Ready, Pre-Device and Mount Point Creation

5.2.2. Creating the RAID Devices and Mount Points

Once you create all of your partitions as software RAID partitions, you must create the RAID device and mount point.

1. On the main partitioning screen, click the RAID button. The RAID Options dialog appears as shown in 그림 5.5. “RAID 옵션” .

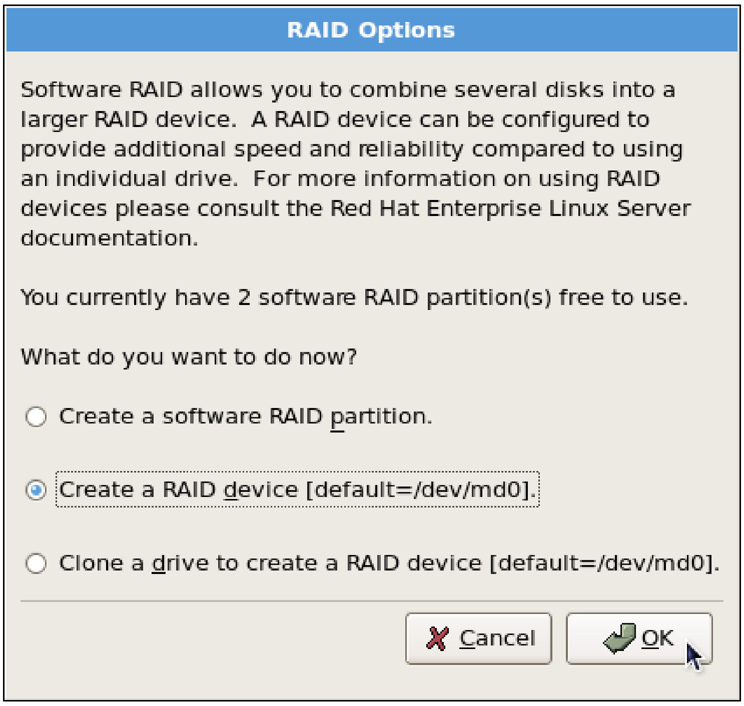

RAID Options

Software RAID allows you to combine several disks into a larger RAID device. A RAID device can be configured to provide additional speed and reliability compared to using an individual drive. For more information on using RAID devices please consult the Red Hat Enterprise Linux Server documentation.

You currently have 2 software RAID partition(s) free to use.

What do you want to do now?

○ Create a software RAID partition.

● Create a RAID device [default=/dev/md0].

○ Clone a drive to create a RAID device [default=/dev/md0].

그림 5.5. RAID 옵션

2. Select the Create a RAID device option, and click OK. As shown in 그림 5.6. “Making a RAID Device and Assigning a Mount Point” , the Make RAID Device dialog appears, allowing you to make a RAID device and assign a mount point.

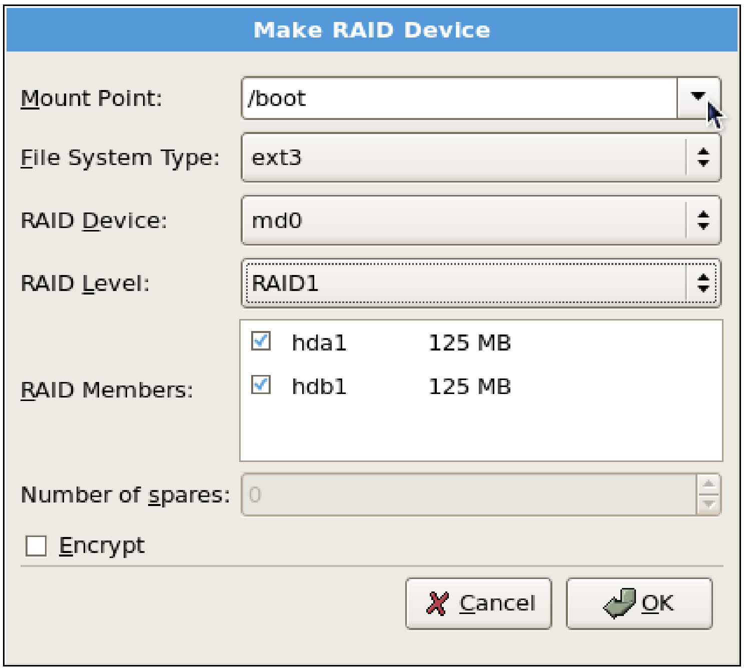

Make RAID Device

Mount Point: /boot

File System Type: ext3

RAID Device: md0

RAID Level: RAID1

RAID Members:
☑ hda1 125 MB
☑ hdb1 125 MB

Number of spares: 0

□ Encrypt

그림 5.6. Making a RAID Device and Assigning a Mount Point

3. Select a mount point from the Mount Point pulldown list.

4. Choose the file system type for the partition from the File System Type pulldown list. At this point you can either configure a dynamic LVM file system or a traditional static ext2/ext3 file

system. For more information on LVM and its configuration during the installation process, refer to 10장. LVM (Logical Volume Manager). If LVM is not required, continue on with the following instructions.

5. From the RAID Device pulldown list, select a device name such as md0.

6. From the RAID Level, choose the required RAID level.

> 알림
>
> If you are making a RAID partition of /boot, you must choose RAID level 1, and it must use one of the first two drives (IDE first, SCSI second). If you are not creating a separate RAID partition of /boot, and you are making a RAID partition for the root file system (that is, /), it must be RAID level 1 and must use one of the first two drives (IDE first, SCSI second).

7. The RAID partitions created appear in the RAID Members list. Select which of these partitions should be used to create the RAID device.

8. If configuring RAID 1 or RAID 5, specify the number of spare partitions in the Number of spares field. If a software RAID partition fails, the spare is automatically used as a replacement. For each spare you want to specify, you must create an additional software RAID partition (in addition to the partitions for the RAID device). Select the partitions for the RAID device and the partition(s) for the spare(s).

9. Click OK to confirm the setup. The RAID device appears in the Drive Summary list.

10. Repeat this chapter's entire process for configuring additional partitions, devices, and mount points, such as the root partition (/), home partition (/home), or swap.

After completing the entire configuration, the figure as shown in 그림 5.7. "Sample RAID Configuration" resembles the default configuration, except for the use of RAID.

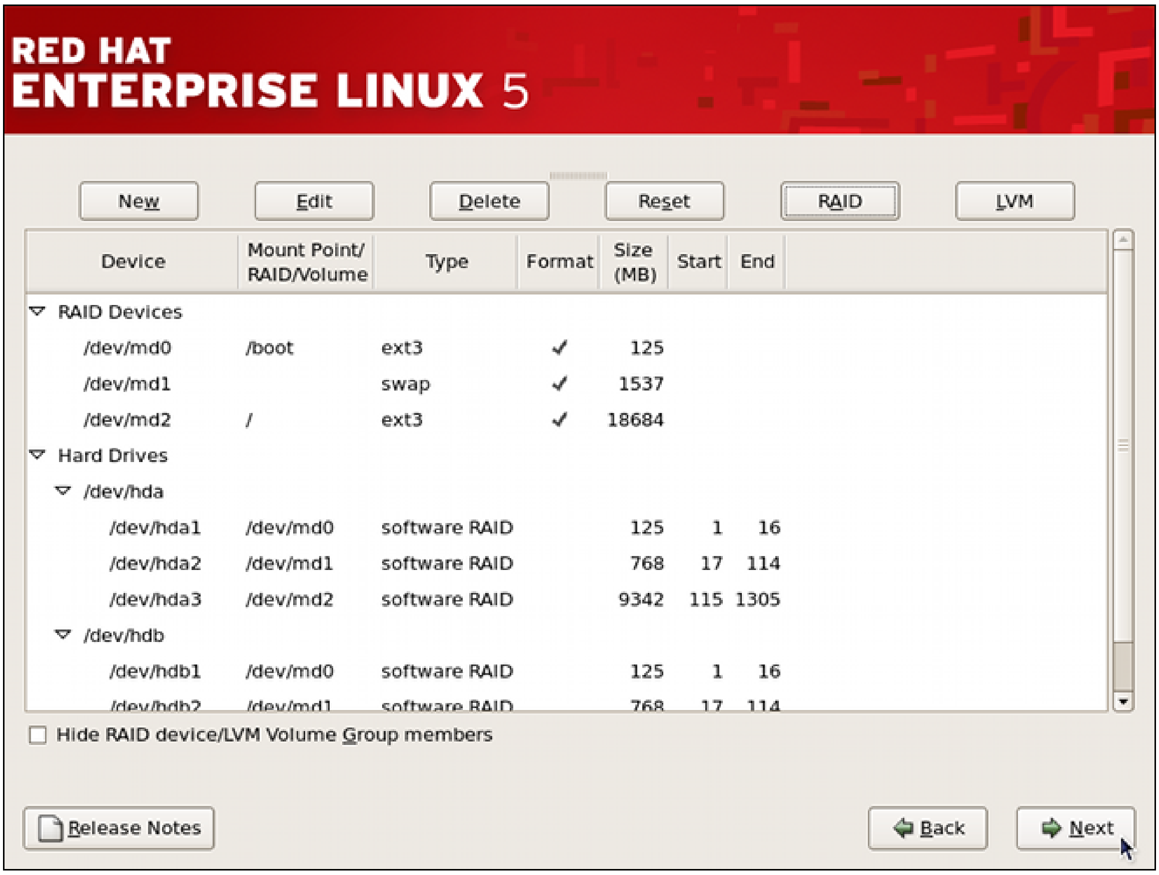

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
▽ RAID Devices						
/dev/md0	/boot	ext3	✓	125		
/dev/md1		swap	✓	1537		
/dev/md2	/	ext3	✓	18684		
▽ Hard Drives						
▽ /dev/hda						
/dev/hda1	/dev/md0	software RAID		125	1	16
/dev/hda2	/dev/md1	software RAID		768	17	114
/dev/hda3	/dev/md2	software RAID		9342	115	1305
▽ /dev/hdb						
/dev/hdb1	/dev/md0	software RAID		125	1	16
/dev/hdb2	/dev/md1	software RAID		768	17	114

☐ Hide RAID device/LVM Volume Group members

그림 5.7. Sample RAID Configuration

The figure as shown in 그림 5.8. "Sample RAID With LVM Configuration" is an example of a RAID and LVM configuration.

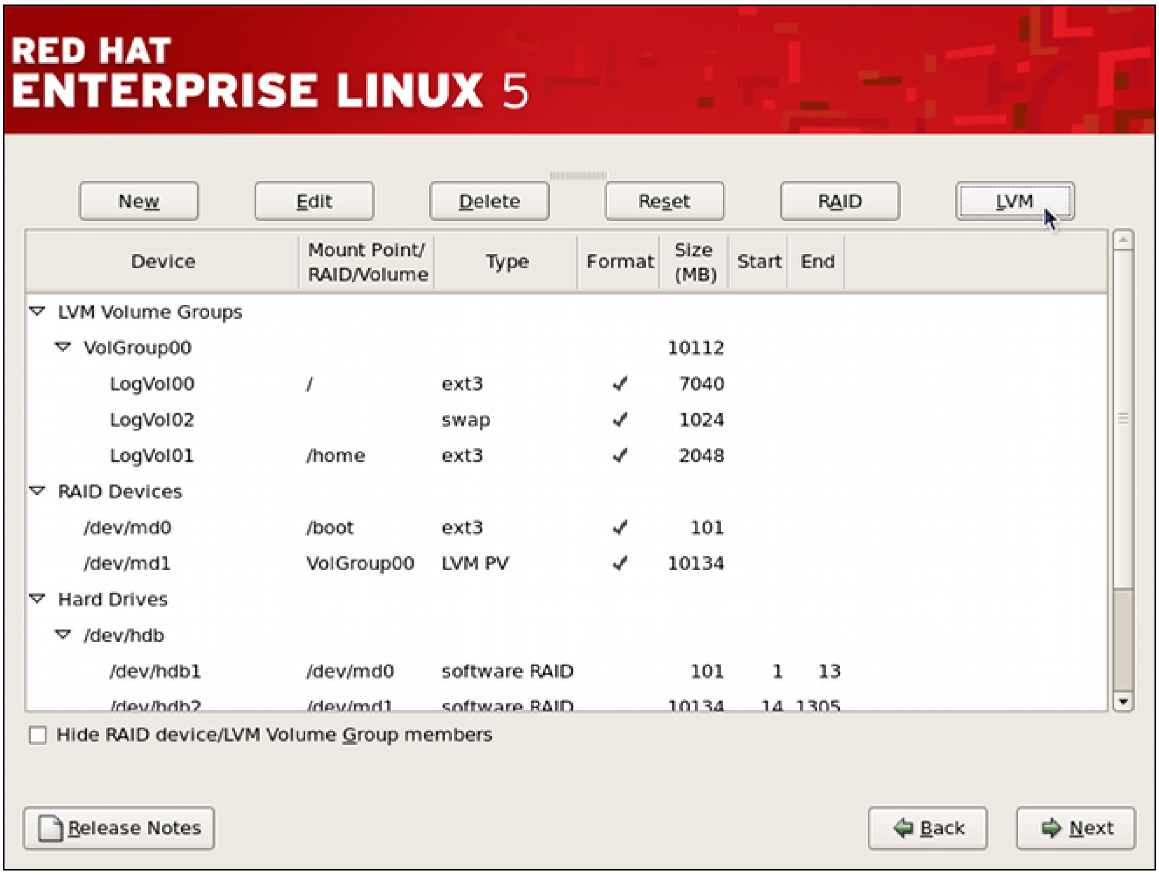

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
▽ LVM Volume Groups						
▽ VolGroup00				10112		
LogVol00	/	ext3	✓	7040		
LogVol02		swap	✓	1024		
LogVol01	/home	ext3	✓	2048		
▽ RAID Devices						
/dev/md0	/boot	ext3	✓	101		
/dev/md1	VolGroup00	LVM PV	✓	10134		
▽ Hard Drives						
▽ /dev/hdb						
/dev/hdb1	/dev/md0	software RAID		101	1	13
/dev/hdb2	/dev/md1	software RAID		10134	14	1305

□ Hide RAID device/LVM Volume Group members

그림 5.8. Sample RAID With LVM Configuration

You can proceed with your installation process by clicking Next. Refer to the Red Hat Enterprise Linux Installation Guide for further instructions.

5.3. Managing Software RAID

This section discusses software RAID configuration and management after the installation, and covers the following topics:

- Reviewing existing software RAID configuration.
- Creating a new RAID device.
- Replacing a faulty device in an array.
- Adding a new device to an existing array.
- Deactivating and removing an existing RAID device.
- Saving the configuration.

All examples in this section use the software RAID configuration from the previous section.

5.3.1. Reviewing RAID Configuration

When a software RAID is in use, basic information about all presently active RAID devices are stored in the /proc/mdstat special file. To list these devices, display the content of this file by typing the following at a shell prompt:

```
cat /proc/mdstat
```

To determine whether a certain device is a RAID device or a component device, run the command in the following form as root:

```
mdadm --query device…
```

In order to examine a RAID device in more detail, use the following command:

```
mdadm --detail raid_device…
```

Similarly, to examine a component device, type:

```
mdadm --examine component_device…
```

While the mdadm --detail command displays information about a RAID device, mdadm --examine only relays information about a RAID device as it relates to a given component device. This distinction is particularly important when working with a RAID device that itself is a component of another RAID device.

The mdadm --query command, as well as both mdadm --detail and mdadm --examine commands allow you to specify multiple devices at once.

예 5.1. Reviewing RAID configuration

Assume the system uses configuration from 그림 5.7. "Sample RAID Configuration" . You can verify that /dev/md0 is a RAID device by typing the following at a shell prompt:

```
~]# mdadm --query /dev/md0
/dev/md0: 125.38MiB raid1 2 devices, 0 spares. Use mdadm --detail for more detail.
/dev/md0: No md super block found, not an md component.
```

As you can see, the above command produces only a brief overview of the RAID device and its configuration. To display more detailed information, use the following command instead:

```
~]# mdadm --detail /dev/md0
/dev/md0:
        Version : 0.90
  Creation Time : Tue Jun 28 16:05:49 2011
     Raid Level : raid1
     Array Size : 128384 (125.40 MiB 131.47 MB)
  Used Dev Size : 128384 (125.40 MiB 131.47 MB)
   Raid Devices : 2
  Total Devices : 2
Preferred Minor : 0
    Persistence : Superblock is persistent

    Update Time : Thu Jun 30 17:06:34 2011
          State : clean
 Active Devices : 2
Working Devices : 2
 Failed Devices : 0
  Spare Devices : 0

           UUID : 49c5ac74:c2b79501:5c28cb9c:16a6dd9f
         Events : 0.6

    Number   Major   Minor   RaidDevice State
```

```
        0       3       1       0       active sync   /dev/hda1
        1       3       65      1       active sync   /dev/hdb1
```

Finally, to list all presently active RAID devices, type:

```
~]$ cat /proc/mdstat
Personalities : [raid0] [raid1]
md0 : active raid1 hdb1[1] hda1[0]
      128384 blocks [2/2] [UU]

md1 : active raid0 hdb2[1] hda2[0]
      1573888 blocks 256k chunks

md2 : active raid0 hdb3[1] hda3[0]
      19132928 blocks 256k chunks

unused devices: <none>
```

5.3.2. Creating a New RAID Device

To create a new RAID device, use the command in the following form as root:

```
mdadm --create raid_device --level=level --raid-devices=number component_device…
```

This is the simplest way to create a RAID array. There are many more options that allow you to specify the number of spare devices, the block size of a stripe array, if the array has a write-intent bitmap, and much more. All these options can have a significant impact on the performance, but are beyond the scope of this document. For more detailed information, refer to the CREATE MODE section of the mdadm(8) manual page.

예 5.2. Creating a new RAID device

Assume that the system has two unused SCSI disk drives available, and that each of these devices has exactly one partition of the same size:

```
~]# ls /dev/sd*
/dev/sda  /dev/sda1  /dev/sdb  /dev/sdb1
```

To create /dev/md3 as a new RAID level 1 array from /dev/sda1 and /dev/sdb1, run the following command:

```
~]# mdadm --create /dev/md3 --level=1 --raid-devices=2 /dev/sda1 /dev/sdb1
mdadm: array /dev/md3 started.
```

5.3.3. Replacing a Faulty Device

To replace a particular device in a software RAID, first make sure it is marked as faulty by running the following command as root:

```
mdadm raid_device --fail component_device
```

Then remove the faulty device from the array by using the command in the following form:

```
mdadm raid_device --remove component_device
```

Once the device is operational again, you can re-add it to the array:

```
mdadm raid_device --add component_device
```

예 5.3. Replacing a faulty device

Assume the system has an active RAID device, /dev/md3, with the following layout (that is, the RAID device created in 예 5.2. “Creating a new RAID device”):

```
~]# mdadm --detail /dev/md3 | tail -n 3
    Number   Major   Minor   RaidDevice State
       0       8       1        0      active sync   /dev/sda1
       1       8      17        1      active sync   /dev/sdb1
```

Imagine the first disk drive fails and needs to be replaced. To do so, first mark the /dev/sdb1 device as faulty:

```
~]# mdadm /dev/md3 --fail /dev/sdb1
mdadm: set /dev/sdb1 faulty in /dev/md3
```

Then remove it from the RAID device:

```
~]# mdadm /dev/md3 --remove /dev/sdb1
mdadm: hot removed /dev/sdb1
```

As soon as the hardware is replaced, you can add the device back to the array by using the following command:

```
~]# mdadm /dev/md3 --add /dev/sdb1
mdadm: added /dev/sdb1
```

5.3.4. Extending a RAID Device

To add a new device to an existing array, use the command in the following form as root:

```
mdadm raid_device --add component_device
```

This will add the device as a spare device. To grow the array to use this device actively, type the following at a shell prompt:

```
mdadm --grow raid_device --raid-devices=number
```

예 5.4. Extending a RAID device

Assume the system has an active RAID device, /dev/md3, with the following layout (that is, the RAID device created in 예 5.2. “Creating a new RAID device”):

```
~]# mdadm --detail /dev/md3 | tail -n 3
    Number   Major   Minor   RaidDevice State
       0       8       1        0      active sync   /dev/sda1
       1       8      17        1      active sync   /dev/sdb1
```

Also assume that a new SCSI disk drive, /dev/sdc, has been added and has exactly one partition. To add it to the /dev/md3 array, type the following at a shell prompt:

```
~]# mdadm /dev/md3 --add /dev/sdc1
mdadm: added /dev/sdc1
```

This will add /dev/sdc1 as a spare device. To change the size of the array to actually use it, type:

```
~]# mdadm --grow /dev/md3 --raid-devices=3
```

5.3.5. Removing a RAID Device

To remove an existing RAID device, first deactivate it by running the following command as root:

```
mdadm --stop raid_device
```

Once deactivated, remove the RAID device itself:

```
mdadm --remove raid_device
```

Finally, zero superblocks on all devices that were associated with the particular array:

```
mdadm --zero-superblock component_device…
```

예 5.5. Removing a RAID device

Assume the system has an active RAID device, /dev/md3, with the following layout (that is, the RAID device created in 예 5.4. “Extending a RAID device”):

```
~]# mdadm --detail /dev/md3 | tail -n 4
    Number   Major   Minor   RaidDevice State
       0       8        1        0      active sync   /dev/sda1
       1       8       17        1      active sync   /dev/sdb1
       2       8       33        2      active sync   /dev/sdc1
```

In order to remove this device, first stop it by typing the following at a shell prompt:

```
~]# mdadm --stop /dev/md3
mdadm: stopped /dev/md3
```

Once stopped, you can remove the /dev/md3 device by running the following command:

```
~]# mdadm --remove /dev/md3
```

Finally, to remove the superblocks from all associated devices, type:

```
~]# mdadm --zero-superblock /dev/sda1 /dev/sdb1 /dev/sdc1
```

5.3.6. Preserving the Configuration

By default, changes made by the mdadm command only apply to the current session, and will not survive a system restart. At boot time, the mdmonitor service reads the content of the /etc/mdadm.conf configuration file to see which RAID devices to start. If the software RAID was configured during the graphical installation process, this file contains directives listed in 표 5.1. “Common mdadm.conf directives” by default.

표 5.1. Common mdadm.conf directives

Option	Description
ARRAY	Allows you to identify a particular array.
DEVICE	Allows you to specify a list of devices to scan for a RAID component (for example, "/dev/hda1"). You can also use the keyword partitions to use all partitions listed in /proc/partitions, or containers to specify an array container.
MAILADDR	Allows you to specify an email address to use in case of an alert.

To list what ARRAY lines are presently in use regardless of the configuration, run the following command as root:

```
mdadm --detail --scan
```

Use the output of this command to determine which lines to add to the /etc/mdadm.conf file. You can also display the ARRAY line for a particular device:

```
mdadm --detail --brief raid_device
```

By redirecting the output of this command, you can add such a line to the configuration file with a single command:

```
mdadm --detail --brief raid_device >> /etc/mdadm.conf
```

> 예 5.6. Preserving the configuration
>
> By default, the /etc/mdadm.conf contains the software RAID configuration created during the system installation:
>
> ```
> # mdadm.conf written out by anaconda
> DEVICE partitions
> MAILADDR root
> ARRAY /dev/md0 level=raid1 num-devices=2 UUID=49c5ac74:c2b79501:5c28cb9c:16a6dd9f
> ARRAY /dev/md1 level=raid0 num-devices=2 UUID=76914c11:5bfa2c00:dc6097d1:a1f4506d
> ARRAY /dev/md2 level=raid0 num-devices=2 UUID=2b5d38d0:aea898bf:92be20e2:f9d893c5
> ```
>
> Assuming you have created the /dev/md3 device as shown in 예 5.2. "Creating a new RAID device" , you can make it persistent by running the following command:
>
> ```
> ~]# mdadm --detail --brief /dev/md3 >> /etc/mdadm.conf
> ```

5.4. Additional Resources

For more information on RAID, refer to the following resources.

5.4.1. Installed Documentation

- mdadm man page — A manual page for the mdadm utility.

- mdadm.conf man page — A manual page that provides a comprehensive list of available /etc/mdadm.conf configuration options.

스왑 공간

6.1. 스왑 공간이란?

Swap space in Linux is used when the amount of physical memory (RAM) is full. If the system needs more memory resources and the RAM is full, inactive pages in memory are moved to the swap space. While swap space can help machines with a small amount of RAM, it should not be considered a replacement for more RAM. Swap space is located on hard drives, which have a slower access time than physical memory.

스왑 공간은 스왑 파티션에 사용되거나 (권장 사항), 스왑 파일을 저장하는데 사용되며, 또는 스왑 파티션과 스왑 파일이 함께 스왑 공간을 차지하는 것도 가능합니다.

In years past, the recommended amount of swap space increased linearly with the amount of RAM in the system. But because the amount of memory in modern systems has increased into the hundreds of gigabytes, it is now recognized that the amount of swap space that a system needs is a function of the memory workload running on that system. However, given that swap space is usually designated at install time, and that it can be difficult to determine beforehand the memory workload of a system, we recommend determining system swap using the following table.

표 6.1. Recommended System Swap Space

Amount of RAM in the System	Recommended Amount of Swap Space
4GB of RAM or less	a minimum of 2GB of swap space
4GB to 16GB of RAM	a minimum of 4GB of swap space
16GB to 64GB of RAM	a minimum of 8GB of swap space
64GB to 256GB of RAM	a minimum of 16GB of swap space
256GB to 512GB of RAM	a minimum of 32GB of swap space

Important

File systems and LVM2 volumes assigned as swap space cannot be in use when being modified. For example, no system processes can be assigned the swap space, as well as no amount of swap should be allocated and used by the kernel. Use the free and cat /proc/swaps commands to verify how much and where swap is in use.

The best way to achieve swap space modifications is to boot your system in rescue mode, and then follow the instructions (for each scenario) in the remainder of this chapter. Refer to the Red Hat Enterprise Linux Installation Guide for instructions on booting into rescue mode. When prompted to mount the file system, select Skip.

6.2. 스왑 공간 추가하기

Sometimes it is necessary to add more swap space after installation. For example, you may upgrade the amount of RAM in your system from 128 MB to 256 MB, but there is only 256 MB of swap space. It might be advantageous to increase the amount of swap space to 512 MB if you perform memory-intense operations or run applications that require a large amount of memory.

You have three options: create a new swap partition, create a new swap file, or extend swap on an existing LVM2 logical volume. It is recommended that you extend an existing logical volume.

6.2.1. Extending Swap on an LVM2 Logical Volume

To extend an LVM2 swap logical volume (assuming /dev/VolGroup00/LogVol01 is the volume you want to extend):

1. Disable swapping for the associated logical volume:

   ```
   swapoff -v /dev/VolGroup00/LogVol01
   ```

2. Resize the LVM2 logical volume by 256 MB:

   ```
   lvm lvresize /dev/VolGroup00/LogVol01 -L +256M
   ```

3. Format the new swap space:

   ```
   mkswap /dev/VolGroup00/LogVol01
   ```

4. Enable the extended logical volume:

   ```
   swapon -va
   ```

5. Test that the logical volume has been extended properly:

   ```
   cat /proc/swaps
   free
   ```

6.2.2. Creating an LVM2 Logical Volume for Swap

To add a swap volume group (assuming /dev/VolGroup00/LogVol02 is the swap volume you want to add):

1. Create the LVM2 logical volume of size 256 MB:

   ```
   lvm lvcreate VolGroup00 -n LogVol02 -L 256M
   ```

2. Format the new swap space:

   ```
   mkswap /dev/VolGroup00/LogVol02
   ```

3. Add the following entry to the /etc/fstab file:

   ```
   /dev/VolGroup00/LogVol02   swap     swap    defaults     0 0
   ```

4. Enable the extended logical volume:

```
swapon -va
```

5. Test that the logical volume has been extended properly:

```
cat /proc/swaps
free
```

6.2.3. Creating a Swap File

스왑 파일을 추가하시려면:

1. Determine the size of the new swap file in megabytes and multiply by 1024 to determine the number of blocks. For example, the block size of a 64 MB swap file is 65536.

2. At a shell prompt as root, type the following command with count being equal to the desired block size:

```
dd if=/dev/zero of=/swapfile bs=1024 count=65536
```

3. 다음 명령을 사용하여 스왑 파일을 설정합니다:

```
mkswap /swapfile
```

4. 스왑 파일을 즉시 활성화하시려면 다음 명령을 입력해 주십시오:

```
swapon /swapfile
```

5. To enable it at boot time, edit /etc/fstab to include the following entry:

```
/swapfile          swap            swap    defaults        0 0
```

   다음에 시스템 부팅시 새로운 스왑 파일이 활성화됩니다.

6. After adding the new swap file and enabling it, verify it is enabled by viewing the output of the command cat /proc/swaps or free.

6.3. 스왑 공간 삭제하기

Sometimes it can be prudent to reduce swap space after installation. For example, say you downgraded the amount of RAM in your system from 1 GB to 512 MB, but there is 2 GB of swap space still assigned. It might be advantageous to reduce the amount of swap space to 1 GB, since the larger 2 GB could be wasting disk space.

You have three options: remove an entire LVM2 logical volume used for swap, remove a swap file, or reduce swap space on an existing LVM2 logical volume.

6.3.1. Reducing Swap on an LVM2 Logical Volume

To reduce an LVM2 swap logical volume (assuming /dev/VolGroup00/LogVol01 is the volume you want to reduce):

1. Disable swapping for the associated logical volume:

```
swapoff -v /dev/VolGroup00/LogVol01
```

2. Reduce the LVM2 logical volume by 512 MB:

```
lvm lvreduce /dev/VolGroup00/LogVol01 -L -512M
```

3. Format the new swap space:

```
mkswap /dev/VolGroup00/LogVol01
```

4. Enable the extended logical volume:

```
swapon -va
```

5. Test that the logical volume has been reduced properly:

```
cat /proc/swaps
free
```

6.3.2. Removing an LVM2 Logical Volume for Swap

The swap logical volume cannot be in use (no system locks or processes on the volume). The easiest way to achieve this is to boot your system in rescue mode. Refer to the Red Hat Enterprise Linux Installation Guide for instructions on booting into rescue mode. When prompted to mount the file system, select Skip.

To remove a swap volume group (assuming /dev/VolGroup00/LogVol02 is the swap volume you want to remove):

1. Disable swapping for the associated logical volume:

```
swapoff -v /dev/VolGroup00/LogVol02
```

2. Remove the LVM2 logical volume of size 512 MB:

```
lvm lvremove /dev/VolGroup00/LogVol02
```

3. Remove the following entry from the /etc/fstab file:

```
/dev/VolGroup00/LogVol02   swap     swap    defaults     0 0
```

4. Test that the logical volume has been removed:

```
cat /proc/swaps
free
```

6.3.3. Removing a Swap File

스왑 파일을 삭제하기 위해서는:

1. At a shell prompt as root, execute the following command to disable the swap file (where /swapfile is the swap file):

   ```
   swapoff -v /swapfile
   ```

2. Remove its entry from the /etc/fstab file.

3. 실제 파일을 삭제하십시오:

   ```
   rm /swapfile
   ```

6.4. 스왑 공간 이동하기

스왑 공간을 한 장소에서 다른 위치로 이동시키기 위해서는, 앞에서 설명된 방법에 따라서 스왑 공간을 삭제하신 후 새로운 장소에서 다시 새로운 스왑 공간을 생성하시기 바랍니다.

디스크 공간 관리

7.1. Standard Partitions using parted

The utility parted allows users to:

- View the existing partition table
- Change the size of existing partitions
- Add partitions from free space or additional hard drives

If you want to view the system's disk space usage or monitor the disk space usage, refer to 40.3절. "파일 시스템".

By default, the parted package is included when installing Red Hat Enterprise Linux. To start parted, log in as root and type the command parted /dev/sda at a shell prompt (where /dev/sda is the device name for the drive you want to configure).

If you want to remove or resize a partition, the device on which that partition resides must not be in use. Creating a new partition on a device which is in use—while possible—is not recommended.

For a device to not be in use, none of the partitions on the device can be mounted, and any swap space on the device must not be enabled.

As well, the partition table should not be modified while it is in use because the kernel may not properly recognize the changes. If the partition table does not match the actual state of the mounted partitions, information could be written to the wrong partition, resulting in lost and overwritten data.

The easiest way to achieve this is to boot your system in rescue mode. When prompted to mount the file system, select Skip.

Alternately, if the drive does not contain any partitions in use (system processes that use or lock the file system from being unmounted), you can unmount them with the umount command and turn off all the swap space on the hard drive with the swapoff command.

표 7.1. "parted commands" contains a list of commonly used parted commands. The sections that follow explain some of these commands and arguments in more detail.

표 7.1. parted commands

명령어	설명
check minor-num	파일 시스템에 대한 간단한 확인 작업을 수행합니다.
cp from to	파일 시스템을 한 파티션에서 다른 파티션으로 복사합니다; from과 to는 파티션의 minor 번호를 의미합니다.
help	사용 가능한 명령어 목록을 보여줍니다.
mklabel label	파티션 테이블에 대한 디스크 레이블을 생성합니다.
mkfs minor-num file-system-type	file-system-type 유형의 파일 시스템을 생성합니다.

명령어	설명
mkpart part-type fs-type start-mb end-mb	새로운 파일 시스템을 생성하지 않고 파티션을 만듭니다.
mkpartfs part-type fs-type start-mb end-mb	파티션을 만들고 특정 파일 시스템을 생성합니다.
move minor-num start-mb end-mb	파티션을 이동합니다.
name minor-num name	Mac과 PC98 디스크레이블 용 파티션만 이름 지정합니다
print	파티션 테이블을 보여줍니다.
quit	Quit parted
rescue start-mb end-mb	잃은 파티션의 크기를 start-mb에서 end-mb로 복구합니다.
resize minor-num start-mb end-mb	파티션의 크기를 start-mb에서 end-mb로 재조정합니다.
rm minor-num	파티션을 삭제합니다.
select device	설정할 다른 장치를 선택합니다.
set minor-num flag state	파티션 상에 플래그(flag)를 설정합니다; state는 on(켜짐) 이나 off(꺼짐) 중 하나를 입력합니다.
toggle [NUMBER [FLAG]	Toggle the state of FLAG on partition NUMBER
unit UNIT	Set the default unit to UNIT

7.1.1. 파티션 테이블 보기

After starting parted, use the command print to view the partition table. A table similar to the following appears:

```
Model: ATA ST3160812AS (scsi)
Disk /dev/sda: 160GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos

Number  Start   End    Size    Type      File system  Flags
 1      32.3kB  107MB  107MB   primary   ext3         boot
 2      107MB   105GB  105GB   primary   ext3
 3      105GB   107GB  2147MB  primary   linux-swap
 4      107GB   160GB  52.9GB  extended               root
 5      107GB   133GB  26.2GB  logical   ext3
 6      133GB   133GB  107MB   logical   ext3
 7      133GB   160GB  26.6GB  logical                lvm
```

The first line contains the disk type, manufacturer, model number and interface, and the second line displays the disk label type. The remaining output below the fourth line shows the partition table.

In the partition table, the Minor number is the partition number. For example, the partition with minor number 1 corresponds to /dev/sda1. The Start and End values are in megabytes. Valid Type are metadata, free, primary, extended, or logical. The Filesystem is the file system type, which can be any of the following:

- ext2
- ext3

- fat16
- fat32
- hfs
- jfs
- linux-swap
- ntfs
- reiserfs
- hp-ufs
- sun-ufs
- xfs

If a Filesystem of a device shows no value, this means that its file system type is unknown.

The Flags column lists the flags set for the partition. Available flags are boot, root, swap, hidden, raid, lvm, or lba.

Tip

To select a different device without having to restart parted, use the select command followed by the device name (for example, /dev/sda). Doing so allows you to view or configure the partition table of a device.

7.1.2. 파티션 생성하기

경고

사용 중인 장치에 파티션을 생성하지 마십시오.

파티션을 생성하시기 전에 복구 모드로 부팅하셔야 합니다. (또는 장치 상 모든 파티션을 마운트 해제하신 후 모든 스왑 공간을 비활성화하셔야 합니다).

Start parted, where /dev/sda is the device on which to create the partition:

```
parted /dev/sda
```

현재 파티션 테이블에 충분한 여유 공간이 있는지 확인하시기 바랍니다:

```
print
```

If there is not enough free space, you can resize an existing partition. Refer to 7.1.4절. “파티션 크기 재조정하기” for details.

7.1.2.1. 파티션 만들기

From the partition table, determine the start and end points of the new partition and what partition type it should be. You can only have four primary partitions (with no extended partition) on a device. If you need more than four partitions, you can have three primary partitions, one extended partition, and multiple logical partitions within the extended. For an overview of disk partitions, refer to the appendix An Introduction to Disk Partitions in the Red Hat Enterprise Linux Installation Guide.

예를 들어 하드 드라이브 상에 1024 메가바이트 부터 2048 메가바이트에 이르는 ext3 파일 시스템을 갖춘 1차 파티션을 생성하기 위해서는, 다음과 같은 명령을 입력하시면 됩니다:

```
mkpart primary ext3 1024 2048
```

Tip

If you use the mkpartfs command instead, the file system is created after the partition is created. However, parted does not support creating an ext3 file system. Thus, if you wish to create an ext3 file system, use mkpart and create the file system with the mkfs command as described later.

The changes start taking place as soon as you press Enter, so review the command before executing to it.

After creating the partition, use the print command to confirm that it is in the partition table with the correct partition type, file system type, and size. Also remember the minor number of the new partition so that you can label it. You should also view the output of

```
cat /proc/partitions
```

명령어의 출력 결과를 보시고 커널이 새로운 파티션을 인식하는지 여부를 확인하시기 바랍니다.

7.1.2.2. Formatting the Partition

새로 생성된 파티션은 아직 파일 시스템을 갖추고 있지 않습니다. 이제 다음 명령을 사용하여 파일 시스템을 생성하시기 바랍니다:

```
mkfs -t ext3 /dev/sda6
```

경고

파티션을 포맷하시면 현재 파티션 상에 저장된 모든 자료가 삭제될 것입니다.

7.1.2.3. 파티션 이름 붙이기 (labeling)

Next, give the partition a label. For example, if the new partition is /dev/sda6 and you want to label it /work:

```
e2label /dev/sda6 /work
```

기본 값으로 설치 프로그램은 각 파티션마다 고유한 이름을 갖도록 파티션의 마운트 지점을 이름으로 사용합니다. 하지만 여러분이 원하시는 이름으로 변경하실 수 있습니다.

7.1.2.4. 마운트 지점 생성하기

마운트 지점을 생성하기 위해서는, 루트로 로그인하신 후 다음 명령을 입력하십시오:

```
mkdir /work
```

7.1.2.5. Add to /etc/fstab

As root, edit the /etc/fstab file to include the new partition. The new line should look similar to the following:

```
LABEL=/work            /work                 ext3     defaults         1 2
```

The first column should contain LABEL= followed by the label you gave the partition. The second column should contain the mount point for the new partition, and the next column should be the file system type (for example, ext3 or swap). If you need more information about the format, read the man page with the command man fstab.

If the fourth column is the word defaults, the partition is mounted at boot time. To mount the partition without rebooting, as root, type the command:

```
mount /work
```

7.1.3. 파티션 제거하기

경고

사용 중인 장치에 위치한 파티션을 삭제하지 마십시오.

파티션을 삭제하시기 전에 복구 모드로 부팅하셔야 합니다. (또는 장치 상 모든 파티션을 마운트 해제하신 후 모든 스왑 공간을 비활성화하셔야 합니다).

Start parted, where /dev/sda is the device on which to remove the partition:

```
parted /dev/sda
```

현재 파티션 테이블에서 삭제할 파티션의 minor 번호를 확인하십시오:

```
print
```

Remove the partition with the command rm. For example, to remove the partition with minor number 3:

```
rm 3
```

The changes start taking place as soon as you press Enter, so review the command before committing to it.

After removing the partition, use the print command to confirm that it is removed from the partition table. You should also view the output of

```
cat /proc/partitions
```

명령어의 출력 결과를 보시고 커널이 해당 파티션이 삭제된 것을 인식하는지 확인해 주십시오.

The last step is to remove it from the /etc/fstab file. Find the line that declares the removed partition, and remove it from the file.

7.1.4. 파티션 크기 재조정하기

경고

사용 중인 장치에 위치한 파티션의 크기를 재조정하지 마십시오.

파티션의 크기를 재조정하시기 전에 복구 모드로 부팅하셔야 합니다. (또는 장치 상 모든 파티션을 마운트 해제하신 후 모든 스왑 공간을 비활성화하셔야 합니다).

Start parted, where /dev/sda is the device on which to resize the partition:

```
parted /dev/sda
```

현재 파티션 테이블에서 크기를 재조정할 파티션의 시작 지점과 마지막 지점을 비롯하여 크기를 재조정할 파티션의 minor 번호를 확인하십시오:

```
print
```

To resize the partition, use the resize command followed by the minor number for the partition, the starting place in megabytes, and the end place in megabytes. For example:

```
resize 3 1024 2048
```

경고

A partition cannot be made larger than the space available on the device

After resizing the partition, use the print command to confirm that the partition has been resized correctly, is the correct partition type, and is the correct file system type.

After rebooting the system into normal mode, use the command df to make sure the partition was mounted and is recognized with the new size.

7.2. LVM Partition Management

The following commands can be found by issuing lvm help at a command prompt.

표 7.2. LVM commands

명령어	설명
dumpconfig	Dump the active configuration
formats	List the available metadata formats
help	Display the help commands
lvchange	Change the attributes of logical volume(s)
lvcreate	Create a logical volume
lvdisplay	Display information about a logical volume
lvextend	Add space to a logical volume
lvmchange	Due to use of the device mapper, this command has been deprecated
lvmdiskscan	List devices that may be used as physical volumes
lvmsadc	Collect activity data
lvmsar	Create activity report
lvreduce	Reduce the size of a logical volume
lvremove	Remove logical volume(s) from the system
lvrename	Rename a logical volume
lvresize	Resize a logical volume
lvs	Display information about logical volumes
lvscan	List all logical volumes in all volume groups
pvchange	Change attributes of physical volume(s)
pvcreate	Initialize physical volume(s) for use by LVM
pvdata	Display the on-disk metadata for physical volume(s)
pvdisplay	Display various attributes of physical volume(s)

명령어	설명
pvmove	Move extents from one physical volume to another
pvremove	Remove LVM label(s) from physical volume(s)
pvresize	Resize a physical volume in use by a volume group
pvs	Display information about physical volumes
pvscan	List all physical volumes
segtypes	List available segment types
vgcfgbackup	Backup volume group configuration
vgcfgrestore	Restore volume group configuration
vgchange	Change volume group attributes
vgck	Check the consistency of a volume group
vgconvert	Change volume group metadata format
vgcreate	Create a volume group
vgdisplay	Display volume group information
vgexport	Unregister a volume group from the system
vgextend	Add physical volumes to a volume group
vgimport	Register exported volume group with system
vgmerge	Merge volume groups
vgmknodes	Create the special files for volume group devices in /dev/
vgreduce	Remove a physical volume from a volume group
vgremove	Remove a volume group
vgrename	Rename a volume group
vgs	Display information about volume groups
vgscan	Search for all volume groups
vgsplit	Move physical volumes into a new volume group
version	Display software and driver version information

디스크 사용량 할당하기

Disk space can be restricted by implementing disk quotas which alert a system administrator before a user consumes too much disk space or a partition becomes full.

Disk quotas can be configured for individual users as well as user groups. This makes it possible to manage the space allocated for user-specific files (such as email) separately from the space allocated to the projects a user works on (assuming the projects are given their own groups).

In addition, quotas can be set not just to control the number of disk blocks consumed but to control the number of inodes (data structures that contain information about files in UNIX file systems). Because inodes are used to contain file-related information, this allows control over the number of files that can be created.

The quota RPM must be installed to implement disk quotas.

> 알림
>
> For more information on installing RPM packages, refer to II부. 패키지 관리.

8.1. 디스크 사용량 제한 설정하기

디스크 사용량을 할당하시려면, 다음과 같은 과정을 따르시기 바랍니다:

1. Enable quotas per file system by modifying the /etc/fstab file.
2. Remount the file system(s).
3. Create the quota database files and generate the disk usage table.
4. Assign quota policies.

다음 부분에서는 앞에서 언급된 과정을 보다 자세하게 설명해 보겠습니다.

8.1.1. 디스크 사용량 할당 활성화하기

As root, using a text editor, edit the /etc/fstab file. Add the usrquota and/or grpquota options to the file systems that require quotas:

```
/dev/VolGroup00/LogVol00 /          ext3    defaults        1 1
LABEL=/boot              /boot      ext3    defaults        1 2
none                     /dev/pts   devpts  gid=5,mode=620  0 0
none                     /dev/shm   tmpfs   defaults        0 0
none                     /proc      proc    defaults        0 0
none                     /sys       sysfs   defaults        0 0
/dev/VolGroup00/LogVol02 /home      ext3    defaults,usrquota,grpquota  1 2
/dev/VolGroup00/LogVol01 swap       swap    defaults        0 0 . . .
```

In this example, the /home file system has both user and group quotas enabled.

알림

The following examples assume that a separate /home partition was created during the installation of Red Hat Enterprise Linux. The root (/) partition can be used for setting quota policies in the /etc/fstab file.

8.1.2. 파일 시스템 재마운트하기

After adding the usrquota and/or grpquota options, remount each file system whose fstab entry has been modified. If the file system is not in use by any process, use one of the following methods:

- Issue the umount command followed by the mount command to remount the file system.(See the man page for both umount and mount for the specific syntax for mounting and unmounting various filesystem types.)
- Issue the mount -o remount <file-system> command (where <file-system> is the name of the file system) to remount the file system. For example, to remount the /home file system, the command to issue is mount -o remount /home.

If the file system is currently in use, the easiest method for remounting the file system is to reboot the system.

8.1.3. Creating the Quota Database Files

After each quota-enabled file system is remounted, the system is capable of working with disk quotas. However, the file system itself is not yet ready to support quotas. The next step is to run the quotacheck command.

The quotacheck command examines quota-enabled file systems and builds a table of the current disk usage per file system. The table is then used to update the operating system's copy of disk usage. In addition, the file system's disk quota files are updated.

To create the quota files (aquota.user and aquota.group) on the file system, use the -c option of the quotacheck command. For example, if user and group quotas are enabled for the /home file system, create the files in the /home directory:

```
quotacheck -cug /home
```

The -c option specifies that the quota files should be created for each file system with quotas enabled, the -u option specifies to check for user quotas, and the -g option specifies to check for group quotas.

만일 -u 옵션과 -g 옵션이 지정되지 않았다면, 사용자 사용량 할당 파일만 생성됩니다. -g 옵션만 지정된 경우에는, 그룹 사용량 할당 파일만 만들어 집니다.

파일이 생성된 후, 다음과 같은 명령을 실행하여 사용량 할당이 활성화된 파일 시스템마다 현재 디스크 사용량을 보여주는 표를 생성하시기 바랍니다:

```
quotacheck -avug
```

이 명령에서 사용된 옵션들은 다음과 같습니다:

- a — 디스크 사용량 할당이 활성화되고, 로컬에서 마운트된 모든 파일 시스템을 확인합니다.
- v — 사용량 할당 확인 작업이 진행 과정을 상세한 상태 정보로 보여줍니다.
- u — 사용자 디스크 사용량 할당 정보를 체크합니다.
- g — 그룹 디스크 사용량 할당 정보를 체크합니다.

After quotacheck has finished running, the quota files corresponding to the enabled quotas (user and/or group) are populated with data for each quota-enabled locally-mounted file system such as /home.

8.1.4. 사용자 당 디스크 사용량 할당하기

The last step is assigning the disk quotas with the edquota command.

개인 사용자 당 사용량 할당을 설정하시려면, 쉘 프롬프트에서 루트로 로그인 하신 후 다음과 같은 명령을 입력하십시오:

```
edquota username
```

Perform this step for each user who needs a quota. For example, if a quota is enabled in /etc/fstab for the /home partition (/dev/VolGroup00/LogVol02 in the example below) and the command edquota testuser is executed, the following is shown in the editor configured as the default for the system:

```
Disk quotas for user testuser (uid 501):
Filesystem                blocks     soft     hard     inodes    soft     hard
/dev/VolGroup00/LogVol02  440436        0        0      37418       0        0
```

알림

The text editor defined by the EDITOR environment variable is used by edquota. To change the editor, set the EDITOR environment variable in your ~/.bash_profile file to the full path of the editor of your choice.

The first column is the name of the file system that has a quota enabled for it. The second column shows how many blocks the user is currently using. The next two columns are used to set soft and hard block limits for the user on the file system. The inodes column shows how many inodes the user is currently using. The last two columns are used to set the soft and hard inode limits for the user on the file system.

The hard block limit is the absolute maximum amount of disk space that a user or group can use. Once this limit is reached, no further disk space can be used.

The soft block limit defines the maximum amount of disk space that can be used. However, unlike the hard limit, the soft limit can be exceeded for a certain amount of time. That time is known as the grace period. The grace period can be expressed in seconds, minutes, hours, days, weeks, or months.

0라고 설정된 값이 있다면, 제한이 설정되지 않은 것입니다. 텍스트 편집기를 사용하여 원하시는 제한 값으로 변경하시기 바랍니다. 예로 들면:

```
Disk quotas for user testuser (uid 501):
Filesystem                blocks     soft      hard    inodes   soft    hard
/dev/VolGroup00/LogVol02  440436     500000    550000   37418      0       0
```

해당 사용자에 대한 디스크 사용량이 할당되었는지 확인해 보시려면, 다음 명령을 입력하십시오:

```
quota testuser
```

8.1.5. 그룹 당 디스크 사용량 할당하기

Quotas can also be assigned on a per-group basis. For example, to set a group quota for the devel group (the group must exist prior to setting the group quota), use the command:

```
edquota -g devel
```

이 명령을 입력하시면 텍스트 편집기에서 해당 그룹에 대한 기존 디스크 사용량이 나타납니다:

```
Disk quotas for group devel (gid 505):
Filesystem                blocks    soft      hard     inodes    soft     hard
/dev/VolGroup00/LogVol02  440400       0         0      37418       0        0
```

Modify the limits, then save the file.

그룹 디스크 사용량이 설정되었는지 여부를 확인해 보시려면, 다음 명령을 사용하시기 바랍니다:

```
quota -g devel
```

8.1.6. Setting the Grace Period for Soft Limits

If soft limits are set for a given quota (whether inode or block and for either users or groups) the grace period, or amount of time a soft limit can be exceeded, should be set with the command:

```
edquota -t
```

While other edquota commands operate on a particular user's or group's quota, the -t option operates on every filesystem with quotas enabled.

8.2. 디스크 사용량 할당 관리하기

If quotas are implemented, they need some maintenance — mostly in the form of watching to see if the quotas are exceeded and making sure the quotas are accurate.

Of course, if users repeatedly exceed their quotas or consistently reach their soft limits, a system administrator has a few choices to make depending on what type of users they are and how much disk space impacts their work. The administrator can either help the user determine how to use less disk space or increase the user's disk quota.

8.2.1. 활성화와 비활성화

It is possible to disable quotas without setting them to 0. To turn all user and group quotas off, use the following command:

```
quotaoff -vaug
```

If neither the -u or -g options are specified, only the user quotas are disabled. If only -g is specified, only group quotas are disabled. The -v switch causes verbose status information to display as the command executes.

To enable quotas again, use the quotaon command with the same options.

For example, to enable user and group quotas for all file systems, use the following command:

```
quotaon -vaug
```

To enable quotas for a specific file system, such as /home, use the following command:

```
quotaon -vug /home
```

만일 -u 옵션이나 -g 옵션이 지정되지 않는다면, 사용자 디스크 사용량 할당만이 활성화 됩니다. -g 옵션만 지정된 경우에는 그룹 디스크 사용량 할당만이 활성화 됩니다.

8.2.2. 디스크 사용량 보고하기

Creating a disk usage report entails running the repquota utility. For example, the command repquota /home produces this output:

```
*** Report for user quotas on device /dev/mapper/VolGroup00-LogVol02
Block grace time: 7days; Inode grace time: 7days
                        Block limits                File limits
User            used    soft    hard  grace    used  soft  hard  grace
----------------------------------------------------------------------
root      --      36       0       0              4     0     0
kristin   --     540       0       0            125     0     0
testuser  --  440400  500000  550000          37418     0     0
```

To view the disk usage report for all (option -a) quota-enabled file systems, use the command:

```
repquota -a
```

While the report is easy to read, a few points should be explained. The -- displayed after each user is a quick way to determine whether the block or inode limits have been exceeded. If either soft limit is exceeded, a + appears in place of the corresponding -; the first - represents the block limit, and the second represents the inode limit.

The grace columns are normally blank. If a soft limit has been exceeded, the column contains a time specification equal to the amount of time remaining on the grace period. If the grace period has expired, none appears in its place.

8.2.3. 정확한 디스크 할당 사용량 지키기

Whenever a file system is not unmounted cleanly (due to a system crash, for example), it is necessary to run quotacheck. However, quotacheck can be run on a regular basis, even if the system has not crashed. Safe methods for periodically running quotacheck include:

Ensuring quotacheck runs on next reboot

> **Best method for most systems**
>
> This method works best for (busy) multiuser systems which are periodically rebooted.

As root, place a shell script into the /etc/cron.daily/ or /etc/cron.weekly/ directory—or schedule one using the crontab -e command—that contains the touch /forcequotacheck command. This creates an empty forcequotacheck file in the root directory, which the system init script looks for at boot time. If it is found, the init script runs quotacheck. Afterward, the init script removes the /forcequotacheck file; thus, scheduling this file to be created periodically with cron ensures that quotacheck is run during the next reboot.

Refer to 37장. Automated Tasks for more information about configuring cron.

Running quotacheck in single user mode

An alternative way to safely run quotacheck is to (re-)boot the system into single-user mode to prevent the possibility of data corruption in quota files and run:

```
~]# quotaoff -vaug /<file_system>
~]# quotacheck -vaug /<file_system>
~]# quotaon -vaug /<file_system>
```

Running quotacheck on a running system

If necessary, it is possible to run quotacheck on a machine during a time when no users are logged in, and thus have no open files on the file system being checked. Run the command quotacheck -vaug <file_system> ; this command will fail if quotacheck cannot remount the given <file_system> as read-only. Note that, following the check, the file system will be remounted read-write.

> **Do not run quotacheck on a live file system**
>
> Running quotacheck on a live file system mounted read-write is not recommended due to the possibility of quota file corruption.

Refer to 37장. Automated Tasks for more information about configuring cron.

8.3. 추가 자료

디스크 사용량 할당에 대한 보다 많은 정보를 원하신다면, 다음 자료를 참조하시기 바랍니다.

8.3.1. 설치된 문서 자료

- The quotacheck, edquota, repquota, quota, quotaon, and quotaoff man pages

8.3.2. 관련 서적

- Red Hat Enterprise Linux Introduction to System Administration; Red Hat, Inc. — Available at http://www.redhat.com/docs/ and on the Documentation CD, this manual contains background information on storage management (including disk quotas) for new Red Hat Enterprise Linux system administrators.

Access Control Lists

Files and directories have permission sets for the owner of the file, the group associated with the file, and all other users for the system. However, these permission sets have limitations. For example, different permissions cannot be configured for different users. Thus, Access Control Lists (ACLs) were implemented.

The Red Hat Enterprise Linux 5 kernel provides ACL support for the ext3 file system and NFS-exported file systems. ACLs are also recognized on ext3 file systems accessed via Samba.

Along with support in the kernel, the acl package is required to implement ACLs. It contains the utilities used to add, modify, remove, and retrieve ACL information.

The cp and mv commands copy or move any ACLs associated with files and directories.

9.1. 파일 시스템 마운트하기

파일이나 디렉토리에 ACL을 사용하기 전에, 그 파일과 디렉토리의 파티션은 ACL 지원을 사용하여 마운트되어야 합니다. 만일 지역 ext3 파일 시스템이라면, 다음 명령을 사용하여 마운트할 수 있습니다:

```
mount -t ext3 -o acl <device-name> <partition>
```

예를 들면:

```
mount -t ext3 -o acl /dev/VolGroup00/LogVol02 /work
```

Alternatively, if the partition is listed in the /etc/fstab file, the entry for the partition can include the acl option:

```
LABEL=/work      /work       ext3     acl        1 2
```

만일 Samba를 통하여 ext3 파일 시스템에 접속한 경우 이 파일 시스템에 ACL이 활성화되어 있다면, Samba가 --with-acl-support 옵션을 사용하여 컴파일되었기 때문에 ACL을 사용 가능합니다. Samba 공유를 사용하거나 마운팅할시 특별한 플래그(flag)를 사용할 필요가 없습니다.

9.1.1. NFS

기본 값으로 NFS 서버에 의해 export된 파일 시스템은 ACL을 지원하며 NFS 클라이언트는 ACL을 읽고 사용할 수 있습니다.

To disable ACLs on NFS shares when configuring the server, include the no_acl option in the /etc/exports file. To disable ACLs on an NFS share when mounting it on a client, mount it with the no_acl option via the command line or the /etc/fstab file.

9.2. Access ACL 설정하기

두가지 종류의 ACL이 있습니다: access ACLs과 기본 ACLs. access ACL은 특정 파일이나 디렉토리에 사용되는 접근 제어 목록을 말합니다. 기본 ACL은 오직 디렉토리에만 사용됩니다; 만일 디렉

토리 내의 한 파일에 access ACL이 없다면, 이 파일은 디렉토리의 기본 ACL 규칙을 사용합니다. 기본 ACL은 옵션입니다.

ACL은 다음과 같이 설정 가능합니다:

1. 사용자 당 설정

2. 그룹 당 설정

3. 유효한 접근 권리 마스크를 사용한 설정

4. 그 파일의 사용자 그룹에 속하지 않는 사용자들에 대한 설정

The setfacl utility sets ACLs for files and directories. Use the -m option to add or modify the ACL of a file or directory:

```
setfacl -m <rules> <files>
```

Rules (<rules>) must be specified in the following formats. Multiple rules can be specified in the same command if they are separated by commas.

u:<uid>:<perms>
사용자에 대한 access ACL을 설정합니다. 사용자명이나 UID를 지정하실 수 있습니다. 사용자는 시스템 상 어느 사용자라도 가능합니다.

g:<gid>:<perms>
그룹에 대한 access ACL을 설정합니다. 그룹명이나 GID를 지정할 수 있습니다. 그룹은 시스템 상 어느 그룹이라도 가능합니다.

m:<perms>
유효한 접근 권한 마스크를 설정합니다. 이 마스크는 소유 그룹의 허가와 모든 사용자와 그룹 항목을 조합한 것입니다.

o:<perms>
파일의 그룹에 속한 사용자가 아닌 다른 사용자에 대한 access ACL을 설정합니다.

White space is ignored. Permissions (<perms>) must be a combination of the characters r, w, and x for read, write, and execute.

If a file or directory already has an ACL, and the setfacl command is used, the additional rules are added to the existing ACL or the existing rule is modified.

For example, to give read and write permissions to user andrius:

```
setfacl -m u:andrius:rw /project/somefile
```

To remove all the permissions for a user, group, or others, use the -x option and do not specify any permissions:

```
setfacl -x <rules> <files>
```

예를 들어 UID 500인 사용자의 모든 권한을 삭제하시려면:

```
setfacl -x u:500 /project/somefile
```

9.3. 기본 ACL 설정하기

To set a default ACL, add d: before the rule and specify a directory instead of a file name.

For example, to set the default ACL for the /share/ directory to read and execute for users not in the user group (an access ACL for an individual file can override it):

```
setfacl -m d:o:rx /share
```

9.4. ACL 보기

To determine the existing ACLs for a file or directory, use the getfacl command. In the example below, the getfacl is used to determine the existing ACLs for a file.

```
getfacl home/john/picture.png
```

The above command returns the following output:

```
# file: home/john/picture.png
# owner: john
# group: john
user::rw-
group::r--
other::r--
```

If a directory with a default ACL is specified, the default ACL is also displayed as illustrated below.

```
[john@main /]$ getfacl home/sales/
# file: home/sales/
# owner: john
# group: john
user::rw-
user:barryg:r--
group::r--
mask::r--
other::r--
default:user::rwx
default:user:john:rwx
default:group::r-x
default:mask::rwx
default:other::r-x
```

9.5. ACL을 가진 파일 시스템 압축 저장하기

경고

The tar and dump commands do not backup ACLs.

The star utility is similar to the tar utility in that it can be used to generate archives of files; however, some of its options are different. Refer to 표 9.1. "Command Line Options for star" for a listing of more commonly used options. For all available options, refer to the star man page. The star package is required to use this utility.

표 9.1. Command Line Options for star

옵션	설명
-c	아카이브 파일을 생성합니다.
-n	파일을 압축 해제하지 않습니다; -x 옵션과 함께 사용하여 파일을 압축 해제할 수 있습니다.
-r	아카이브에서 파일을 대체합니다. 파일들은 동일한 경로와 파일명을 가진 파일들을 대체하여 아카이브 파일 마지막에 기록됩니다.
-t	아카이브 파일의 내용을 보여줍니다.
-u	아카이브 파일을 업데이트합니다. 파일들이 아카이브에 존재하지 아허나 파일들이 아카이브에 있는 동일한 이름의 파일 보다 최신의 것을 경우, 아카이브 마지막에 기록됩니다. 이 옵션은 아카이브가 파일이거나 특별 테이프 (블록 크기가 정해지지 않은 테이프)인 경우에만 작동합니다.
-x	아카이브에서 파일을 압축 해제합니다. -U 옵션과 함께 사용되면 파일 시스템에 존재하는 상응하는 파일 보다 아카이브에 저장된 파일이 이전 것일 경우, 그 파일은 압축 해제되지 않습니다.
-help	대부분의 중요한 옵션을 보여줍니다.
-xhelp	자주 사용되지 않는 옵션들을 보여줍니다.
-/	아카이브에서 파일을 압축 해제할 때 파일명 앞에 위치한 슬래쉬를 제거하지 마십시오. 기본 값으로, 파일이 압축 해제될 때 슬래쉬가 제거됩니다.
-acl	파일과 디렉토리와 연관된 ACL을 생성, 압축 해제, 압축하거나 복구할 때 사용됩니다.

9.6. 이전 시스템과의 호환성

If an ACL has been set on any file on a given file system, that file system has the ext_attr attribute. This attribute can be seen using the following command:

```
tune2fs -l <filesystem-device>
```

A file system that has acquired the ext_attr attribute can be mounted with older kernels, but those kernels do not enforce any ACLs which have been set.

Versions of the e2fsck utility included in version 1.22 and higher of the e2fsprogs package (including the versions in Red Hat Enterprise Linux 2.1 and 4) can check a file system with the ext_attr attribute. Older versions refuse to check it.

9.7. 추가 자료

보다 자세한 정보를 원하신다면, 다음 자료들을 참조하시기 바랍니다.

9.7.1. 설치된 문서 자료

- acl man page — Description of ACLs
- getfacl man page — Discusses how to get file access control lists
- setfacl man page — Explains how to set file access control lists
- star man page — Explains more about the star utility and its many options

9.7.2. 유용한 웹사이트

- http://acl.bestbits.at/ — ACL 웹사이트

LVM (Logical Volume Manager)

10.1. What is LVM?

LVM is a tool for logical volume management which includes allocating disks, striping, mirroring and resizing logical volumes.

With LVM, a hard drive or set of hard drives is allocated to one or more physical volumes. LVM physical volumes can be placed on other block devices which might span two or more disks.

The physical volumes are combined into logical volumes, with the exception of the /boot partition. The /boot partition cannot be on a logical volume group because the boot loader cannot read it. If the root (/) partition is on a logical volume, create a separate /boot partition which is not a part of a volume group.

Since a physical volume cannot span over multiple drives, to span over more than one drive, create one or more physical volumes per drive.

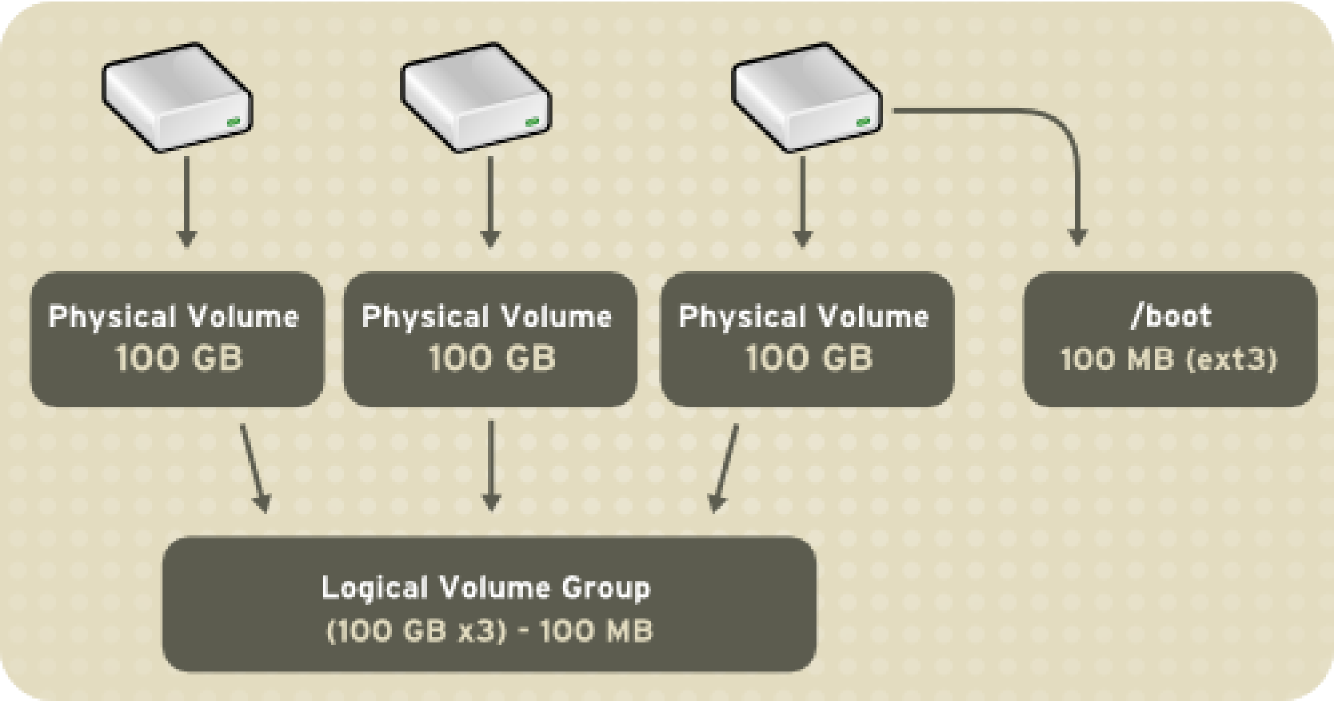

그림 10.1. Logical Volumes

The volume groups can be divided into logical volumes, which are assigned mount points, such as /home and / and file system types, such as ext2 or ext3. When "partitions" reach their full capacity, free space from the volume group can be added to the logical volume to increase the size of the partition. When a new hard drive is added to the system, it can be added to the volume group, and partitions that are logical volumes can be increased in size.

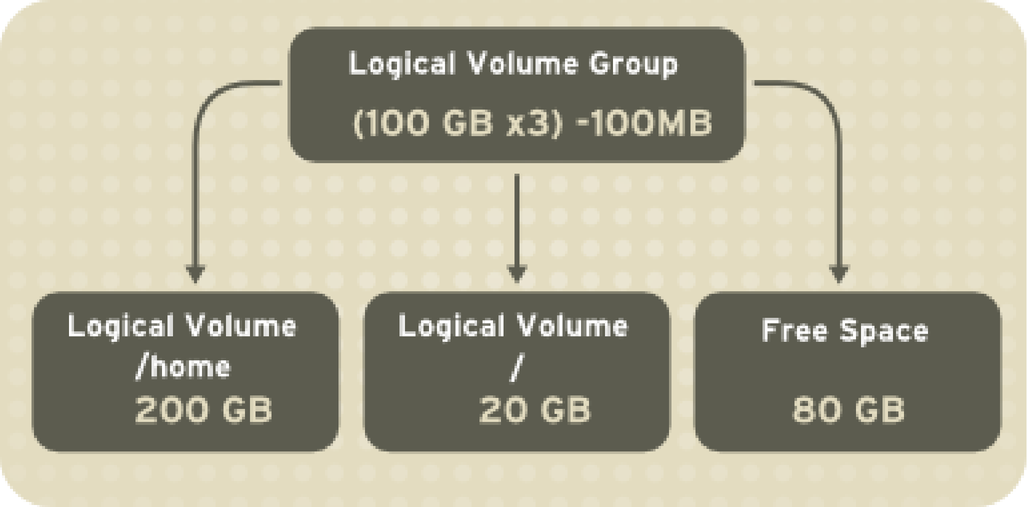

그림 10.2. Logical Volumes

On the other hand, if a system is partitioned with the ext3 file system, the hard drive is divided into partitions of defined sizes. If a partition becomes full, it is not easy to expand the size of the partition. Even if the partition is moved to another hard drive, the original hard drive space has to be reallocated as a different partition or not used.

To learn how to configure LVM during the installation process, refer to 10.2절. "LVM 설정" .

10.1.1. What is LVM2?

LVM version 2, or LVM2, is the default for Red Hat Enterprise Linux 5, which uses the device mapper driver contained in the 2.6 kernel. LVM2 can be upgraded from versions of Red Hat Enterprise Linux running the 2.4 kernel.

10.2. LVM 설정

LVM can be configured during the graphical installation process, the text-based installation process, or during a kickstart installation. You can use the system-config-lvm utility to create your own LVM configuration post-installation. The next two sections focus on using Disk Druid during installation to complete this task. The third section introduces the LVM utility (system-config-lvm) which allows you to manage your LVM volumes in X windows or graphically.

Read 10.1절. "What is LVM?" first to learn about LVM. An overview of the steps required to configure LVM include:

- Creating physical volumes from the hard drives.
- Creating volume groups from the physical volumes.
- Creating logical volumes from the volume groups and assign the logical volumes mount points.

Two 9.1 GB SCSI drives (/dev/sda and /dev/sdb) are used in the following examples. They detail how to create a simple configuration using a single LVM volume group with associated logical volumes during installation.

10.3. Automatic Partitioning

On the Disk Partitioning Setup screen, select Remove linux partitions on selected drives and create default layout from the pulldown list.

For Red Hat Enterprise Linux, LVM is the default method for disk partitioning. If you do not wish to have LVM implemented, or if you require RAID partitioning, manual disk partitioning through Disk Druid is required.

The following properties make up the automatically created configuration:

- The /boot partition resides on its own non-LVM partition. In the following example, it is the first partition on the first drive (/dev/sda1). Bootable partitions cannot reside on LVM logical volumes.
- A single LVM volume group (VolGroup00) is created, which spans all selected drives and all remaining space available. In the following example, the remainder of the first drive (/dev/sda2), and the entire second drive (/dev/sdb1) are allocated to the volume group.
- Two LVM logical volumes (LogVol00 and LogVol01) are created from the newly created spanned volume group. In the following example, the recommended swap space is automatically calculated and assigned to LogVol01, and the remainder is allocated to the root file system, LogVol00.

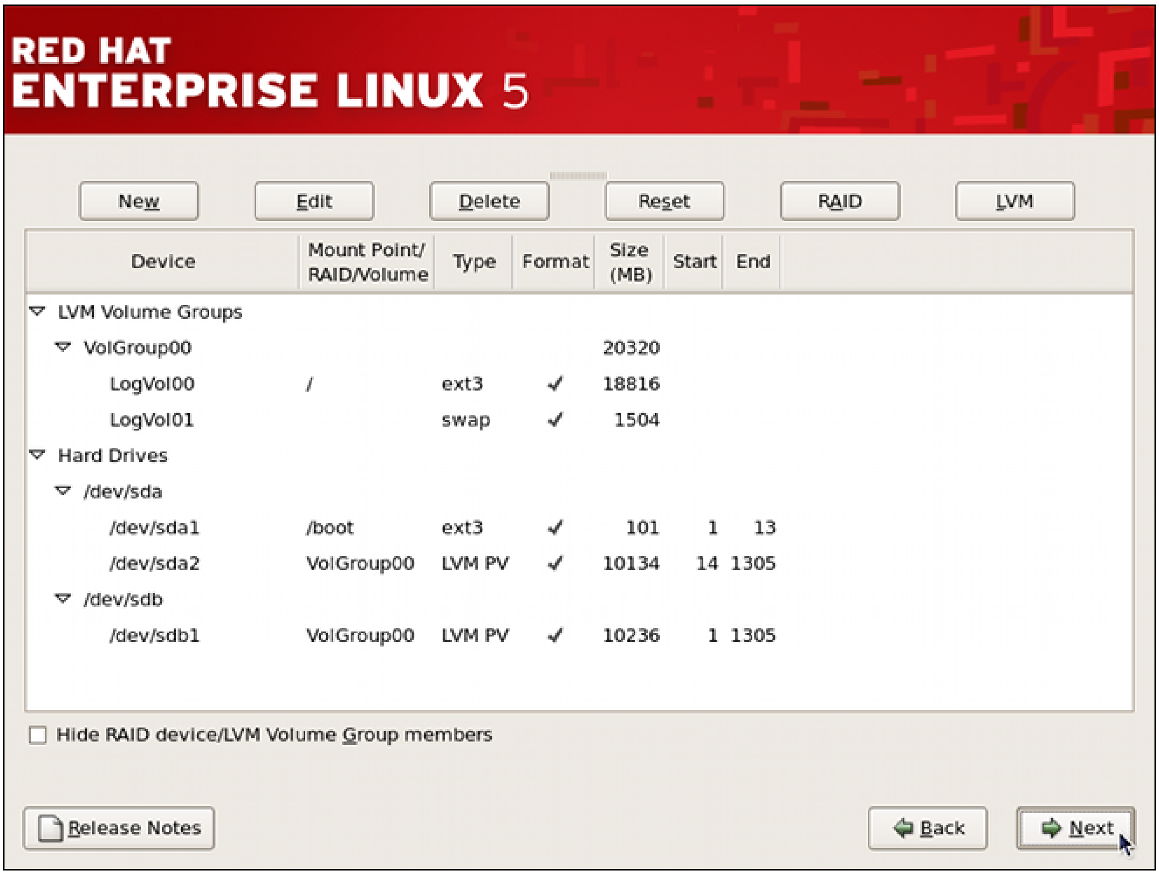

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
LVM Volume Groups						
VolGroup00				20320		
LogVol00	/	ext3	✓	18816		
LogVol01		swap	✓	1504		
Hard Drives						
/dev/sda						
/dev/sda1	/boot	ext3	✓	101	1	13
/dev/sda2	VolGroup00	LVM PV	✓	10134	14	1305
/dev/sdb						
/dev/sdb1	VolGroup00	LVM PV	✓	10236	1	1305

□ Hide RAID device/LVM Volume Group members

그림 10.3. Automatic LVM Configuration With Two SCSI Drives

알림

If enabling quotas are of interest to you, it may be best to modify the automatic configuration to include other mount points, such as /home or /var, so that each file system has its own independent quota configuration limits.

In most cases, the default automatic LVM partitioning is sufficient, but advanced implementations could warrant modification or manual configuration of the partition tables.

알림

If you anticipate future memory upgrades, leaving some free space in the volume group would allow for easy future expansion of the swap space logical volume on the system; in which case, the automatic LVM configuration should be modified to leave available space for future growth.

10.4. Manual LVM Partitioning

The following section explains how to manually configure LVM for Red Hat Enterprise Linux. Because there are numerous ways to manually configure a system with LVM, the following example is similar to the default configuration done in 10.3절. “Automatic Partitioning” .

On the Disk Partitioning Setup screen, select Create custom layout from the pulldown list and click the Next button in the bottom right corner of the screen.

10.4.1. Creating the /boot Partition

In a typical situation, the disk drives are new, or formatted clean. The following figure, 그림 10.4. “Two Blank Drives, Ready for Configuration” , shows both drives as raw devices with no partitioning configured.

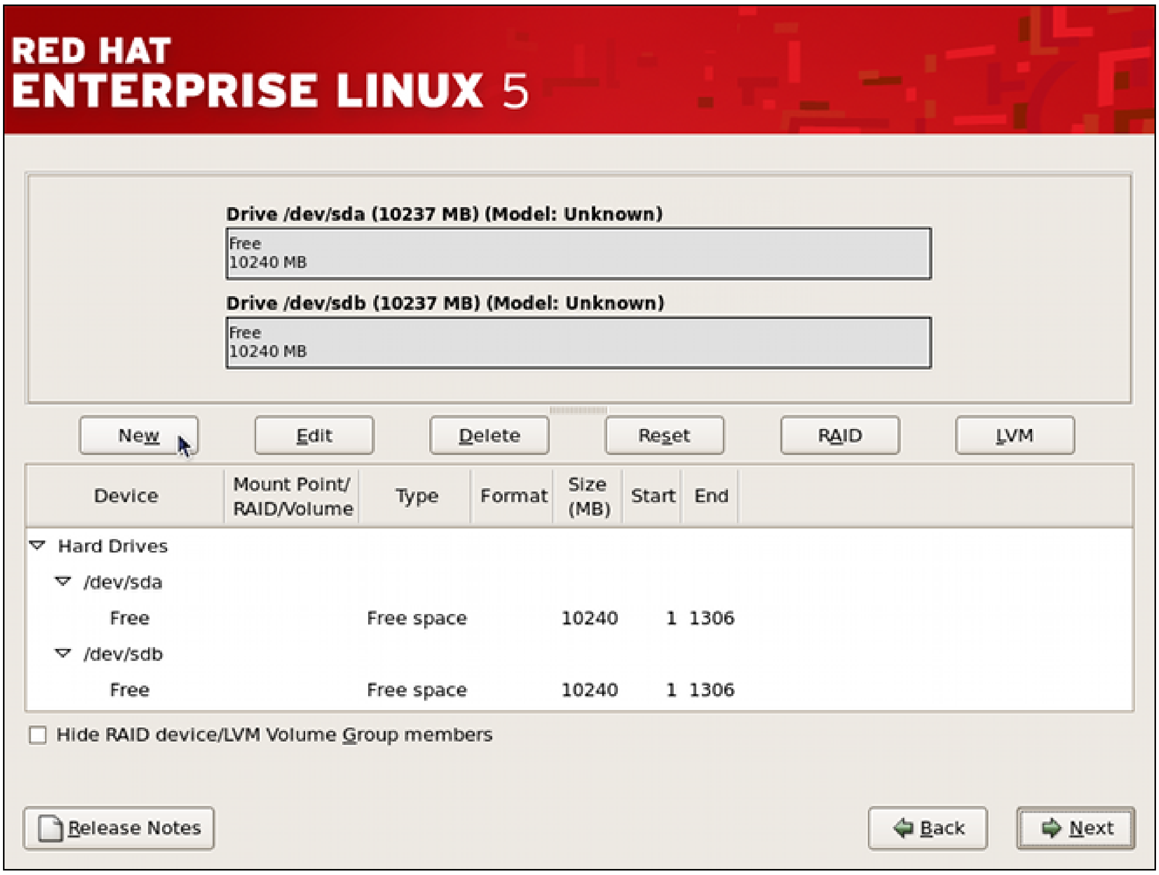

Drive /dev/sda (10237 MB) (Model: Unknown)

Free
10240 MB

Drive /dev/sdb (10237 MB) (Model: Unknown)

Free
10240 MB

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
▽ Hard Drives						
▽ /dev/sda						
Free		Free space		10240	1	1306
▽ /dev/sdb						
Free		Free space		10240	1	1306

□ Hide RAID device/LVM Volume Group members

그림 10.4. Two Blank Drives, Ready for Configuration

경고

The /boot partition cannot reside on an LVM volume because the GRUB boot loader cannot read it.

1. Select New.
2. Select /boot from the Mount Point pulldown menu.
3. Select ext3 from the File System Type pulldown menu.
4. Select only the sda checkbox from the Allowable Drives area.
5. Leave 100 (the default) in the Size (MB) menu.
6. Leave the Fixed size (the default) radio button selected in the Additional Size Options area.
7. Select Force to be a primary partition to make the partition be a primary partition. A primary partition is one of the first four partitions on the hard drive. If unselected, the partition is created as a logical partition. If other operating systems are already on the system, unselecting this option should be considered. For more information on primary versus logical/extended partitions, refer to the appendix section of the Red Hat Enterprise Linux Installation Guide.

Refer to 그림 10.5. "Creation of the Boot Partition" to verify your inputted values:

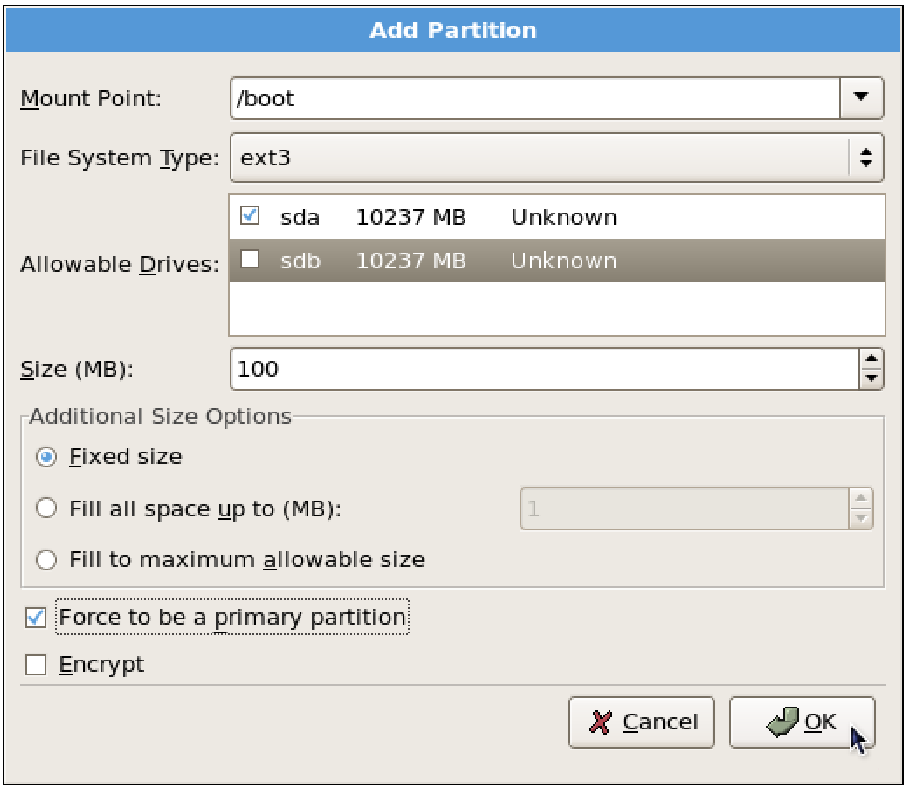

Mount Point: /boot

File System Type: ext3

Allowable Drives:

☑	sda	10237 MB	Unknown
□	sdb	10237 MB	Unknown

Size (MB): 100

Additional Size Options

● Fixed size

○ Fill all space up to (MB): 1

○ Fill to maximum allowable size

☑ Force to be a primary partition

□ Encrypt

그림 10.5. Creation of the Boot Partition

Click OK to return to the main screen. The following figure displays the boot partition correctly set:

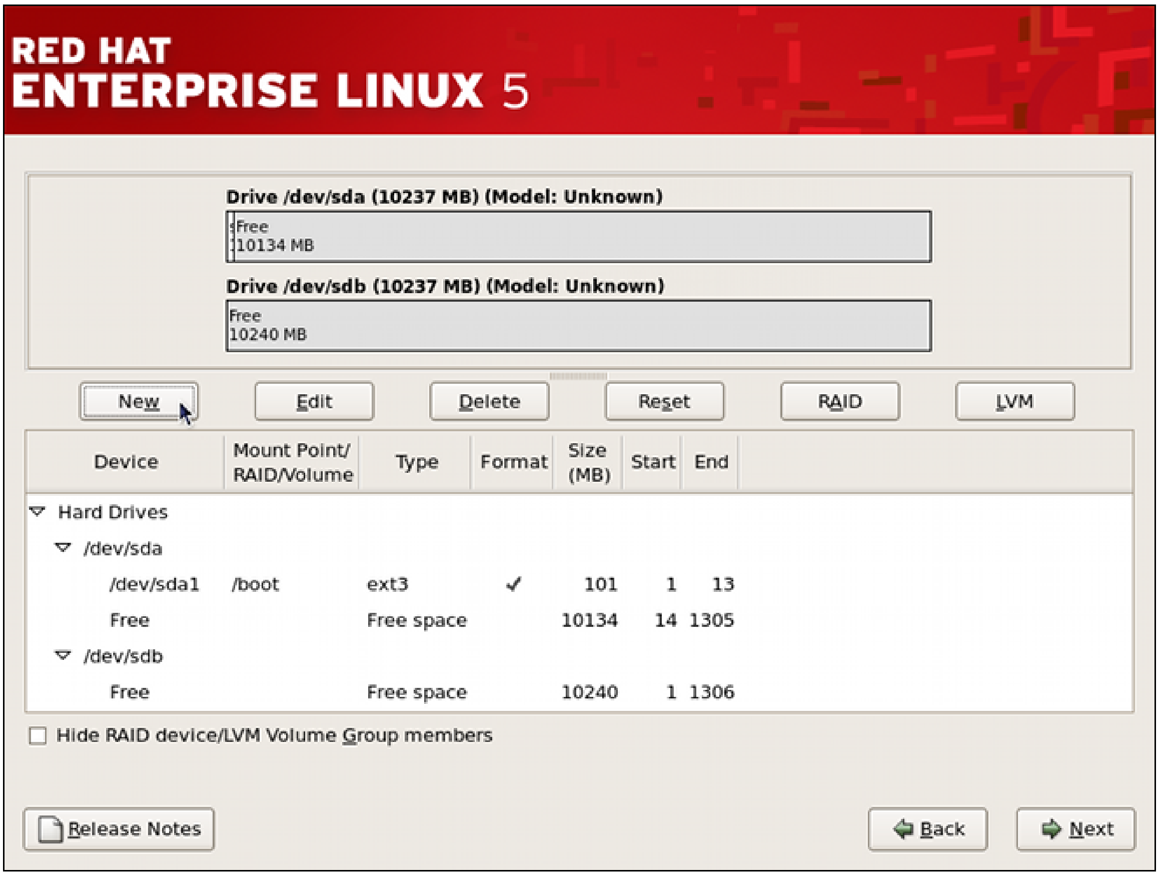

Drive /dev/sda (10237 MB) (Model: Unknown)

Free
10134 MB

Drive /dev/sdb (10237 MB) (Model: Unknown)

Free
10240 MB

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
▽ Hard Drives						
▽ /dev/sda						
/dev/sda1	/boot	ext3	✓	101	1	13
Free		Free space		10134	14	1305
▽ /dev/sdb						
Free		Free space		10240	1	1306

□ Hide RAID device/LVM Volume Group members

그림 10.6. The /boot Partition Displayed

10.4.2. Creating the LVM Physical Volumes

Once the boot partition is created, the remainder of all disk space can be allocated to LVM partitions. The first step in creating a successful LVM implementation is the creation of the physical volume(s).

1. Select New.

2. Select physical volume (LVM) from the File System Type pulldown menu as shown in 그림 10.7. "Creating a Physical Volume" .

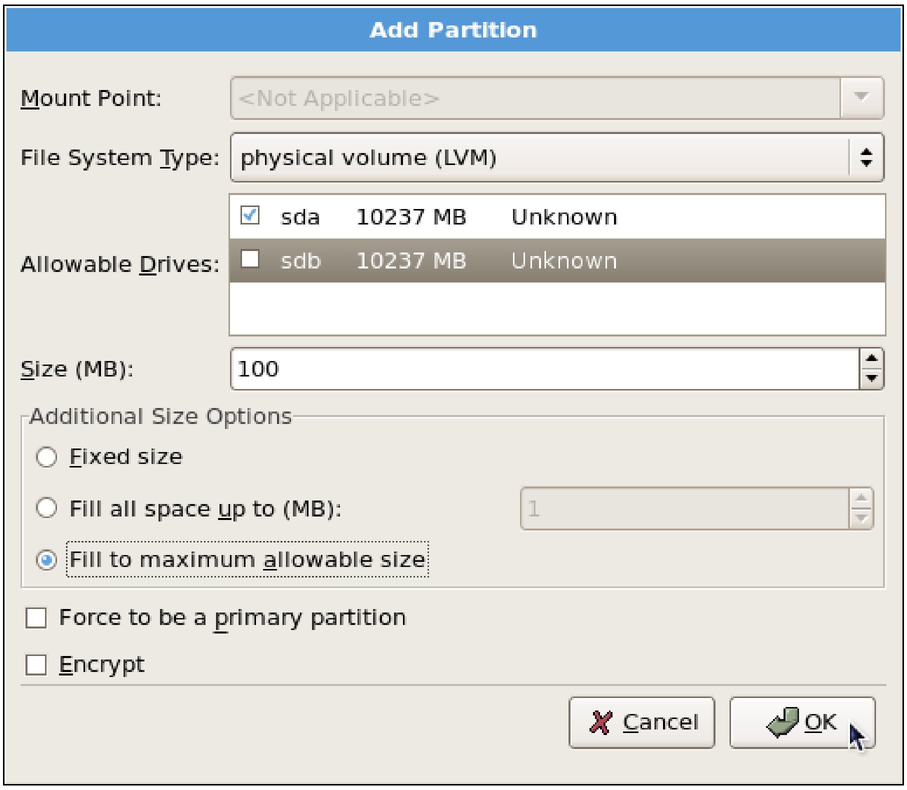

Add Partition

Mount Point: <Not Applicable>

File System Type: physical volume (LVM)

Allowable Drives:
☑ sda 10237 MB Unknown
□ sdb 10237 MB Unknown

Size (MB): 100

Additional Size Options
○ Fixed size
○ Fill all space up to (MB): 1
● Fill to maximum allowable size

□ Force to be a primary partition

□ Encrypt

그림 10.7. Creating a Physical Volume

3. You cannot enter a mount point yet (you can once you have created all your physical volumes and then all volume groups).

4. A physical volume must be constrained to one drive. For Allowable Drives, select the drive on which the physical volume are created. If you have multiple drives, all drives are selected, and you must deselect all but one drive.

5. 원하시는 물리적 볼륨의 크기를 입력하십시오.

6. Select Fixed size to make the physical volume the specified size, select Fill all space up to (MB) and enter a size in MBs to give range for the physical volume size, or select Fill to maximum allowable size to make it grow to fill all available space on the hard disk. If you make more than one growable, they share the available free space on the disk.

7. Select Force to be a primary partition if you want the partition to be a primary partition.

8. Click OK to return to the main screen.

Repeat these steps to create as many physical volumes as needed for your LVM setup. For example, if you want the volume group to span over more than one drive, create a physical volume on each of the drives. The following figure shows both drives completed after the repeated process:

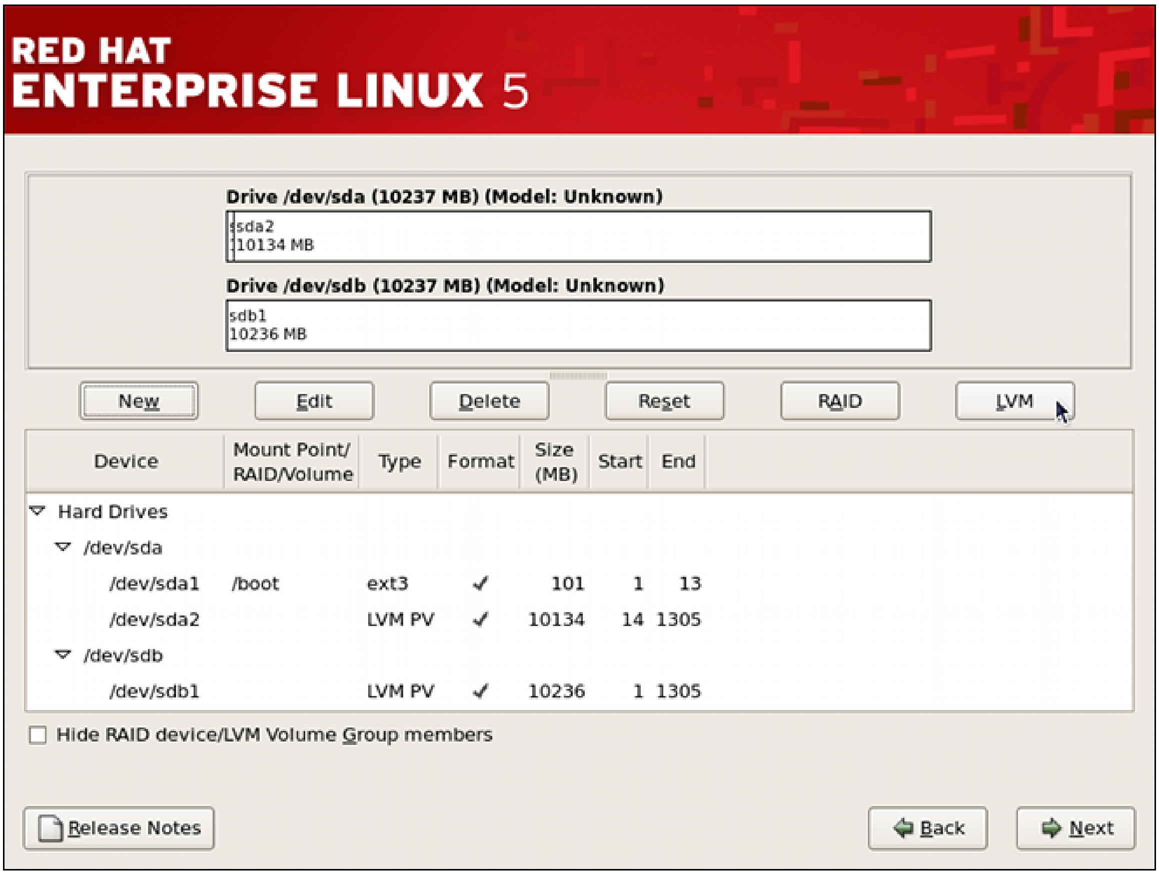

Drive /dev/sda (10237 MB) (Model: Unknown)

sda2
10134 MB

Drive /dev/sdb (10237 MB) (Model: Unknown)

sdb1
10236 MB

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
▽ Hard Drives						
▽ /dev/sda						
/dev/sda1	/boot	ext3	✓	101	1	13
/dev/sda2		LVM PV	✓	10134	14	1305
▽ /dev/sdb						
/dev/sdb1		LVM PV	✓	10236	1	1305

□ Hide RAID device/LVM Volume Group members

그림 10.8. Two Physical Volumes Created

10.4.3. Creating the LVM Volume Groups

Once all the physical volumes are created, the volume groups can be created:

1. Click the LVM button to collect the physical volumes into volume groups. A volume group is basically a collection of physical volumes. You can have multiple logical volumes, but a physical volume can only be in one volume group.

> 알림
>
> There is overhead disk space reserved in the volume group. The volume group size is slightly less than the total of physical volume sizes.

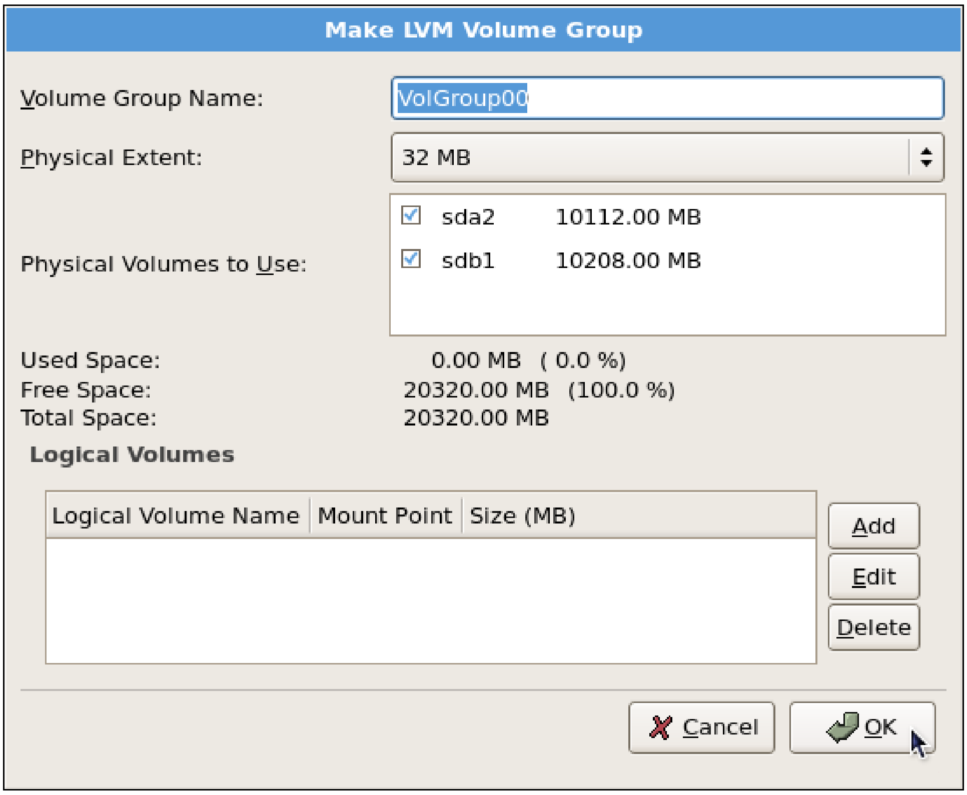

Make LVM Volume Group

Volume Group Name: VolGroup00

Physical Extent: 32 MB

Physical Volumes to Use:

☑	sda2	10112.00 MB
☑	sdb1	10208.00 MB

Used Space:	0.00 MB (0.0 %)
Free Space:	20320.00 MB (100.0 %)
Total Space:	20320.00 MB

Logical Volumes

Logical Volume Name	Mount Point	Size (MB)

그림 10.9. Creating an LVM Volume Group

2. Change the Volume Group Name if desired.

3. All logical volumes inside the volume group must be allocated in physical extent (PE) units. A physical extent is an allocation unit for data.

4. 볼륨 그룹에 사용할 물리적 볼륨을 선택해 주십시오.

10.4.4. Creating the LVM Logical Volumes

Create logical volumes with mount points such as /, /home, and swap space. Remember that /boot cannot be a logical volume. To add a logical volume, click the Add button in the Logical Volumes section. A dialog window as shown in 그림 10.10. "Creating a Logical Volume" appears.

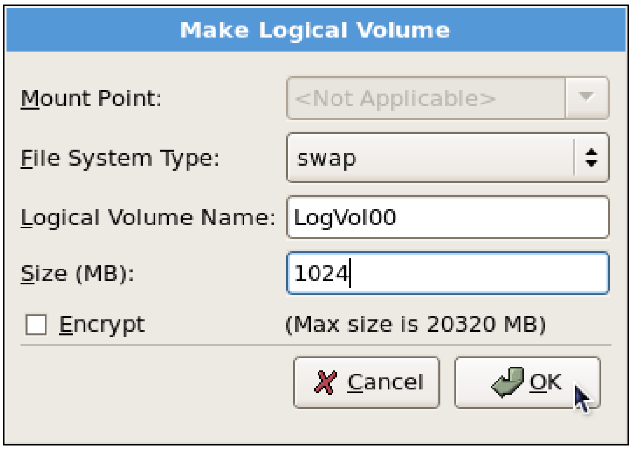

Make Logical Volume

Mount Point: <Not Applicable>

File System Type: swap

Logical Volume Name: LogVol00

Size (MB): 1024

□ Encrypt (Max size is 20320 MB)

그림 10.10. Creating a Logical Volume

생성하시려는 볼륨 그룹 마다 이 과정을 반복하십시오.

Tip

You may want to leave some free space in the volume group so you can expand the logical volumes later. The default automatic configuration does not do this, but this manual configuration example does — approximately 1 GB is left as free space for future expansion.

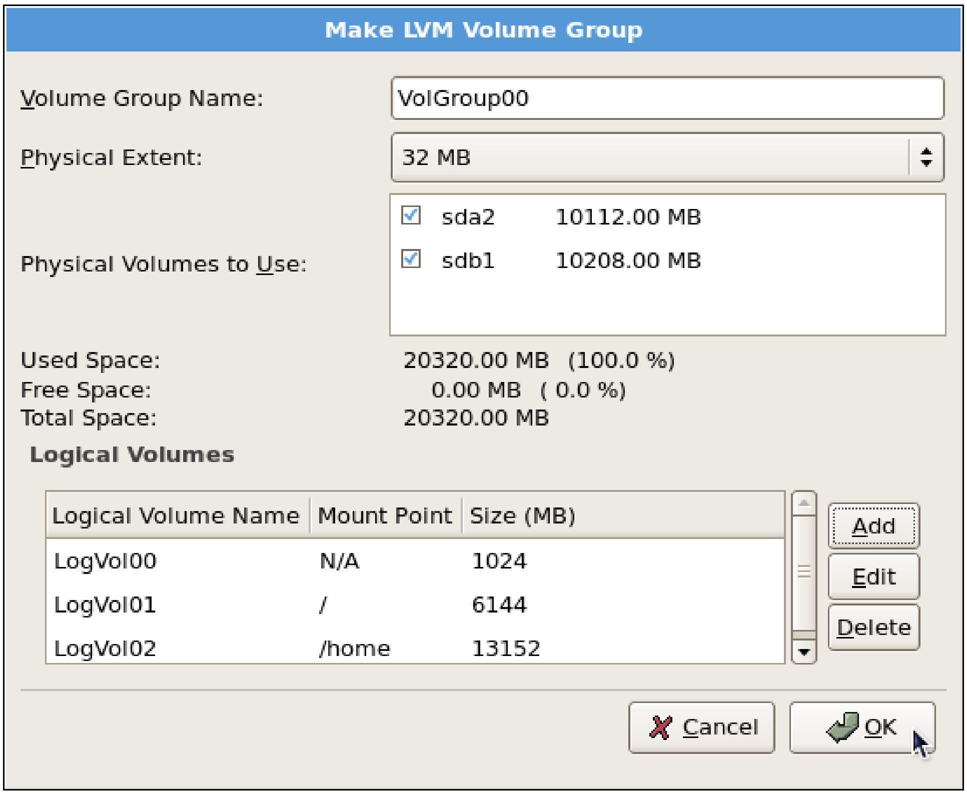

Make LVM Volume Group

Volume Group Name: VolGroup00

Physical Extent: 32 MB

Physical Volumes to Use:

☑ sda2 10112.00 MB

☑ sdb1 10208.00 MB

Used Space: 20320.00 MB (100.0 %)
Free Space: 0.00 MB (0.0 %)
Total Space: 20320.00 MB

Logical Volumes

Logical Volume Name	Mount Point	Size (MB)
LogVol00	N/A	1024
LogVol01	/	6144
LogVol02	/home	13152

그림 10.11. Pending Logical Volumes

Click OK to apply the volume group and all associated logical volumes.

The following figure shows the final manual configuration:

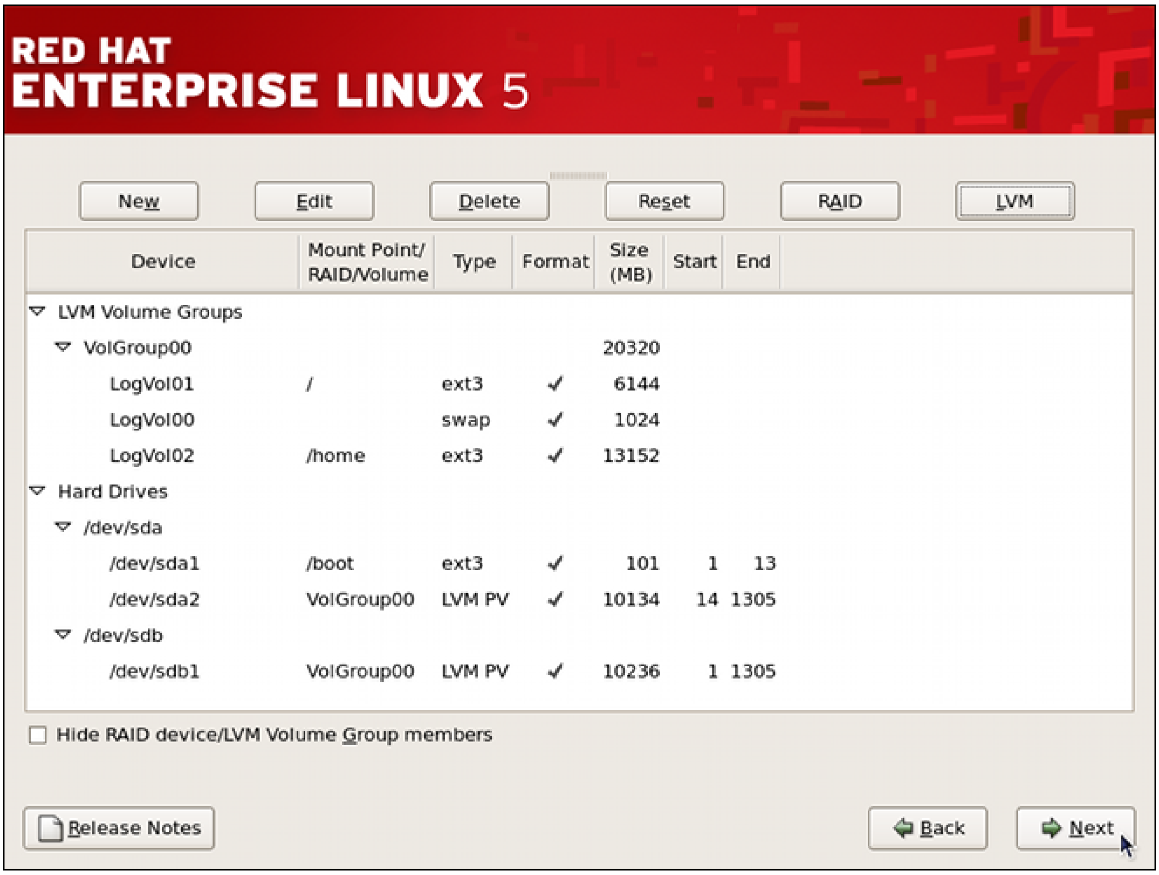

Device	Mount Point/ RAID/Volume	Type	Format	Size (MB)	Start	End
LVM Volume Groups						
VolGroup00				20320		
LogVol01	/	ext3	✓	6144		
LogVol00		swap	✓	1024		
LogVol02	/home	ext3	✓	13152		
Hard Drives						
/dev/sda						
/dev/sda1	/boot	ext3	✓	101	1	13
/dev/sda2	VolGroup00	LVM PV	✓	10134	14	1305
/dev/sdb						
/dev/sdb1	VolGroup00	LVM PV	✓	10236	1	1305

□ Hide RAID device/LVM Volume Group members

그림 10.12. Final Manual Configuration

10.5. Using the LVM utility system-config-lvm

The LVM utility allows you to manage logical volumes within X windows or graphically. You can access the application by selecting from your menu panel System > Administration > Logical Volume Management. Alternatively you can start the Logical Volume Management utility by typing system-config-lvm from a terminal.

In the example used in this section, the following are the details for the volume group that was created during the installation:

```
/boot - (Ext3) file system. Displayed under 'Uninitialized Entities'. (DO NOT initialize this partition).
LogVol00 - (LVM) contains the (/) directory (312 extents).
LogVol02 - (LVM) contains the (/home) directory (128 extents).
LogVol03 - (LVM) swap (28 extents).
```

The logical volumes above were created in disk entity /dev/hda2 while /boot was created in /dev/hda1. The system also consists of 'Uninitialized Entities' which are illustrated in 그림 10.17. “Uninitialized Entities” . The figure below illustrates the main window in the LVM utility. The logical and the physical views of the above configuration are illustrated below. The three logical volumes exist on the same physical volume (hda2).

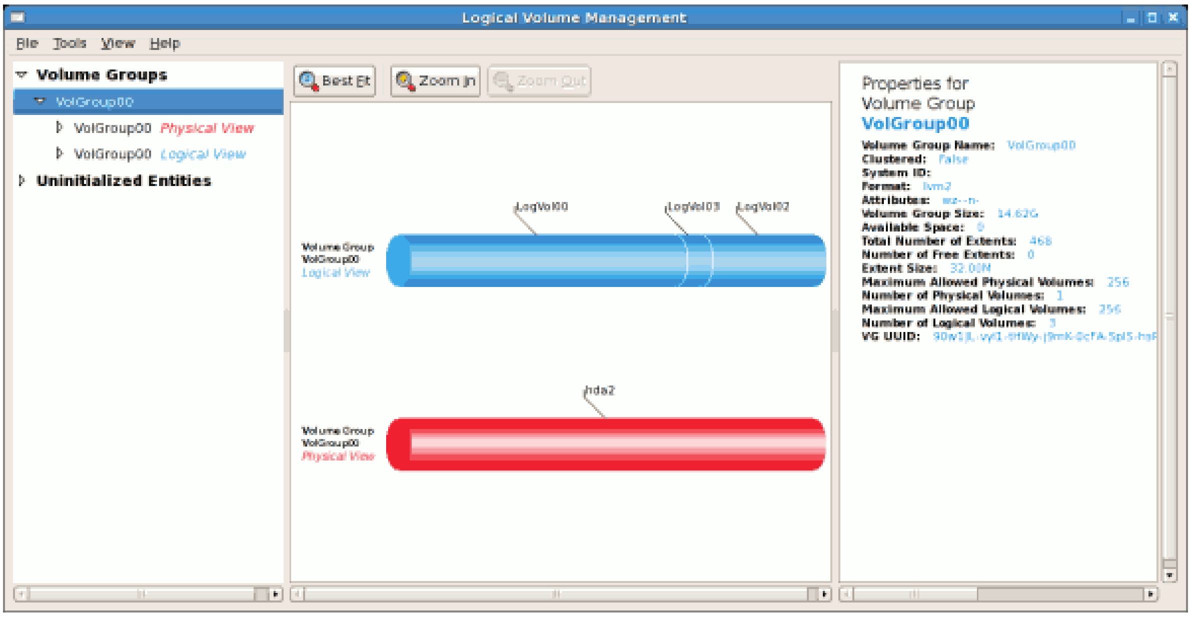

Properties for
Volume Group
VolGroup00

Volume Group Name: VolGroup00
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 14.62G
Available Space: 0
Total Number of Extents: 468
Number of Free Extents: 0
Extent Size: 32.00M
Maximum Allowed Physical Volumes: 256
Number of Physical Volumes: 1
Maximum Allowed Logical Volumes: 256
Number of Logical Volumes: 3
VG UUID: 90w1jL-vyl1-tHWy-j9mK-0cFA-Spl5-hqF

그림 10.13. Main LVM Window

The figure below illustrates the physical view for the volume. In this window, you can select and remove a volume from the volume group or migrate extents from the volume to another volume group. Steps to migrate extents are discussed in 그림 10.22. "Migrate Extents" .

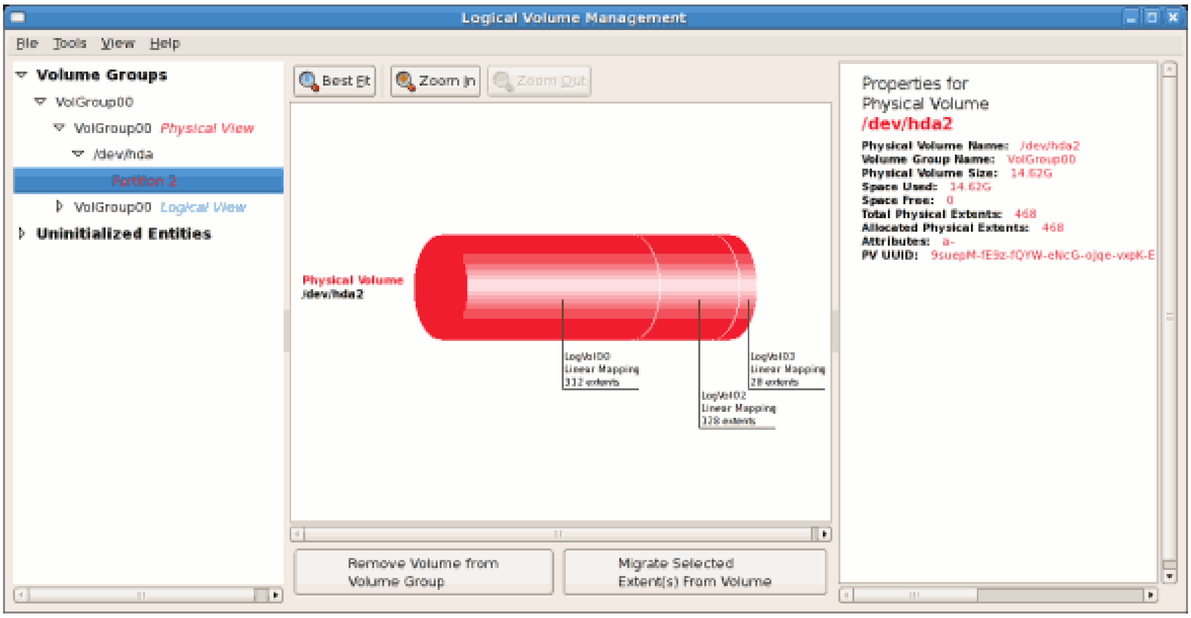

Properties for
Physical Volume
/dev/hda2

Physical Volume Name: /dev/hda2
Volume Group Name: VolGroup00
Physical Volume Size: 14.62G
Space Used: 14.62G
Space Free: 0
Total Physical Extents: 468
Allocated Physical Extents: 468
Attributes: a-
PV UUID: 9suepM-fE9z-fQYW-eNcG-ojqe-vxpK-E

그림 10.14. Physical View Window

The figure below illustrates the logical view for the selected volume group. The logical volume size is also indicated with the individual logical volume sizes illustrated.

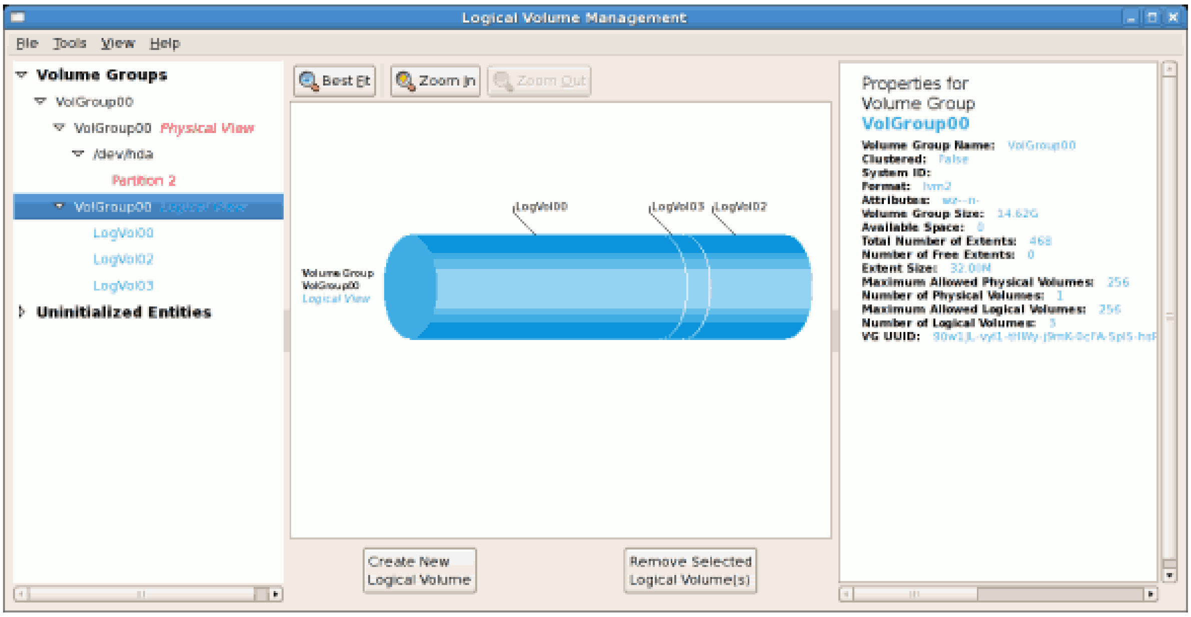

Properties for
Volume Group
VolGroup00

Volume Group Name: VolGroup00
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 14.62G
Available Space: 0
Total Number of Extents: 468
Number of Free Extents: 0
Extent Size: 32.00M
Maximum Allowed Physical Volumes: 256
Number of Physical Volumes: 1
Maximum Allowed Logical Volumes: 256
Number of Logical Volumes: 3
VG UUID: 90w1JL-vyl1-4lWy-j9mK-0cfA-Sp5-hsF

그림 10.15. Logical View Window

On the left side column, you can select the individual logical volumes in the volume group to view more details about each. In this example the objective is to rename the logical volume name for 'LogVol03' to 'Swap'. To perform this operation select the respective logical volume and click on the Edit Properties button. This will display the Edit Logical Volume window from which you can modify the Logical volume name, size (in extents) and also use the remaining space available in a logical volume group. The figure below illustrates this.

Please note that this logical volume cannot be changed in size as there is currently no free space in the volume group. If there was remaining space, this option would be enabled (see 그림 10.31. "Edit logical volume"). Click on the OK button to save your changes (this will remount the volume). To cancel your changes click on the Cancel button. To revert to the last snapshot settings click on the Revert button. A snapshot can be created by clicking on the Create Snapshot button on the LVM utility window. If the selected logical volume is in use by the system (for example) the / (root) directory, this task will not be successful as the volume cannot be unmounted.

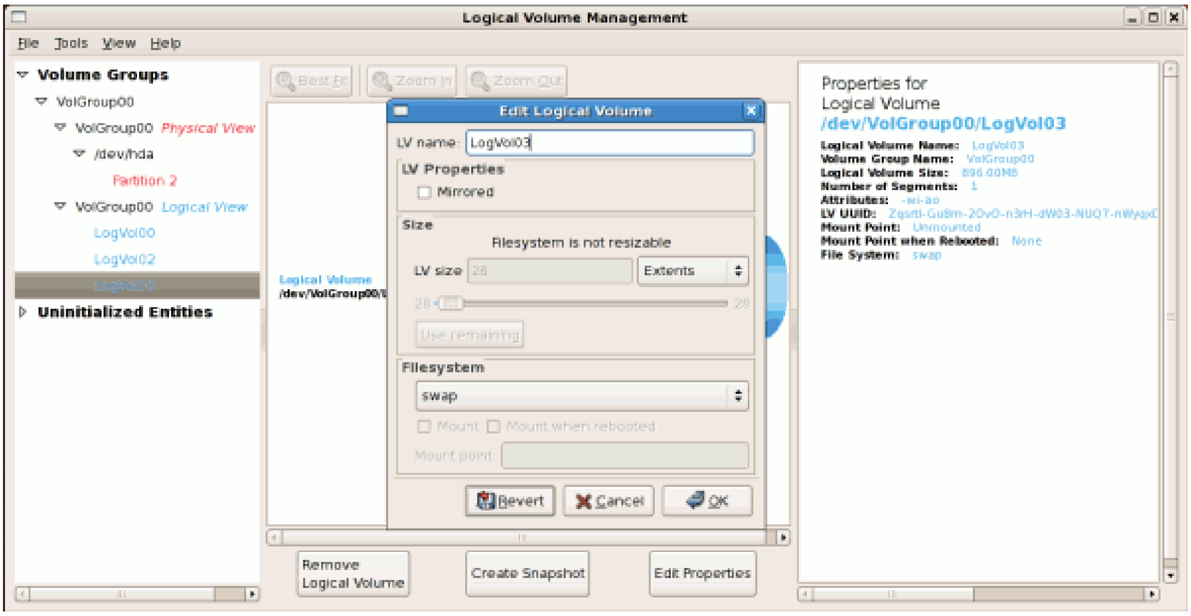

LV name: LogVol03
□ Mirrored
LV size: 28 Extents
28 28
swap
□ Mount □ Mount when rebooted

Properties for
Logical Volume
/dev/VolGroup00/LogVol03

Logical Volume Name: LogVol03
Volume Group Name: VolGroup00
Logical Volume Size: 896.00MB
Number of Segments: 1
Attributes: -wi-ao
LV UUID: Zgsrtl-Gu8m-2OvO-n3rH-dW03-NUQ7-nWyqxC
Mount Point: Unmounted
Mount Point when Rebooted: None
File System: swap

그림 10.16. Edit Logical Volume

10.5.1. Utilizing uninitialized entities

'Uninitialized Entities' consist of unpartitioned space and non LVM file systems. In this example partitions 3, 4, 5, 6 and 7 were created during installation and some unpartitioned space was left on the hard disk. Please view each partition and ensure that you read the 'Properties for Disk Entity' on the right column of the window to ensure that you do not delete critical data. In this example partition 1 cannot be initialized as it is /boot. Uninitialized entities are illustrated below.

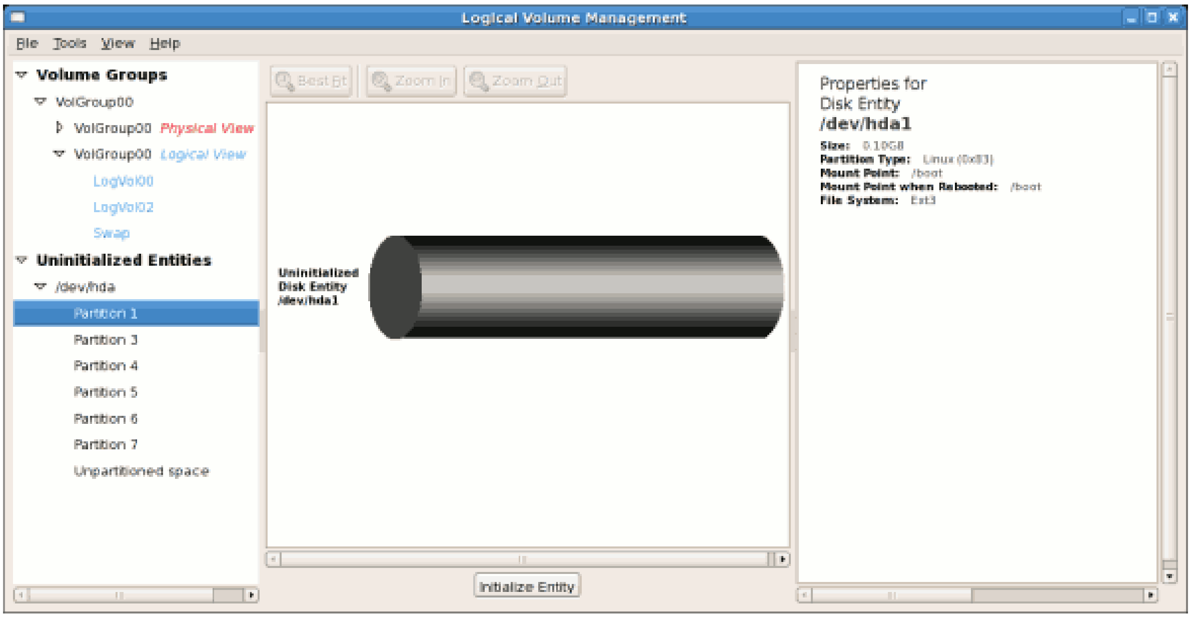

Properties for
Disk Entity
/dev/hda1

Size: 0.10GB
Partition Type: Linux (0x83)
Mount Point: /boot
Mount Point when Rebooted: /boot
File System: Ext3

그림 10.17. Uninitialized Entities

In this example, partition 3 will be initialized and added to an existing volume group. To initialize a partition or unpartioned space, select the partition and click on the Initialize Entity button. Once initialized, a volume will be listed in the 'Unallocated Volumes' list.

10.5.2. Adding Unallocated Volumes to a volume group

Once initialized, a volume will be listed in the 'Unallocated Volumes' list. The figure below illustrates an unallocated partition (Partition 3). The respective buttons at the bottom of the window allow you to:

- create a new volume group,
- add the unallocated volume to an existing volume group,
- remove the volume from LVM.

To add the volume to an existing volume group, click on the Add to Existing Volume Group button.

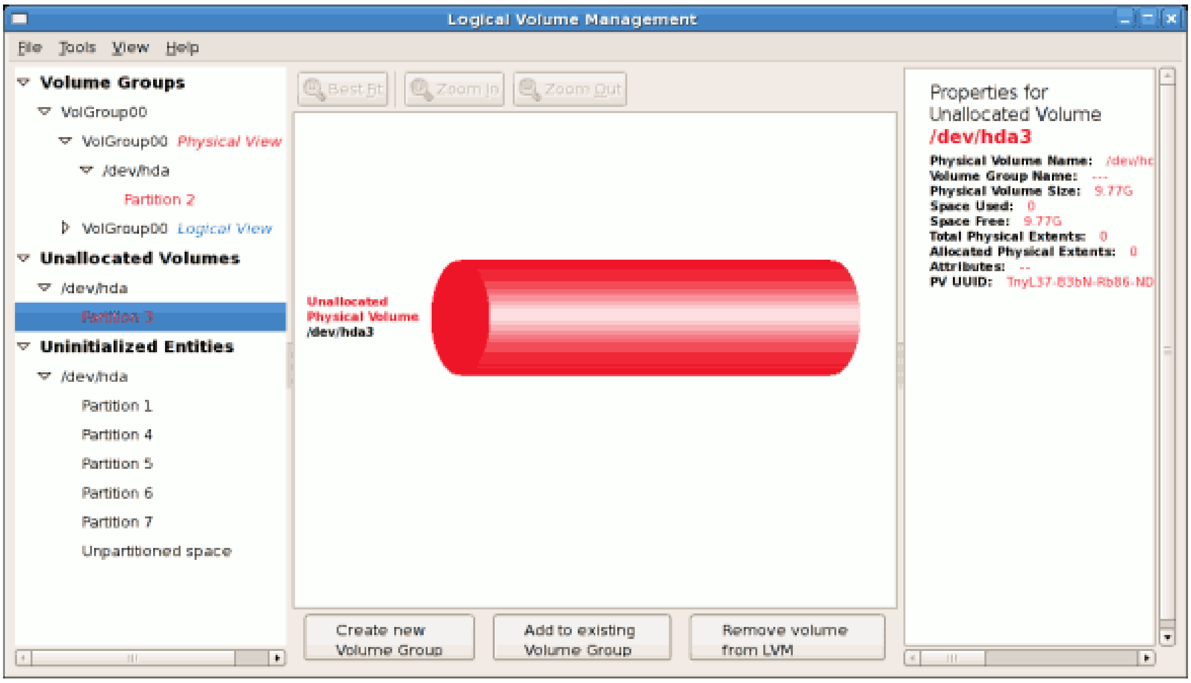

Properties for Unallocated Volume /dev/hda3

Physical Volume Name: /dev/hc
Volume Group Name: ---
Physical Volume Size: 9.77G
Space Used: 0
Space Free: 9.77G
Total Physical Extents: 0
Allocated Physical Extents: 0
Attributes: --
PV UUID: TnyL37-B3bN-Rb86-ND

그림 10.18. Unallocated Volumes

Clicking on the Add to Existing Volume Group button will display a pop up window listing the existing volume groups to which you can add the physical volume you are about to initialize. A volume group may span across one or more hard disks. In this example only one volume group exists as illustrated below.

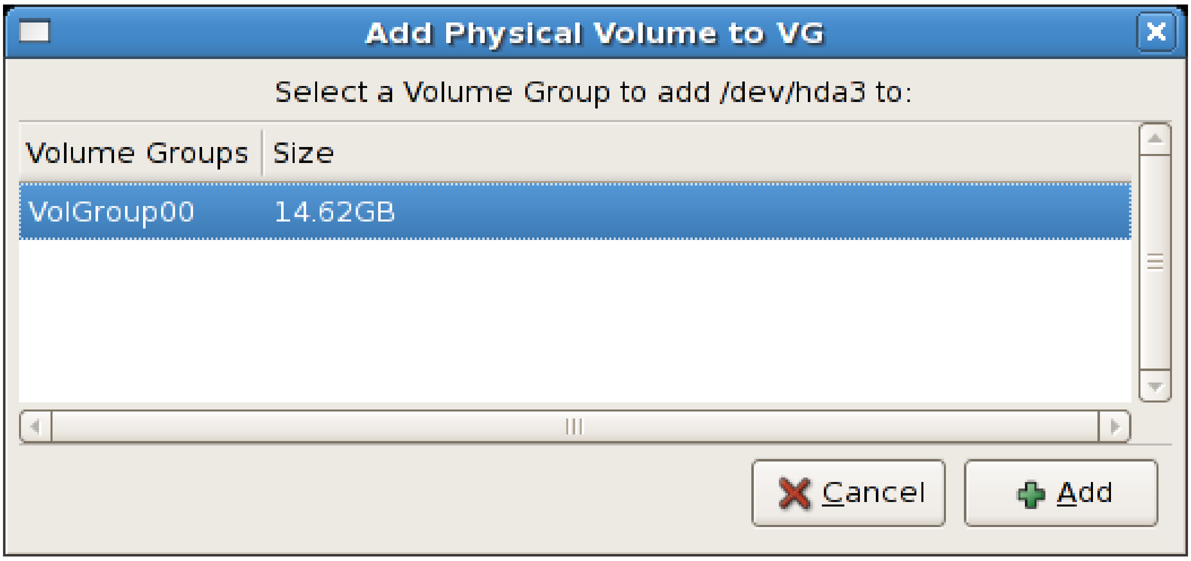

Select a Volume Group to add /dev/hda3 to:

Volume Groups	Size
VolGroup00	14.62GB

그림 10.19. Add physical volume to volume group

Once added to an existing volume group the new logical volume is automatically added to the unused space of the selected volume group. You can use the unused space to:

- create a new logical volume (click on the Create New Logical Volume(s) button,
- select one of the existing logical volumes and increase the extents (see 10.5.6절. "Extending a volume group"),
- select an existing logical volume and remove it from the volume group by clicking on the Remove Selected Logical Volume(s) button. Please note that you cannot select unused space to perform this operation.

The figure below illustrates the logical view of 'VolGroup00' after adding the new volume group.

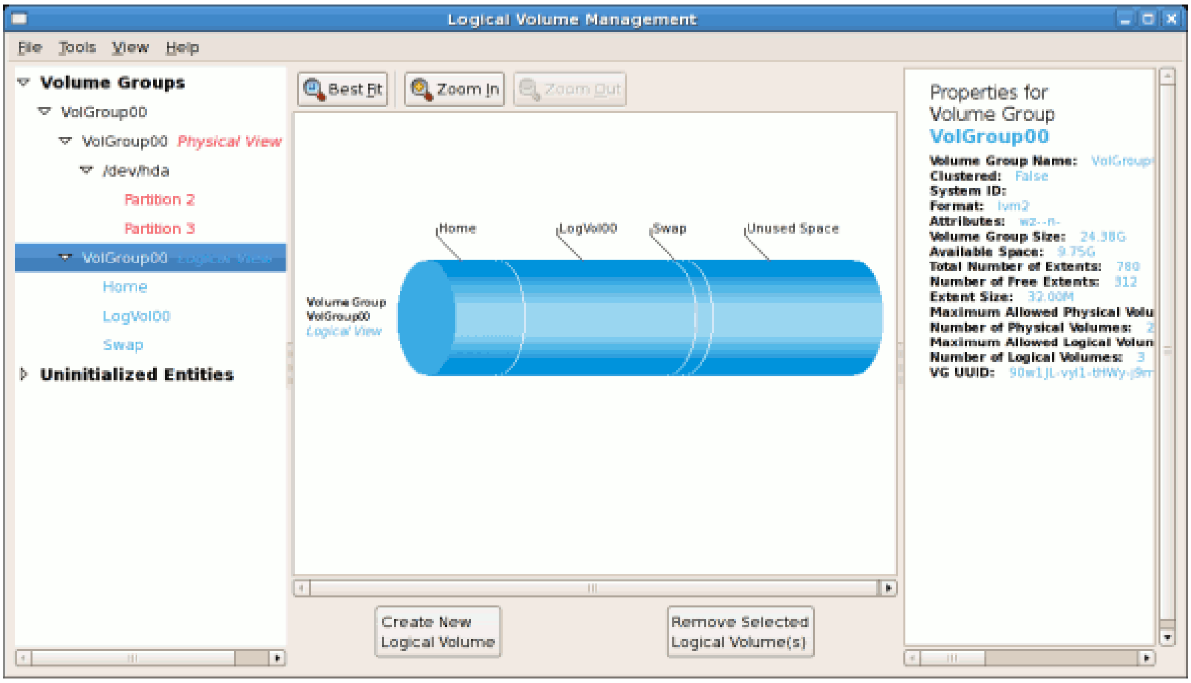

Properties for
Volume Group
VolGroup00

Volume Group Name: VolGroup
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 24.38G
Available Space: 9.75G
Total Number of Extents: 780
Number of Free Extents: 312
Extent Size: 32.00M
Maximum Allowed Physical Volu
Number of Physical Volumes: 2
Maximum Allowed Logical Volun
Number of Logical Volumes: 3
VG UUID: 90w1JL-vyl1-tHWy-j9m

그림 10.20. Logical view of volume group

In the figure below, the uninitialized entities (partitions 3, 5, 6 and 7) were added to 'VolGroup00'.

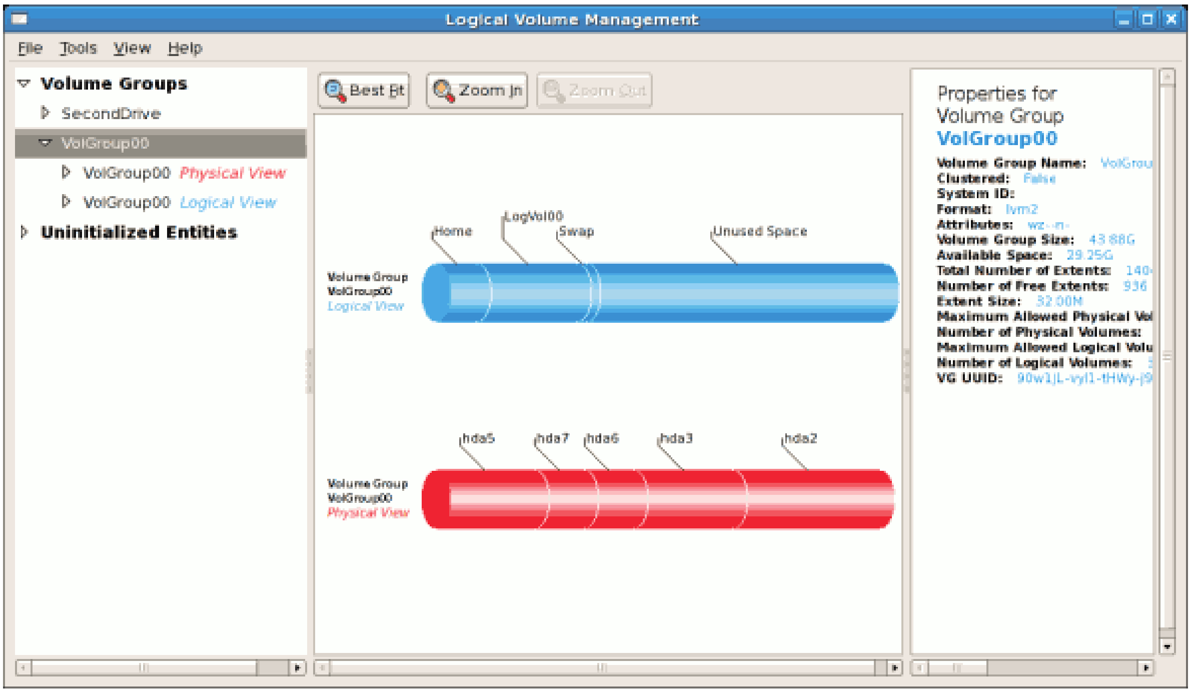

Properties for
Volume Group
VolGroup00

Volume Group Name: VolGrou
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 43.88G
Available Space: 29.25G
Total Number of Extents: 140
Number of Free Extents: 936
Extent Size: 32.00M
Maximum Allowed Physical Vol
Number of Physical Volumes:
Maximum Allowed Logical Volu
Number of Logical Volumes: 3
VG UUID: 90w1JL-vyl1-tHWy-j9

그림 10.21. Logical view of volume group

10.5.3. Migrating extents

To migrate extents from a physical volume, select the volume and click on the Migrate Selected Extent(s) From Volume button. Please note that you need to have a sufficient number of free extents to migrate extents within a volume group. An error message will be displayed if you do not have a sufficient number of free extents. To resolve this problem, please extend your volume group (see 10.5.6절. “Extending a volume group”). If a sufficient number of free extents is detected in the volume group, a pop up window will be displayed from which you can select the destination for the extents or automatically let LVM choose the physical volumes (PVs) to migrate them to. This is illustrated below.

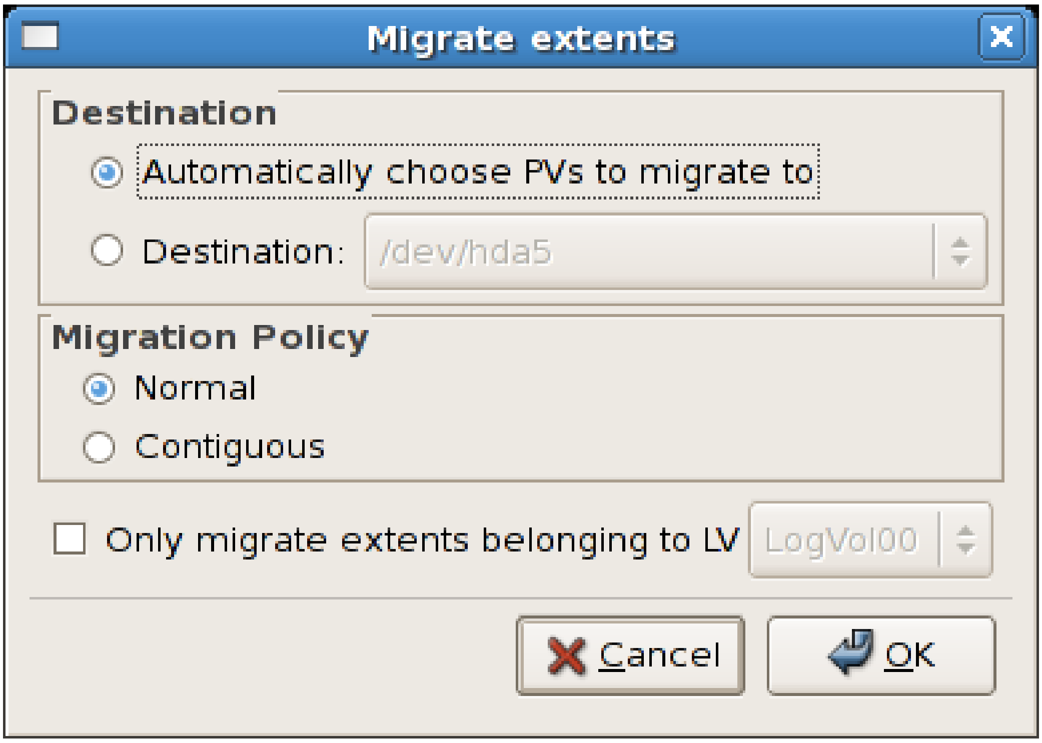

Migrate extents

Destination

● Automatically choose PVs to migrate to

○ Destination: /dev/hda5

Migration Policy

● Normal

○ Contiguous

☐ Only migrate extents belonging to LV LogVol00

그림 10.22. Migrate Extents

The figure below illustrates a migration of extents in progress. In this example, the extents were migrated to 'Partition 3'.

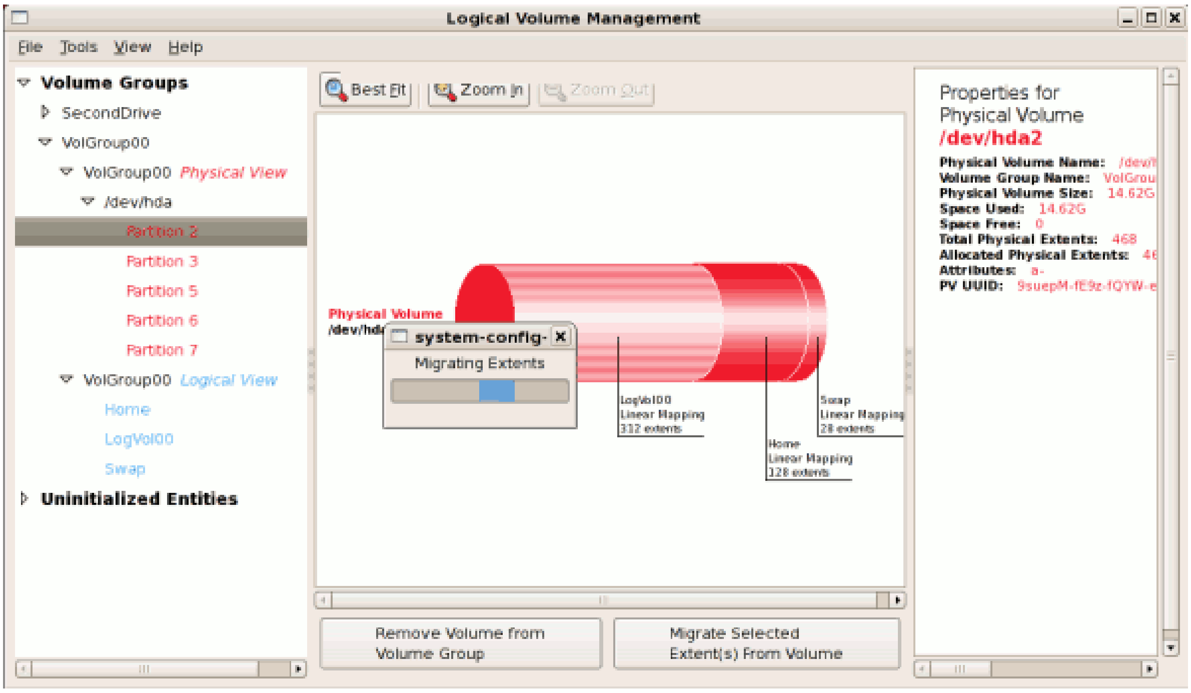

Logical Volume Management

- Volume Groups
  - SecondDrive
  - VolGroup00
    - VolGroup00 *Physical View*
      - /dev/hda
        - Partition 2
        - Partition 3
        - Partition 5
        - Partition 6
        - Partition 7
    - VolGroup00 *Logical View*
      - Home
      - LogVol00
      - Swap
- Uninitialized Entities

system-config-

Migrating Extents

Properties for Physical Volume **/dev/hda2**

Physical Volume Name: /dev/
Volume Group Name: VolGrou
Physical Volume Size: 14.62G
Space Used: 14.62G
Space Free: 0
Total Physical Extents: 468
Allocated Physical Extents: 46
Attributes: a-
PV UUID: 9suepM-fE9z-fOYW-e

그림 10.23. Migrating extents in progress

Once the extents have been migrated, unused space is left on the physical volume. The figure below illustrates the physical and logical view for the volume group. Please note that the extents of

LogVol00 which were initially in hda2 are now in hda3. Migrating extents allows you to move logical volumes in case of hard disk upgrades or to manage your disk space better.

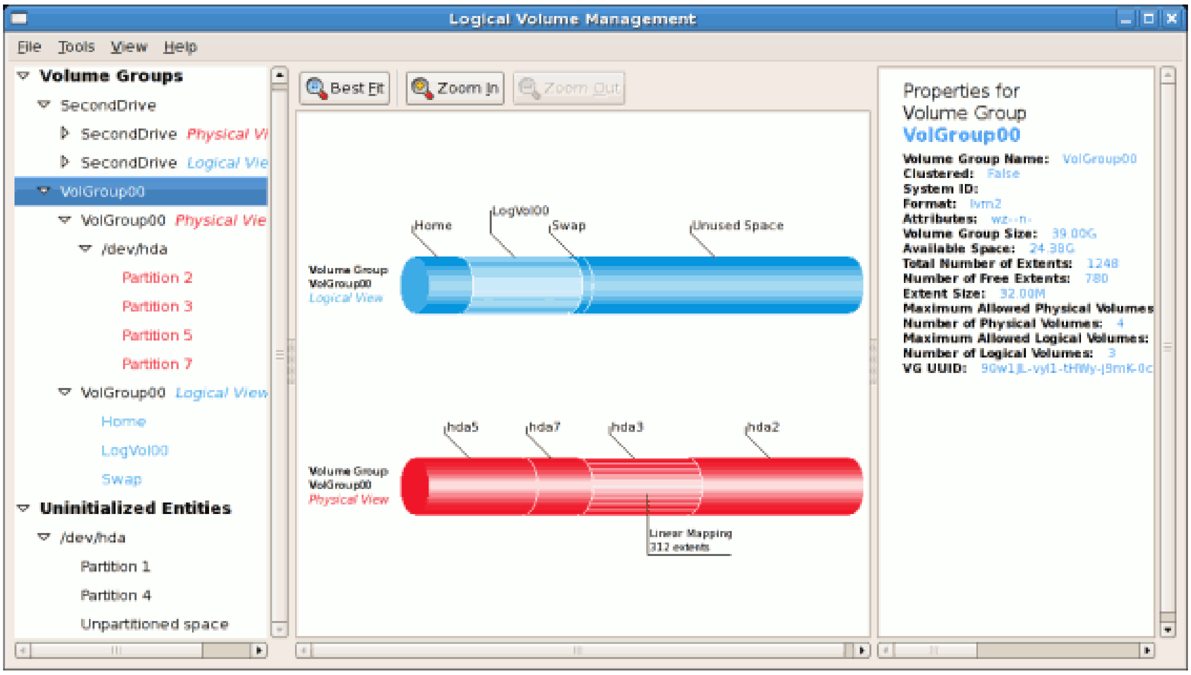

Properties for
Volume Group
VolGroup00

Volume Group Name: VolGroup00
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 39.00G
Available Space: 24.38G
Total Number of Extents: 1248
Number of Free Extents: 780
Extent Size: 32.00M
Maximum Allowed Physical Volumes
Number of Physical Volumes: 4
Maximum Allowed Logical Volumes:
Number of Logical Volumes: 3
VG UUID: 90w1JL-vyl1-tHWy-j9mK-0c

그림 10.24. Logical and physical view of volume group

10.5.4. Adding a new hard disk using LVM

In this example, a new IDE hard disk was added. The figure below illustrates the details for the new hard disk. From the figure below, the disk is uninitialized and not mounted. To initialize a partition, click on the Initialize Entity button. For more details, see 10.5.1절. "Utilizing uninitialized entities" . Once initialized, LVM will add the new volume to the list of unallocated volumes as illustrated in 그림 10.26. "Create new volume group" .

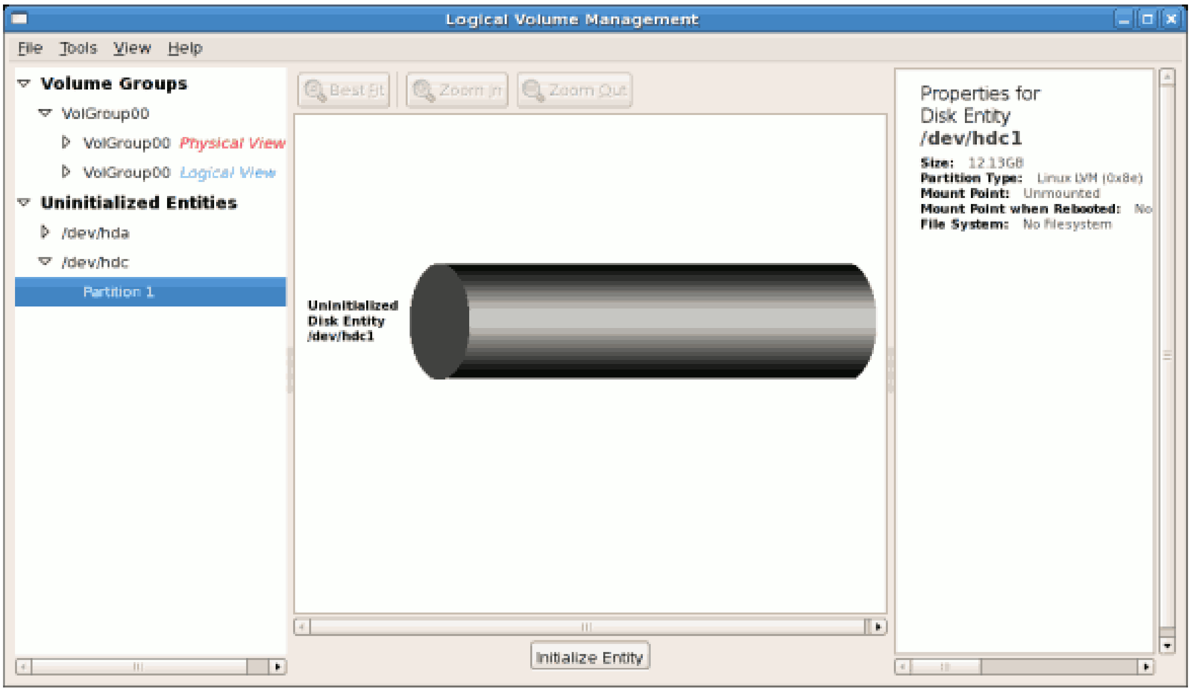

Properties for
Disk Entity
/dev/hdc1

Size: 12.13GB
Partition Type: Linux LVM (0x8e)
Mount Point: Unmounted
Mount Point when Rebooted: No
File System: No filesystem

그림 10.25. Uninitialized hard disk

10.5.5. Adding a new volume group

Once initialized, LVM will add the new volume to the list of unallocated volumes where you can add it to an existing volume group or create a new volume group. You can also remove the volume from LVM. The volume if removed from LVM will be listed in the list of 'Uninitialized Entities' as illustrated in 그림 10.25. "Uninitialized hard disk" . In this example, a new volume group was created as illustrated below.

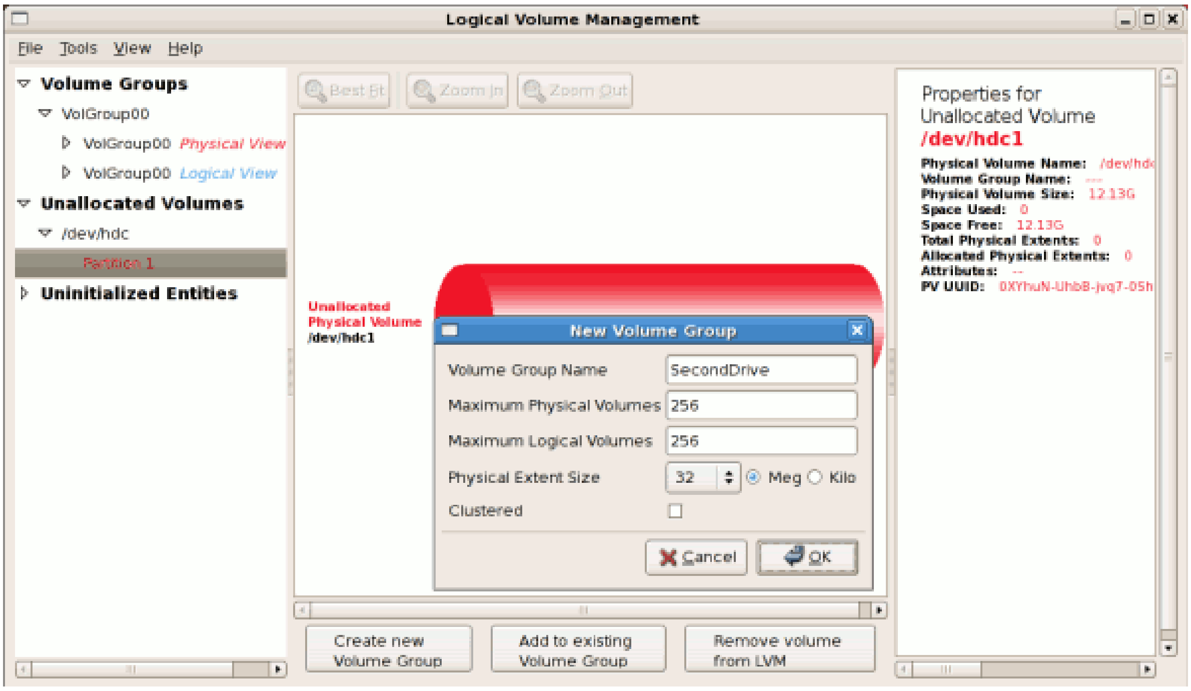

Volume Group Name: SecondDrive
Maximum Physical Volumes: 256
Maximum Logical Volumes: 256
Physical Extent Size: 32 ● Meg ○ Kilo
Clustered □

Properties for Unallocated Volume /dev/hdc1

Physical Volume Name: /dev/hdc
Volume Group Name: ---
Physical Volume Size: 12.13G
Space Used: 0
Space Free: 12.13G
Total Physical Extents: 0
Allocated Physical Extents: 0
Attributes: --
PV UUID: 0XYhuN-UhbB-jvq7-05h

그림 10.26. Create new volume group

Once created a new volume group will be displayed in the list of existing volume groups as illustrated below. The logical view will display the new volume group with unused space as no logical volumes have been created. To create a logical volume, select the volume group and click on the Create New Logical Volume button as illustrated below. Please select the extents you wish to use on the volume group. In this example, all the extents in the volume group were used to create the new logical volume.

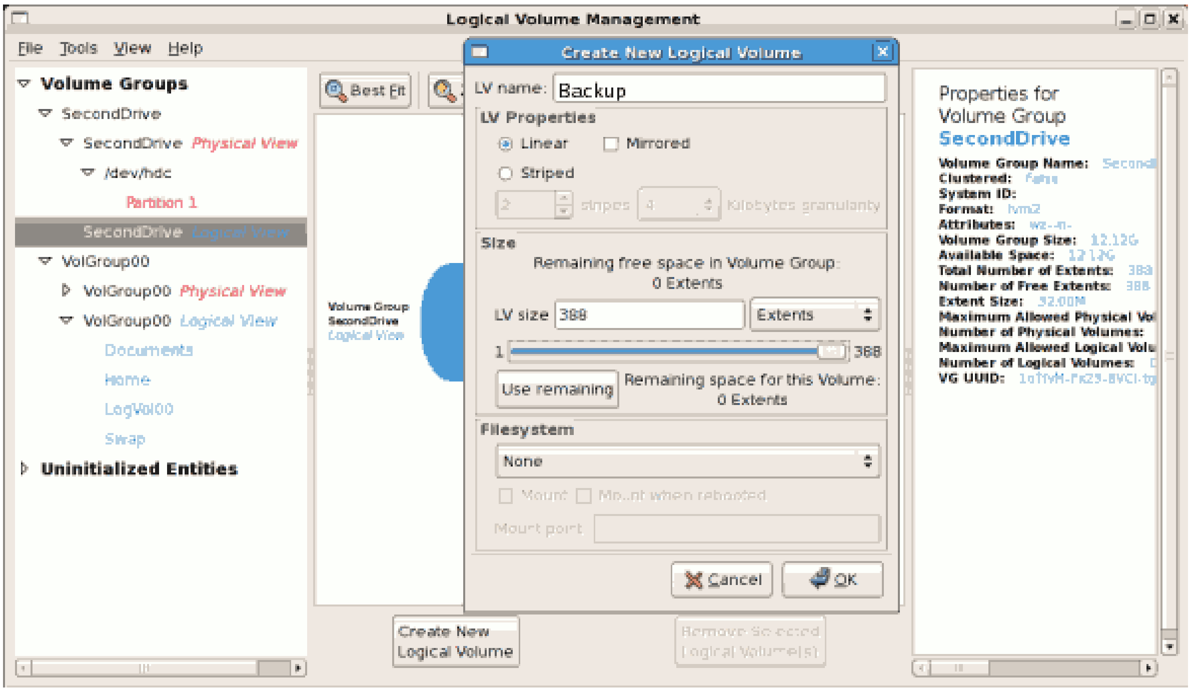

Create New Logical Volume

LV name: Backup

LV Properties

● Linear □ Mirrored

○ Striped

2 stripes 4 Kilobytes granularity

Size

Remaining free space in Volume Group:
0 Extents

LV size 388 Extents

1 388

Use remaining

Remaining space for this Volume:
0 Extents

Filesystem

None

□ Mount □ Mount when rebooted

Mount point

Cancel OK

Properties for
Volume Group
SecondDrive

Volume Group Name: Second
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 12.12G
Available Space: 12.12G
Total Number of Extents: 388
Number of Free Extents: 388
Extent Size: 32.00M
Maximum Allowed Physical Vol
Number of Physical Volumes:
Maximum Allowed Logical Volu
Number of Logical Volumes: 0
VG UUID: 1oTfvM-Fx23-8VCI-tg

그림 10.27. Create new logical volume

The figure below illustrates the physical view of the new volume group. The new logical volume named 'Backups' in this volume group is also listed.

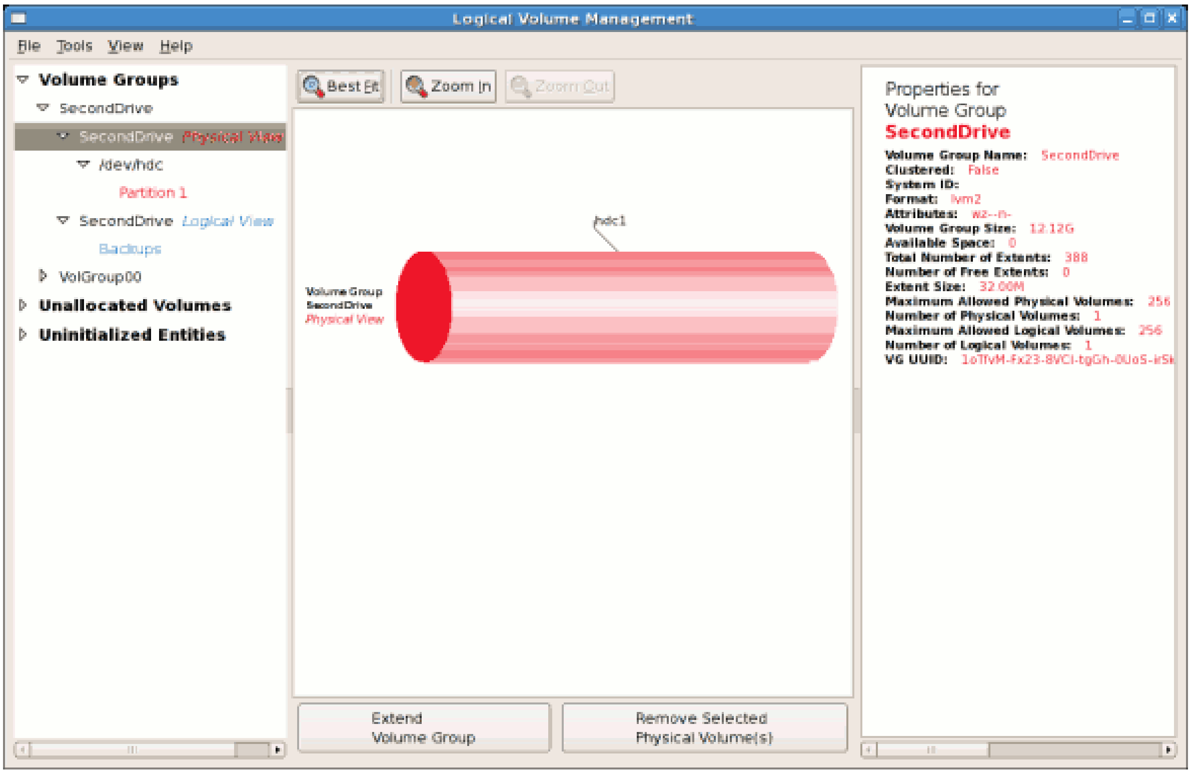

Properties for
Volume Group
SecondDrive

Volume Group Name: SecondDrive
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 12.12G
Available Space: 0
Total Number of Extents: 388
Number of Free Extents: 0
Extent Size: 32.00M
Maximum Allowed Physical Volumes: 256
Number of Physical Volumes: 1
Maximum Allowed Logical Volumes: 256
Number of Logical Volumes: 1
VG UUID: 1oTfvM-Fx23-8VCI-tgGh-0UoS-irSk

그림 10.28. Physical view of new volume group

10.5.6. Extending a volume group

In this example, the objective was to extend the new volume group to include an uninitialized entity (partition). This was to increase the size or number of extents for the volume group. To extend the volume group, click on the Extend Volume Group button. This will display the 'Extend

Volume Group' window as illustrated below. On the 'Extend Volume Group' window, you can select disk entities (partitions) to add to the volume group. Please ensure that you check the contents of any 'Uninitialized Disk Entities' (partitions) to avoid deleting any critical data (see 그림 10.25. "Uninitialized hard disk"). In the example, the disk entity (partition) /dev/hda6 was selected as illustrated below.

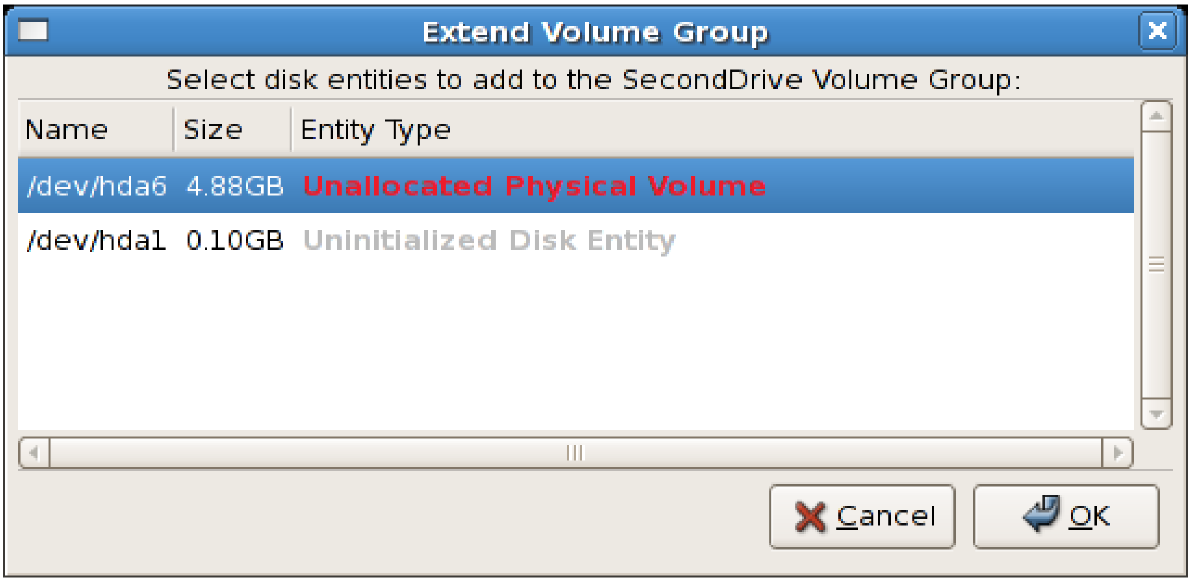

Select disk entities to add to the SecondDrive Volume Group:

Name	Size	Entity Type
/dev/hda6	4.88GB	Unallocated Physical Volume
/dev/hda1	0.10GB	Uninitialized Disk Entity

그림 10.29. Select disk entities

Once added, the new volume will be added as 'Unused Space' in the volume group. The figure below illustrates the logical and physical view of the volume group after it was extended.

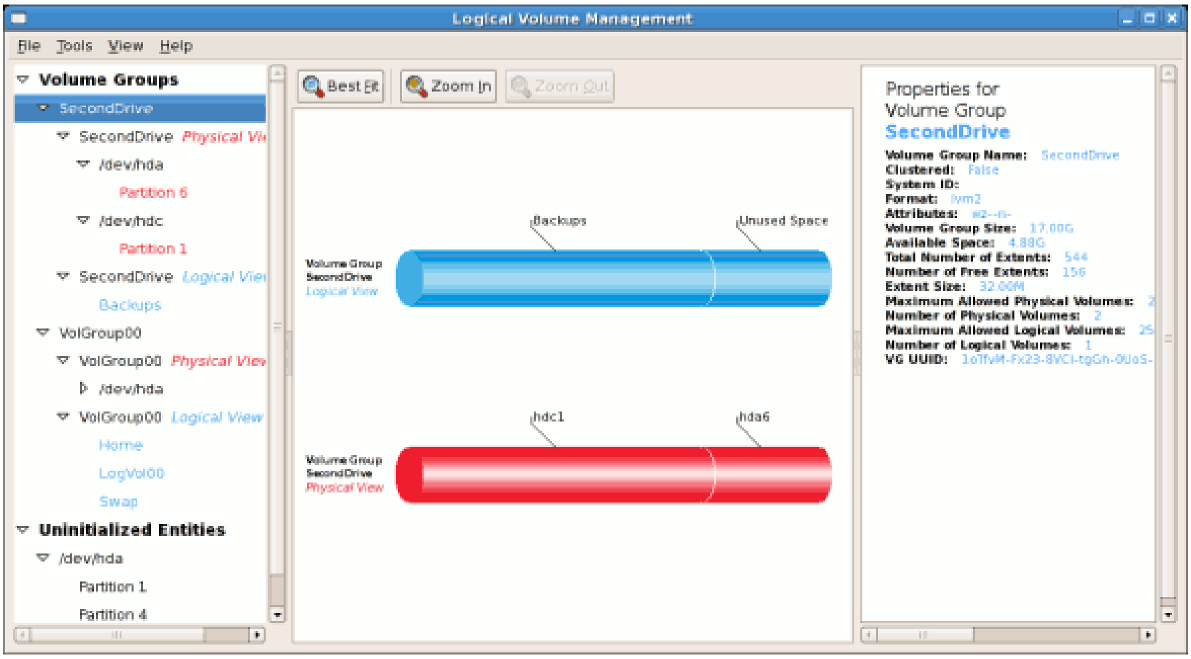

Properties for Volume Group SecondDrive

Volume Group Name: SecondDrive
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 17.00G
Available Space: 4.88G
Total Number of Extents: 544
Number of Free Extents: 156
Extent Size: 32.00M
Maximum Allowed Physical Volumes: 2
Number of Physical Volumes: 2
Maximum Allowed Logical Volumes: 25
Number of Logical Volumes: 1
VG UUID: 1oTfvM-Fx23-8VCI-tgGh-0UoS-

그림 10.30. Logical and physical view of an extended volume group

10.5.7. Editing a Logical Volume

The LVM utility allows you to select a logical volume in the volume group and modify its name, size and specify filesystem options. In this example, the logical volume named 'Backups" was extended onto the remaining space for the volume group.

Clicking on the Edit Properties button will display the 'Edit Logical Volume' popup window from which you can edit the properties of the logical volume. On this window, you can also mount the volume after making the changes and mount it when the system is rebooted. Please note that you should indicate the mount point. If the mount point you specify does not exist, a popup window will be displayed prompting you to create it. The 'Edit Logical Volume' window is illustrated below.

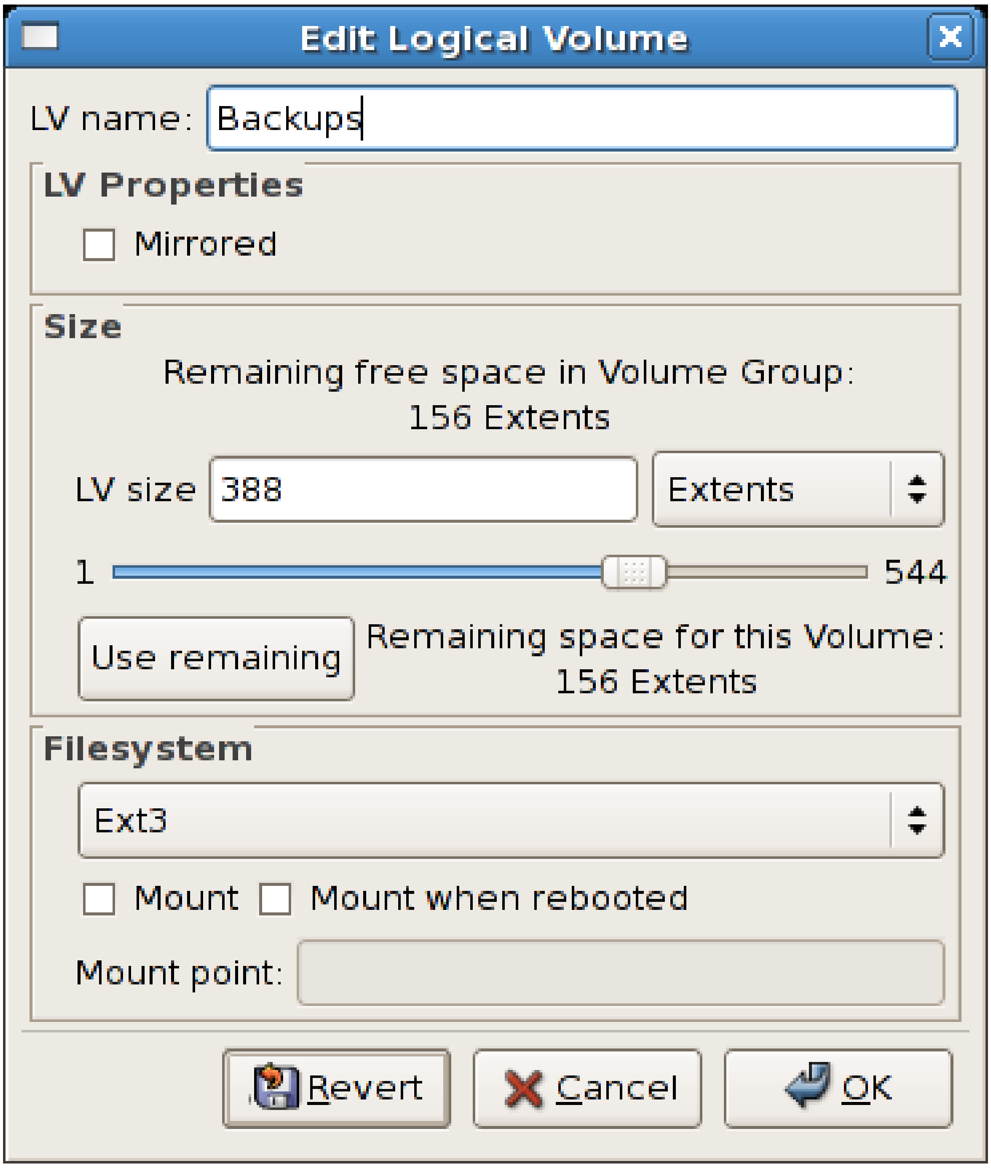

LV name: Backups

□ Mirrored

Remaining free space in Volume Group: 156 Extents

LV size 388 Extents

Remaining space for this Volume: 156 Extents

Ext3

□ Mount □ Mount when rebooted

Mount point:

그림 10.31. Edit logical volume

If you wish to mount the volume, select the 'Mount' checkbox indicating the preferred mount point. To mount the volume when the system is rebooted, select the 'Mount when rebooted' checkbox. In this example, the new volume will be mounted in /mnt/backups. This is illustrated in the figure below.

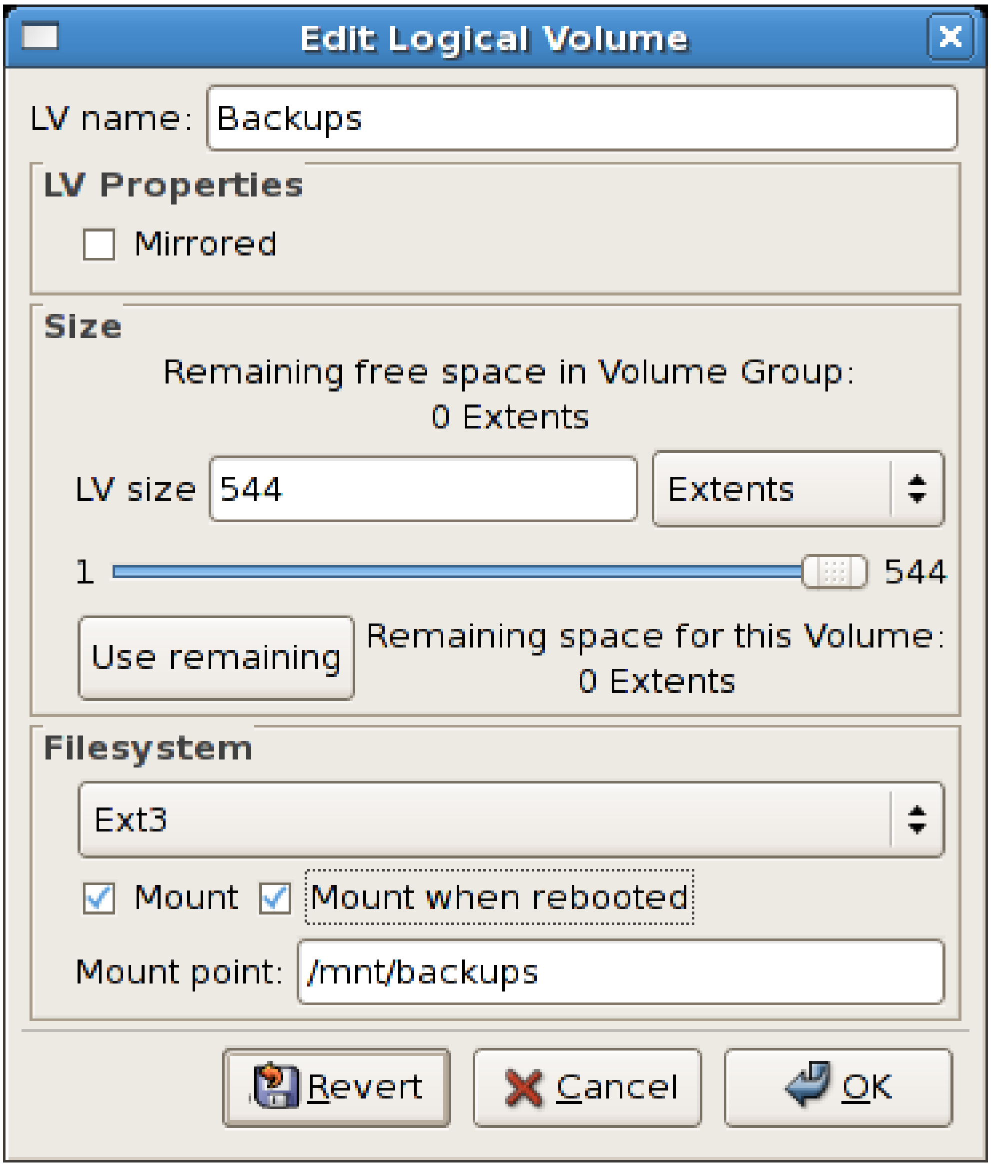

Edit Logical Volume

LV name: Backups

LV Properties

□ Mirrored

Size

Remaining free space in Volume Group:
0 Extents

LV size 544 Extents

1 544

Remaining space for this Volume:
0 Extents

Filesystem

Ext3

☑ Mount ☑ Mount when rebooted

Mount point: /mnt/backups

그림 10.32. Edit logical volume - specifying mount options

The figure below illustrates the logical and physical view of the volume group after the logical volume was extended to the unused space. Please note in this example that the logical volume named 'Backups' spans across two hard disks. A volume can be striped across two or more physical devices using LVM.

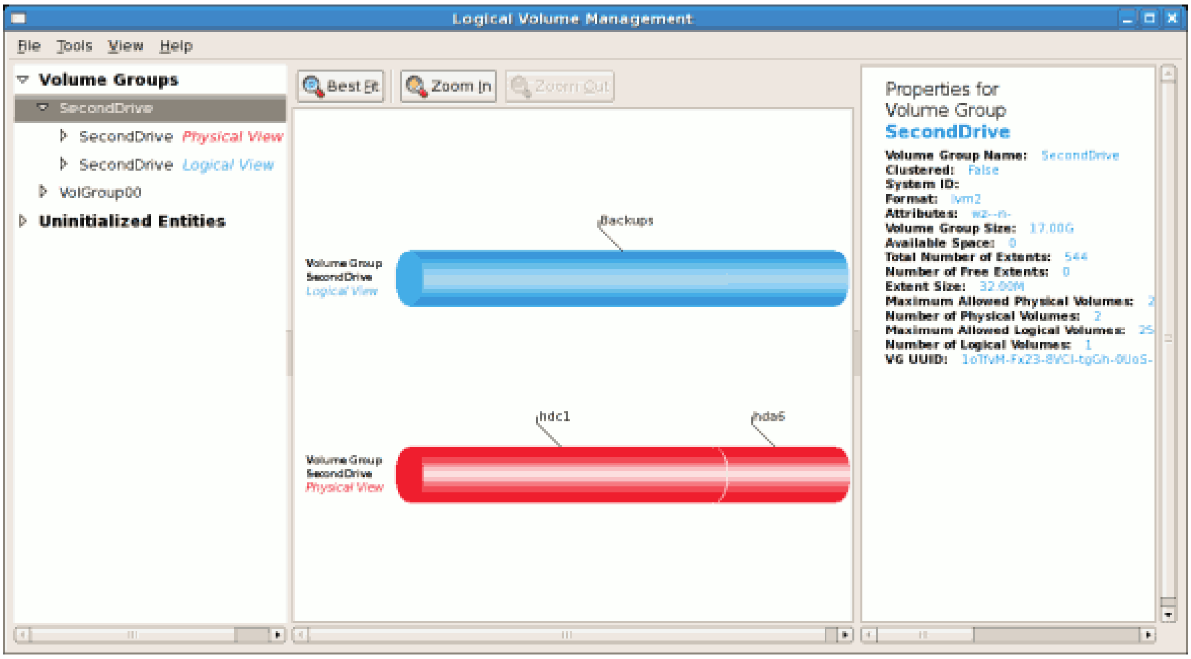

Properties for
Volume Group
SecondDrive

Volume Group Name: SecondDrive
Clustered: False
System ID:
Format: lvm2
Attributes: wz--n-
Volume Group Size: 17.00G
Available Space: 0
Total Number of Extents: 544
Number of Free Extents: 0
Extent Size: 32.00M
Maximum Allowed Physical Volumes: 2
Number of Physical Volumes: 2
Maximum Allowed Logical Volumes: 25
Number of Logical Volumes: 1
VG UUID: 1oTfvM-Fx23-8VCl-tgGh-0UoS-

그림 10.33. Edit logical volume

10.6. Additional Resources

Use these sources to learn more about LVM.

10.6.1. Installed Documentation

- rpm -qd lvm2 — This command shows all the documentation available from the lvm package, including man pages.

- lvm help — This command shows all LVM commands available.

10.6.2. Useful Websites

- http://sources.redhat.com/lvm2 — LVM2 webpage, which contains an overview, link to the mailing lists, and more.

- http://tldp.org/HOWTO/LVM-HOWTO/ — LVM HOWTO from the Linux Documentation Project.

부 II. 패키지 관리

Red Hat Enterprise Linux 시스템 상에 있는 모든 소프트웨어는 설치, 업그레이드, 삭제할 수 있는 RPM 패키지로 나뉘어집니다. 다음 부분에서는 그래픽 및 명령행 도구를 사용하여 Red Hat Enterprise Linux 시스템에 있는 RPM 패키지를 관리하는 방법에 관해 설명합니다.

RPM을 사용한 패키지 관리

The RPM Package Manager (RPM) is an open packaging system, which runs on Red Hat Enterprise Linux as well as other Linux and UNIX systems. Red Hat, Inc. encourages other vendors to use RPM for their own products. RPM is distributed under the terms of the GPL.

The utility works only with packages built for processing by the rpm package. For the end user, RPM makes system updates easy. Installing, uninstalling, and upgrading RPM packages can be accomplished with short commands. RPM maintains a database of installed packages and their files, so you can invoke powerful queries and verifications on your system. If you prefer a graphical interface, you can use the Package Management Tool to perform many RPM commands. Refer to 12 장. Package Management Tool for details.

> **중요**
>
> 패키지를 설치할 때 사용자의 운영체제 및 구조에 호환되는지 확인하시기 바랍니다. 주로 패키지 이름으로 호환 여부를 확인할 수 있습니다.

During upgrades, RPM handles configuration files carefully, so that you never lose your customizations — something that you cannot accomplish with regular .tar.gz files.

개발자는 RPM을 사용하여 소프트웨어 소스 코드를 일반 사용자가 사용할 수 있는 소스와 바이너리 패키지로 구성할 수 있습니다. 이 과정은 매우 단순하며 단독 파일과 생성한 옵션 패치를 주로 사용합니다. 이렇게 기존 소스와 개발 지시 사항이 포함된 패치를 확연히 구분함으로써 새로운 버전의 소프트웨어가 배포될 때마다 쉽게 패키지를 관리하실 수 있습니다.

> **알림**
>
> RPM은 시스템을 변경하므로 RPM 패키지 설치, 제거 및 업그레이드할 때 반드시 루트로 로그인해야 합니다.

11.1. RPM 설계 목표

RPM의 사용법을 이해하려면 우선 RPM의 설계 목적을 이해하시는 것이 도움이 될 것입니다:

업그레이드 기능

RPM을 사용하면 시스템을 완전히 재설치를 하지 않고 개별 구성 요소를 업그레이드할 수 있습니다. RPM에 기반한 새로운 운영체제(예, Red Hat Enterprise Linux) 버전이 배포될 때 다른 패키징 시스템에 기반한 운영체제에서와 같이 다시 설치할 필요가 없습니다. RPM은 지능화되고 완전히 자동화된 인-플레이스 업그레이드(in-place upgrade)를 수행합니다. 패키지의 구성 파일은 업그레이드가 실행되는 동안 모두 보존되므로 사용자 설정 또한 그대로 보존됩니다. 같은 RPM 파일을 사용하여 시스템에서 패키지를 설치하고 업그레이드하기 때문에 특별한 업그레이드 파일이 필요하지 않습니다.

뛰어난 질의 기능

RPM은 뛰어난 질의(query) 옵션을 제공하도록 설계되었습니다. 따라서, 전체 데이터베이스에 저장된 패키지나 특정 파일을 검색하실 수 있으며, 어떠한 파일이 어느 패키지에 담겨 있는지

와 그 패키지의 출처를 쉽게 알아낼 수 있습니다. RPM 패키지에 포함된 파일들은 압축된 아카이브 형식으로 구성되어 있으며, 패키지에 대한 유용한 정보와 내용을 포함하고 있는 사용자 정의 바이너리 헤더 덕분에 개별 패키지를 쉽고 빠르게 질의하실 수 있습니다.

시스템 검증 기능

RPM의 뛰어난 기능 중 하나는 패키지를 검증할 수 있는 능력입니다. 일부 패키지에 필요한 중요한 파일을 삭제되었는지 걱정되면 간단히 패키지를 검증해 보시기 바랍니다. 이상한 점이 발견되면 알려주고, 그 때마다 패키지를 다시 설치할 수 있습니다. 수정된 구성 파일은 재설치 과정에서 모두 보존됩니다.

기존 소스

가장 중요한 설계 목적은 해당 소프트웨어의 개발자에 의하여 배포된 "기존" 소프트웨어 소스를 사용할 수 있게 하는 것이었습니다. RPM에는 기존 소스와 함께 사용된 패치 및 완전한 개발 지시 사항이 포함되어 있습니다. 이것은 여러 가지 이유에서 매우 중요한 장점으로 볼 수 있습니다. 예를 들어, 새로운 버전의 프로그램이 출시될 경우, 프로그램을 컴파일하기 위해 처음부터 다시 시작하실 필요가 없습니다. 패치를 살펴보신 후 어떠한 작업이 필요한지 쉽게 알 수 있습니다. 이 기능을 사용하여 소프트웨어를 제대로 구축하고자 만들어진 컴파일된 기본값과 변경 사항들을 쉽게 보실 수 있습니다.

소스를 원래대로 보존하는 목적은 개발자에게만 중요하다고 느끼실 수도 있지만 결론적으로 최종 사용자에게도 더 높은 수준의 소프트웨어를 가져다 줍니다.

11.2. RPM 사용

RPM has five basic modes of operation (not counting package building): installing, uninstalling, upgrading, querying, and verifying. This section contains an overview of each mode. For complete details and options, try rpm --help or man rpm. You can also refer to 11.5절. "추가 자료" for more information on RPM.

11.2.1. RPM 패키지 검색

RPM 패키지를 사용하기 전에 RPM 패키지가 어디에 있는지 확인해야 합니다. 인터넷 검색을 하면 많은 RPM 레포지토리를 찾을 수 있지만 Red Hat에서 개발된 패키지는 다음 위치에서 찾을 수 있습니다:

- Red Hat Enterprise Linux CD-ROM
- The Red Hat Errata Page available at http://www.redhat.com/apps/support/errata/[1]
- Red Hat Network — Refer to 14장. Product Subscriptions and Entitlements for more details on Red Hat Network.

11.2.2. 설치

RPM packages typically have file names like foo-1.0-1.i386.rpm. The file name includes the package name (foo), version (1.0), release (1), and architecture (i386). To install a package, log in as root and type the following command at a shell prompt:

```
rpm -ivh foo-1.0-1.i386.rpm
```

[1] http://rhn.redhat.com/errata/

Alternatively, the following command can also be used:

```
rpm -Uvh foo-1.0-1.i386.rpm
```

성공적으로 설치되면 다음 내용이 출력됩니다:

```
Preparing...                ############################################ [100%]
   1:foo                     ############################################ [100%]
```

화면에서 보는 것과 같이, RPM은 패키지 이름을 출력하고 패키지가 설치되는 동안 진행 상황을 해시 마크로 표시하여 보여줍니다.

패키지 설치 또는 업그레이드 시 패키지 서명이 자동으로 확인됩니다. 서명은 패키지가 인증된 단체로부터 서명되었는지 확인합니다. 예를 들어, 서명 검증이 실패할 경우, 다음과 같은 오류 메시지가 표시됩니다:

```
error: V3 DSA signature: BAD, key ID 0352860f
```

만일 새 헤더-전용 서명일 경우, 다음과 같은 오류 메시지가 나타납니다:

```
error: Header V3 DSA signature: BAD, key ID 0352860f
```

If you do not have the appropriate key installed to verify the signature, the message contains the word NOKEY such as:

```
warning: V3 DSA signature: NOKEY, key ID 0352860f
```

Refer to 11.3절. "Checking a Package's Signature" for more information on checking a package's signature.

경고

If you are installing a kernel package, you should use rpm -ivh instead. Refer to 42장. Manually Upgrading the Kernel for details.

11.2.2.1. 이미 설치된 패키지

같은 이름과 버전의 패키지가 이미 설치되어 있으면 다음과 같은 메시지가 표시됩니다:

```
Preparing...                ############################################ [100%]
package foo-1.0-1 is already installed
```

However, if you want to install the package anyway, you can use the --replacepkgs option, which tells RPM to ignore the error:

```
rpm -ivh --replacepkgs foo-1.0-1.i386.rpm
```

이 옵션은 RPM에서 설치된 파일이 삭제되었거나 RPM에서 기존 구성 파일을 설치할 때 유용하게 사용됩니다.

11.2.2.2. 파일 충돌

다른 패키지에 의해 이미 설치된 파일을 포함하는 패키지나 같은 패키지의 이전 버전을 설치하려면 다음과 같은 메시지가 나타날 것입니다:

```
Preparing...                ########################################### [100%]
file /usr/bin/foo from install of foo-1.0-1 conflicts with file from package bar-2.0.20
```

To make RPM ignore this error, use the --replacefiles option:

```
rpm -ivh --replacefiles foo-1.0-1.i386.rpm
```

11.2.2.3. 해결되지 않은 의존성 문제

RPM 패키지는 다른 패키지에 종속적일 수 있습니다. 즉, 다른 패키지가 설치되어야 RPM 패키지가 제대로 실행될 수 있다는 것을 의미합니다. 해결되지 않은 의존성을 가진 패키지를 설치하려면 다음과 같은 메시지가 나타날 것입니다:

```
error: Failed dependencies:
        bar.so.2 is needed by foo-1.0-1
Suggested resolutions:
 bar-2.0.20-3.i386.rpm
```

Red Hat Enterprise Linux CD-ROM을 사용하여 패키지를 설치하려면 패키지 간의 의존성 문제를 해결해야 합니다. 필요한 패키지를 Red Hat Enterprise Linux CD-ROM이나 Red Hat Network에서 찾으신 후 다음 명령어를 사용하시기 바랍니다:

```
rpm -ivh foo-1.0-1.i386.rpm bar-2.0.20-3.i386.rpm
```

두 패키지가 성공적으로 설치되면 다음과 같은 결과가 출력될 것입니다:

```
Preparing...                ########################################### [100%]
   1:foo                    ########################################### [ 50%]
   2:bar                    ########################################### [100%]
```

If it does not suggest a package to resolve the dependency, you can try the -q --whatprovides option combination to determine which package contains the required file.

```
rpm -q --whatprovides bar.so.2
```

의존성 문제를 해결하지 않고 설치를 계속 진행하려면 --nodeps 옵션을 사용하시기 바랍니다. (제대로 실행되지 않는 문제가 발생할 수도 있으므로 권장하지 않습니다.)

11.2.3. 제거

패키지 제거는 설치하는 것만큼 간단합니다. 쉘 프롬프트에서 다음 명령을 입력합니다:

```
rpm -e foo
```

알림

Notice that we used the package name foo, not the name of the original package file foo-1.0-1.i386.rpm. To uninstall a package, replace foo with the actual package name of the original package.

제거하려는 패키지에 또 다른 설치된 패키지가 의존하고 있는 경우 패키지 제거 시 의존성 오류가 발생할 수 있습니다. 예를 들어:

```
error: Failed dependencies:
 foo is needed by (installed) bar-2.0.20-3.i386.rpm
```

RPM이 이러한 오류를 무시하고 계속 패키지 삭제 작업을 진행하도록 하시려면 --nodeps 옵션을 사용하시면 됩니다. 하지만, 이 옵션을 사용하시면 패키지가 제대로 실행되지 않을 가능성이 있습니다.

11.2.4. 업그레이드

패키지 업그레이드는 설치하는 것과 비슷합니다. 쉘 프롬프트에서 다음 명령어를 입력합니다:

```
rpm -Uvh foo-2.0-1.i386.rpm
```

As part of upgrading a package, RPM automatically uninstalls any old versions of the foo package. Note that -U will also install a package even when there are no previous versions of the package installed.

Tip

-U 옵션은 이전 패키지를 교체하기 때문에 커널 패키지를 설치할 때는 -U 옵션을 사용하지 않는 것이 좋습니다. 실행 중인 시스템에 영향을 주진 않지만, 시스템을 재시작할 때 새로운 커널을 사용할 수 없으면 대체할 다른 커널이 없으므로 -U 옵션 사용을 지양합니다.

Using the -i option adds the kernel to your GRUB boot menu (/etc/grub.conf). Similarly, removing an old, unneeded kernel removes the kernel from GRUB.

RPM은 구성 파일을 사용하여 지능화된 패키지 업그레이드를 수행합니다. 따라서, 다음과 같은 메시지를 볼 수 있습니다:

```
saving /etc/foo.conf as /etc/foo.conf.rpmsave
```

이 메시지는 사용자가 변경한 구성 파일이 패키지에 있는 새로운 구성 파일과 호환되지 않을 수 있으므로 RPM이 기존 파일을 저장한 후 새로운 구성 파일을 설치했다는 것을 의미합니다. 시스템이 계속해서 제대로 작동할 수 있도록 구성하려면 먼저 두 구성 파일을 비교한 후 차이점을 확인하고 문제를 해결해야 합니다.

이전 버전의 패키지로 업그레이드할 때 (즉, 새로운 버전의 패키지가 이미 설치된 경우), 다음과 같은 메시지를 볼 수 있습니다:

```
package foo-2.0-1 (which is newer than foo-1.0-1) is already installed
```

To force RPM to upgrade anyway, use the --oldpackage option:

```
rpm -Uvh --oldpackage foo-1.0-1.i386.rpm
```

11.2.5. 새로 설치

패키지 새로 설치는 패키지 업그레이드와 비슷합니다. 쉘 프롬프트에서 다음 명령어을 입력합니다:

```
rpm -Fvh foo-1.2-1.i386.rpm
```

RPM's freshen option checks the versions of the packages specified on the command line against the versions of packages that have already been installed on your system. When a newer version of an already-installed package is processed by RPM's freshen option, it is upgraded to the newer version. However, RPM's freshen option does not install a package if no previously-installed package of the same name exists. This differs from RPM's upgrade option, as an upgrade does install packages whether or not an older version of the package was already installed.

RPM의 새로 설치 옵션은 단독 패키지나 패키지 그룹에서 작용합니다. 많은 패키지를 다운로드한 후 시스템에 이미 설치된 패키지만을 업그레이드할 계획이라면, 새로 설치 옵션 사용을 권장합니다. 새로 설치 옵션을 사용하면 이전에 RPM을 사용하여 다운로드 받은 그룹 중에서 원하지 않는 패키지를 직접 삭제하실 필요가 없습니다.

이러한 경우에 간단히 다음과 같은 명령을 사용할 수 있습니다:

```
rpm -Fvh *.rpm
```

RPM은 자동으로 이미 설치된 패키지만을 업그레이드합니다.

11.2.6. 질의

RPM 데이터베이스는 시스템에 설치된 모든 RPM 패키지에 관한 정보를 저장합니다. RPM 데이터베이스는 /var/lib/rpm/ 디렉토리에 저장되며 어느 패키지가 설치되고, 각 패키지 버전이 무엇이며, 설치 후에 패키지의 어느 파일에 어떤 수정 사항이 있었는지에 대한 정보를 찾는 데 사용됩니다.

To query this database, use the -q option. The rpm -q package name command displays the package name, version, and release number of the installed package package name . For example, using rpm -q foo to query installed package foo might generate the following output:

```
foo-2.0-1
```

You can also use the following Package Selection Options with -q to further refine or qualify your query:

- -a — queries all currently installed packages.
- -f <filename> — queries the RPM database for which package owns f<filename> . When specifying a file, specify the absolute path of the file (for example, rpm -qf /bin/ls).

- -p <packagefile> — queries the uninstalled package <packagefile> .

There are a number of ways to specify what information to display about queried packages. The following options are used to select the type of information for which you are searching. These are called Package Query Options.

- -i displays package information including name, description, release, size, build date, install date, vendor, and other miscellaneous information.
- -l displays the list of files that the package contains.
- -s displays the state of all the files in the package.
- -d displays a list of files marked as documentation (man pages, info pages, READMEs, etc.).
- -c displays a list of files marked as configuration files. These are the files you edit after installation to adapt and customize the package to your system (for example, sendmail.cf, passwd, inittab, etc.).

For options that display lists of files, add -v to the command to display the lists in a familiar ls -l format.

11.2.7. 검증

패키지 검증은 패키지에 설치된 파일에 저장된 내용과 기존 패키지의 내용을 비교합니다. 검증 옵션을 사용하면 여러 가지 정보, 즉 개별 파일의 크기, MD5 sum, 권한, 유형, 소유권, 그룹 소유권 등을 비교하게 되며 어떠한 변화가 있을 경우 출력합니다.

The command rpm -V verifies a package. You can use any of the Verify Options listed for querying to specify the packages you wish to verify. A simple use of verifying is rpm -V foo, which verifies that all the files in the foo package are as they were when they were originally installed. For example:

- 특정 파일을 포함하는 패키지를 검증할 때:

  ```
  rpm -Vf /usr/bin/foo
  ```

  예제에서, /usr/bin/foo는 패키지 질의에 사용되는 파일에 대한 절대 경로입니다.

- 시스템에 설치된 모든 패키지 검증:

  ```
  rpm -Va
  ```

- RPM 패키지 파일로 설치된 패키지 검증:

  ```
  rpm -Vp foo-1.0-1.i386.rpm
  ```

  RPM 데이터베이스가 손상되었다고 판단될 때 이 명령어를 사용하여 조사할 수 있습니다.

If everything verified properly, there is no output. If there are any discrepancies, they are displayed. The format of the output is a string of eight characters (a c denotes a configuration file) and then the file name. Each of the eight characters denotes the result of a comparison of one attribute of the file to the value of that attribute recorded in the RPM database. A single period (.) means the test passed. The following characters denote specific discrepancies:

- 5 — MD5 checksum

- S — file size
- L — symbolic link
- T — file modification time
- D — device
- U — user
- G — group
- M — mode (includes permissions and file type)
- ? — unreadable file

문제점이 발견되면 패키지를 제거하거나 재설치할 것인지 또는 다른 방식으로 문제를 해결할 것인지를 잘 결정하셔야 합니다.

11.3. Checking a Package's Signature

If you wish to verify that a package has not been corrupted or tampered with, examine only the md5sum by typing the following command at a shell prompt (where <rpm-file> is the file name of the RPM package):

```
rpm -K --nosignature <rpm-file>
```

The message <rpm-file>: md5 OK is displayed. This brief message means that the file was not corrupted by the download. To see a more verbose message, replace -K with -Kvv in the command.

On the other hand, how trustworthy is the developer who created the package? If the package is signed with the developer's GnuPG key, you know that the developer really is who they say they are.

사용자가 다운로드받은 패키지의 신뢰성 여부를 가려낼 수 있도록 RPM 패키지는 Gnu Privacy Guard(또는 GnuPG)를 사용하여 서명됩니다.

GnuPG는 보안 통신을 위한 도구로서 전자 우편 보안 시스템의 하나인 PGP의 암호화 기술을 대체하는 완전한 기능을 갖춘 프리 소프트웨어입니다. GnuPG를 사용하여 문서의 유효성을 인증하고 다른 수신자와 보내고 받는 데이타를 암호화/해독할 수 있습니다. GnuPG는 또한 PGP 5.x 파일을 해독하고 검증할 수 있습니다.

During installation, GnuPG is installed by default. That way you can immediately start using GnuPG to verify any packages that you receive from Red Hat. Before doing so, you must first import Red Hat's public key.

11.3.1. 키 가져오기

Red Hat 패키지들을 검증하려면 Red Hat GPG 키를 가져와야 합니다. 쉘 프롬프트에서 다음 명령어를 실행하시기 바랍니다:

```
rpm --import /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
```

RPM 검증을 위해 설치된 모든 키 목록을 보시려면 다음 명령을 실행하십시오:

```
rpm -qa gpg-pubkey*
```

Red Hat 키에는 다음과 같은 결과가 출력될 것입니다:

```
gpg-pubkey-37017186-45761324
```

To display details about a specific key, use rpm -qi followed by the output from the previous command:

```
rpm -qi gpg-pubkey-37017186-45761324
```

11.3.2. 패키지 서명 검증

To check the GnuPG signature of an RPM file after importing the builder's GnuPG key, use the following command (replace <rpm-file> with the filename of the RPM package):

```
rpm -K <rpm-file>
```

If all goes well, the following message is displayed: md5 gpg OK. This means that the signature of the package has been verified, and that it is not corrupt.

11.4. RPM 사용법 실습 및 예제

RPM은 시스템을 관리할 뿐만 아니라 문제점을 진단하고 해결하는데 사용되는 유용한 도구입니다. 다음과 같은 몇 가지 예제를 통해 RPM 옵션 사용법을 이해할 수 있습니다.

- 파일을 실수로 삭제했으며 어떤 파일을 삭제했는지 알 수 없는 상황이라고 가정해 봅시다. 전체 시스템을 확인하여 삭제된 파일을 검색하려면 다음과 같은 명령어를 사용할 수 있습니다:

  ```
  rpm -Va
  ```

  일부 파일이 사라졌거나 손상된 것처럼 보이면 패키지를 재설치하거나 제거 후 재설치해야 합니다.

- 작업 중에 알 수 없는 파일을 볼 때가 있습니다. 이 때 다음 명령어를 실행하면 파일이 어느 패키지에 속하는지 찾을 수 있습니다:

  ```
  rpm -qf /usr/bin/ggv
  ```

  출력된 결과는 다음과 비슷하게 나타날 것입니다:

  ```
  ggv-2.6.0-2
  ```

- We can combine the above two examples in the following scenario. Say you are having problems with /usr/bin/paste. You would like to verify the package that owns that program, but you do not know which package owns paste. Enter the following command,

  ```
  rpm -Vf /usr/bin/paste
  ```

해당 패키지가 검증될 것입니다.

- 특정 프로그램에 대한 더 많은 정보를 찾고자 합니까? 다음 명령어로 프로그램을 소유한 패키지에 있는 문서 자료를 찾을 수 있습니다:

```
rpm -qdf /usr/bin/free
```

다음과 같이 출력될 것입니다:

```
/usr/share/doc/procps-3.2.3/BUGS
/usr/share/doc/procps-3.2.3/FAQ
/usr/share/doc/procps-3.2.3/NEWS
/usr/share/doc/procps-3.2.3/TODO
/usr/share/man/man1/free.1.gz
/usr/share/man/man1/pgrep.1.gz
/usr/share/man/man1/pkill.1.gz
/usr/share/man/man1/pmap.1.gz
/usr/share/man/man1/ps.1.gz
/usr/share/man/man1/skill.1.gz
/usr/share/man/man1/slabtop.1.gz
/usr/share/man/man1/snice.1.gz
/usr/share/man/man1/tload.1.gz
/usr/share/man/man1/top.1.gz
/usr/share/man/man1/uptime.1.gz
/usr/share/man/man1/w.1.gz
/usr/share/man/man1/watch.1.gz
/usr/share/man/man5/sysctl.conf.5.gz
/usr/share/man/man8/sysctl.8.gz
/usr/share/man/man8/vmstat.8.gz
```

- 새로운 RPM을 찾았지만 그 기능을 알 수 없을 때는 다음 명령어를 사용하여 RPM에 대한 정보를 찾아볼 수 있습니다:

```
rpm -qip crontabs-1.10-7.noarch.rpm
```

다음과 같이 출력될 것입니다:

```
Name        : crontabs                     Relocations: (not relocatable)
Version     : 1.10                              Vendor: Red Hat, Inc.
Release     : 7                             Build Date: Mon 20 Sep 2004 05:58:10 PM EDT
Install Date: (not installed)               Build Host: tweety.build.redhat.com
Group       : System Environment/Base       Source RPM: crontabs-1.10-7.src.rpm
Size        : 1004                             License: Public Domain
Signature   : DSA/SHA1, Wed 05 Jan 2005 06:05:25 PM EST, Key ID 219180cddb42a60e
Packager    : Red Hat, Inc. <http://bugzilla.redhat.com/bugzilla>
Summary     : Root crontab files used to schedule the execution of programs.
Description : The crontabs package contains root crontab files. Crontab is the
program used to install, uninstall, or list the tables used to drive the
cron daemon. The cron daemon checks the crontab files to see when
particular commands are scheduled to be executed. If commands are
scheduled, then it executes them.
```

- Perhaps you now want to see what files the crontabs RPM installs. You would enter the following:

```
rpm -qlp crontabs-1.10-5.noarch.rpm
```

출력된 결과는 다음과 같이 나타날 것입니다:

```
/etc/cron.daily
/etc/cron.hourly
/etc/cron.monthly
/etc/cron.weekly
/etc/crontab
/usr/bin/run-parts
```

앞에서 설명된 것은 RPM의 많은 기능 중 극히 소수에 불과합니다. RPM을 사용하면 할수록 훨씬 더 많은 기능이 있다는 것을 발견하실 것입니다.

11.5. 추가 자료

RPM은 많은 옵션과 다양한 패키지를 질의, 설치, 업그레이드 및 삭제 방식을 갖춘 매우 복잡한 유틸리티입니다. RPM에 대하여 더 많은 정보를 원하시면 다음과 같은 자료를 참조하시기 바랍니다.

11.5.1. 설치된 문서 자료

- rpm --help — This command displays a quick reference of RPM parameters.
- man rpm — The RPM man page gives more detail about RPM parameters than the rpm --help command.

11.5.2. 유용한 웹사이트

- http://www.rpm.org/ — RPM 웹사이트.
- https://lists.rpm.org/mailman/listinfo/rpm-list[2] — Visit this link to subscribe to the RPM mailing list, which is archived there.

11.5.3. 관련 서적

- The Red Hat RPM Guide by Eric Foster-Johnson is an excellent resource on all details of the RPM package format and the RPM package management utility. It is available online at http://docs.fedoraproject.org/drafts/rpm-guide-en/.

[2] http://www.redhat.com/mailman/listinfo/rpm-list/

Package Management Tool

If you prefer to use a graphical interface to view and manage packages in your system, you can use the Package Management Tool, better known as pirut. This tool allows you to perform basic package management of your system through an easy-to-use interface to remove installed packages or download (and install) packages compatible to your system. It also allows you to view what packages are installed in your system and which ones are available for download from Red Hat Network. In addition, the Package Management Tool also automatically resolves any critical dependencies when you install or remove packages in the same way that the rpm command does.

알림

While the Package Management Tool can automatically resolve dependencies during package installation and removal, it cannot perform a forced install / remove the same way that rpm -e --nodeps or rpm -U --nodeps can.

The X Window System is required to run the Package Management Tool. To start the application, go to Applications (the main menu on the panel) > Add/Remove Software. Alternatively, you can type the commands system-config-packages or pirut at shell prompt.

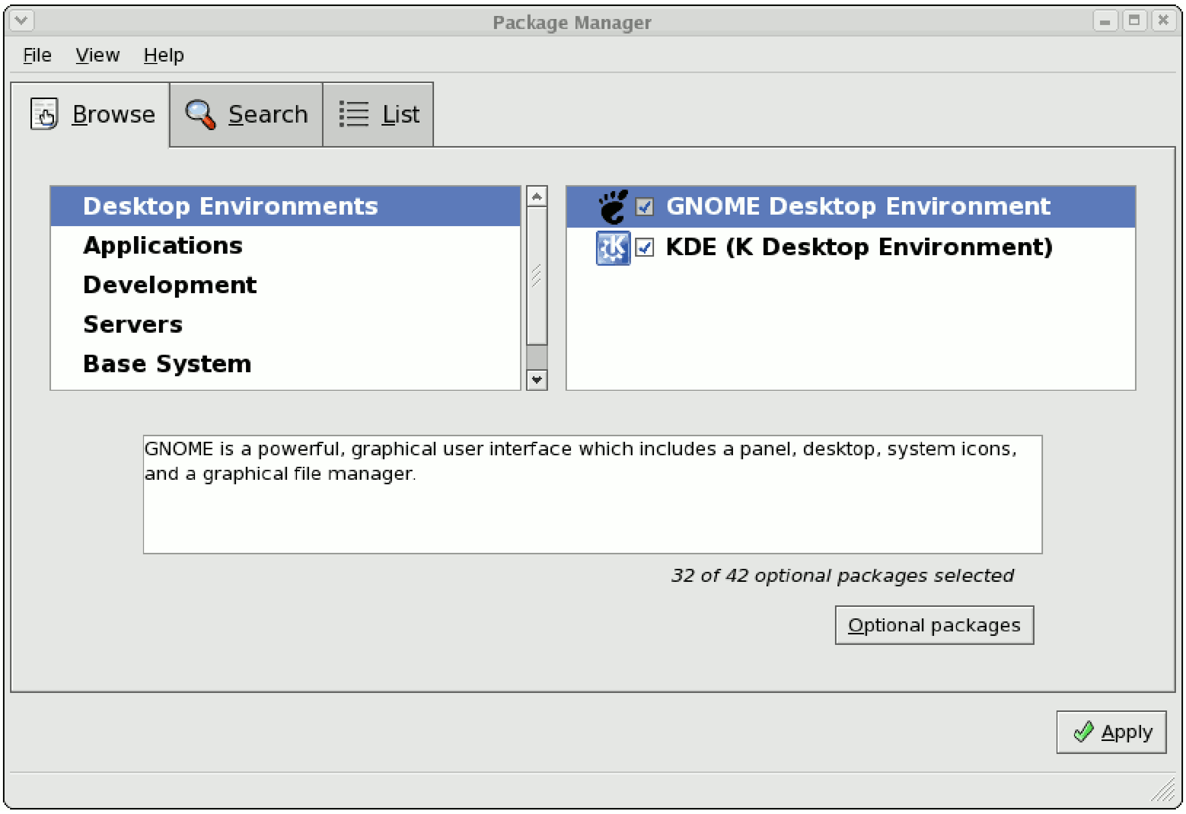

Desktop Environments
Applications
Development
Servers
Base System

☑ GNOME Desktop Environment
☑ KDE (K Desktop Environment)

GNOME is a powerful, graphical user interface which includes a panel, desktop, system icons, and a graphical file manager.

32 of 42 optional packages selected

그림 12.1. Package Management Tool

12.1. 패키지 목록 및 분석

You can use the Package Management Tool to search and list all packages installed in your system, as well as any packages available for you to download. The Browse, Search, and List tabs present different options in viewing, analyzing, installing or removing packages.

The Browse tab allows you to view packages by group. In 그림 12.1. "Package Management Tool", the left window shows the different package group types you can choose from (for example, Desktop Environments, Applications, Development and more). When a package group type is selected, the right window displays the different package groups of that type.

어떤 패키지가 패키지 그룹에 포함되었는 지를 보시려면 옵션 패키지O 버튼을 클릭하시기 바랍니다. 설치된 패키지를 확인하게 됩니다.

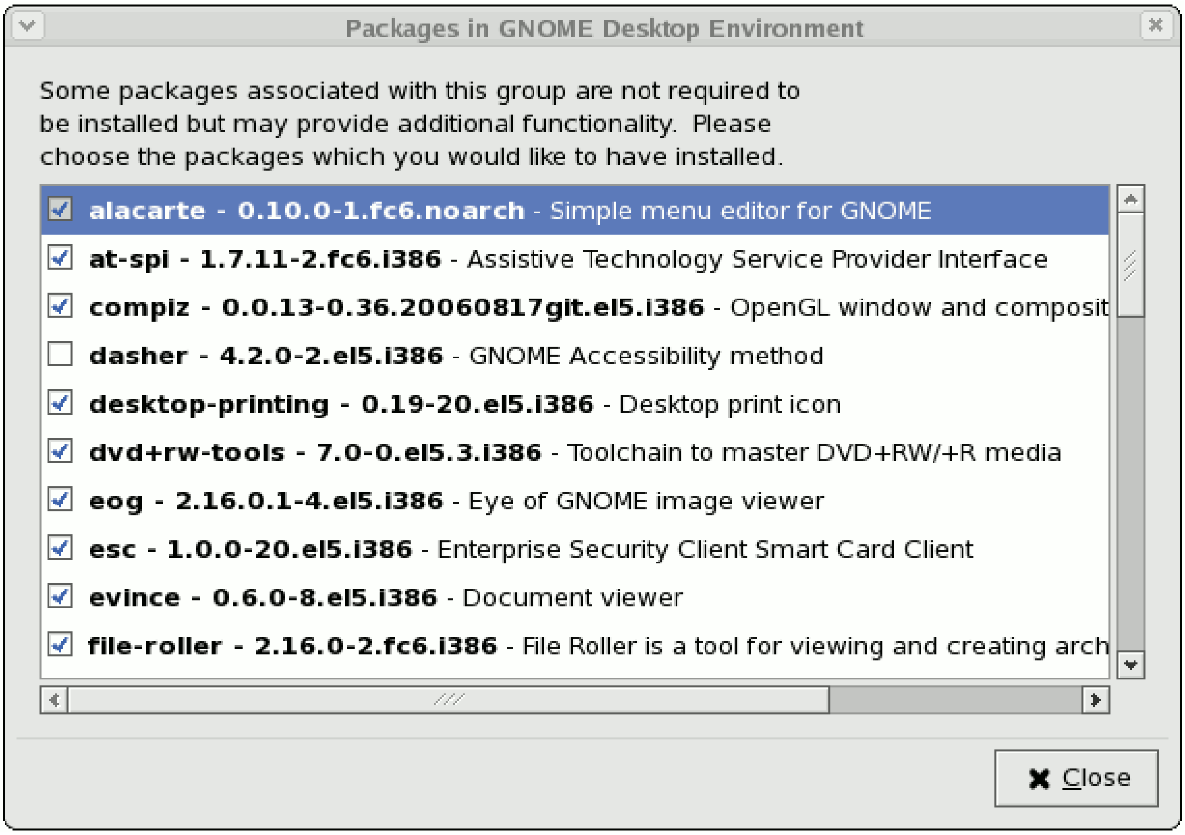

Packages in GNOME Desktop Environment

Some packages associated with this group are not required to be installed but may provide additional functionality. Please choose the packages which you would like to have installed.

☑ **alacarte - 0.10.0-1.fc6.noarch** - Simple menu editor for GNOME
☑ **at-spi - 1.7.11-2.fc6.i386** - Assistive Technology Service Provider Interface
☑ **compiz - 0.0.13-0.36.20060817git.el5.i386** - OpenGL window and composit
□ **dasher - 4.2.0-2.el5.i386** - GNOME Accessibility method
☑ **desktop-printing - 0.19-20.el5.i386** - Desktop print icon
☑ **dvd+rw-tools - 7.0-0.el5.3.i386** - Toolchain to master DVD+RW/+R media
☑ **eog - 2.16.0.1-4.el5.i386** - Eye of GNOME image viewer
☑ **esc - 1.0.0-20.el5.i386** - Enterprise Security Client Smart Card Client
☑ **evince - 0.6.0-8.el5.i386** - Document viewer
☑ **file-roller - 2.16.0-2.fc6.i386** - File Roller is a tool for viewing and creating arch

그림 12.2. Optional Packages

The List tab displays a list of packages installed or available for download. Packages already installed in your system are marked with a green check (✓).

디폴트로 주 화면 위에 있는 모든 패키지A 옵션이 선택되어 모든 패키지를 보여줍니다. 설치된 패키지I 옵션으로 시스템에 설치된 패키지만을 보실 수 있으며 사용 가능한 패키지v 옵션으로 어떤 패키지를 다운로드 및 설치할 수 있는 지를 확인하실 수 있습니다.

검색S 탭으로 특정 패키지 검색에 필요한 키워드를 사용하실 수 있으며 패키지 요약 설명 또한 보실 수 있습니다. 패키지를 선택하신 후 주 화면 아래의 패키지 정보 D 버튼을 클릭하시면 됩니다.

12.2. 패키지 설치 및 삭제

To install a package available for download, click the checkbox beside the package name. When you do so, an installation icon ( ) appears beside its checkbox. This indicates that the package is queued for download and installation. You can select multiple packages to download and install; once you have made your selection, click the Apply button.

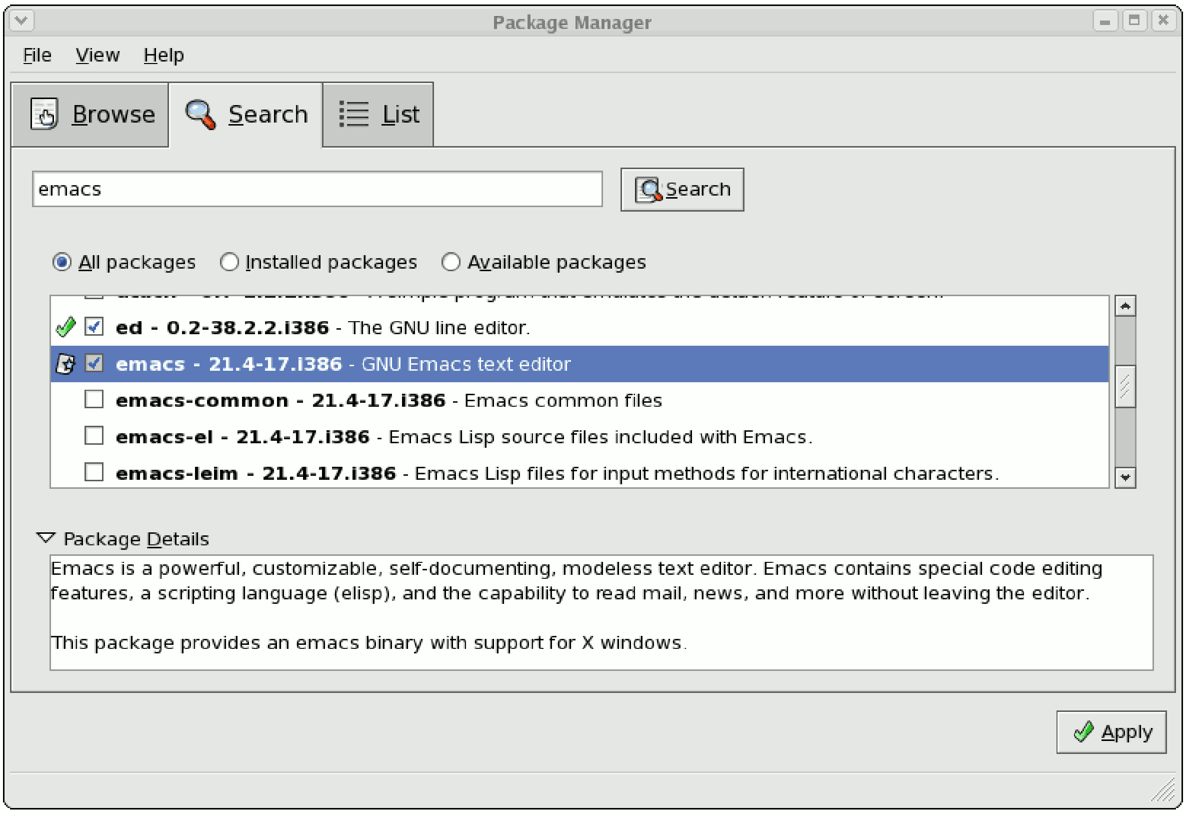

emacs

● All packages ○ Installed packages ○ Available packages

✓ ☑ **ed - 0.2-38.2.2.i386** - The GNU line editor.

☑ **emacs - 21.4-17.i386** - GNU Emacs text editor

☐ **emacs-common - 21.4-17.i386** - Emacs common files

☐ **emacs-el - 21.4-17.i386** - Emacs Lisp source files included with Emacs.

☐ **emacs-leim - 21.4-17.i386** - Emacs Lisp files for input methods for international characters.

Emacs is a powerful, customizable, self-documenting, modeless text editor. Emacs contains special code editing features, a scripting language (elisp), and the capability to read mail, news, and more without leaving the editor.

This package provides an emacs binary with support for X windows.

그림 12.3. Package installation

If there are any package dependencies for your selected downloads, the Package Management Tool will notify you accordingly. Click Details to view what additional packages are needed. To proceed with downloading and installing the package (along with all other dependent packages) click Continue.

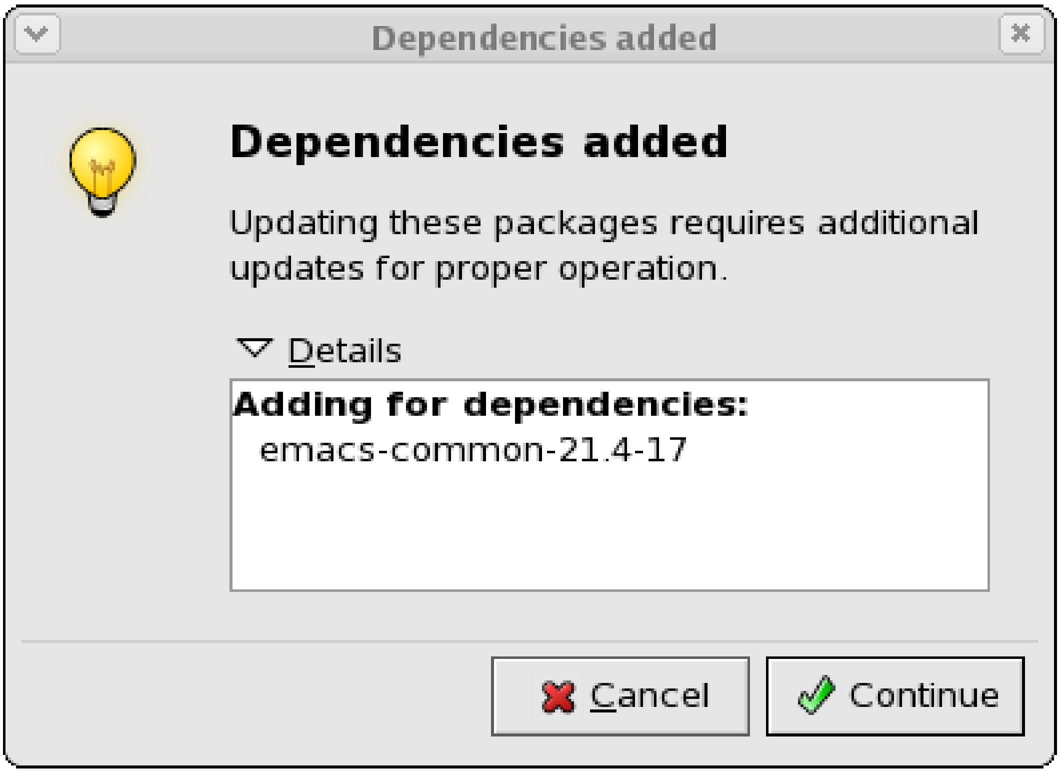

Dependencies added

Updating these packages requires additional updates for proper operation.

Adding for dependencies:
emacs-common-21.4-17

그림 12.4. Package dependencies: installation

Removing a package can be done in a similar manner. To remove a package installed in your system, click the checkbox beside the package name. The green check appearing beside the package name will be replaced by a package removal icon (🗑). This indicates that the package is queued for removal; you can also select multiple packages to be removed at the same time. Once you have selected the packages you want to remove, click the Apply button.

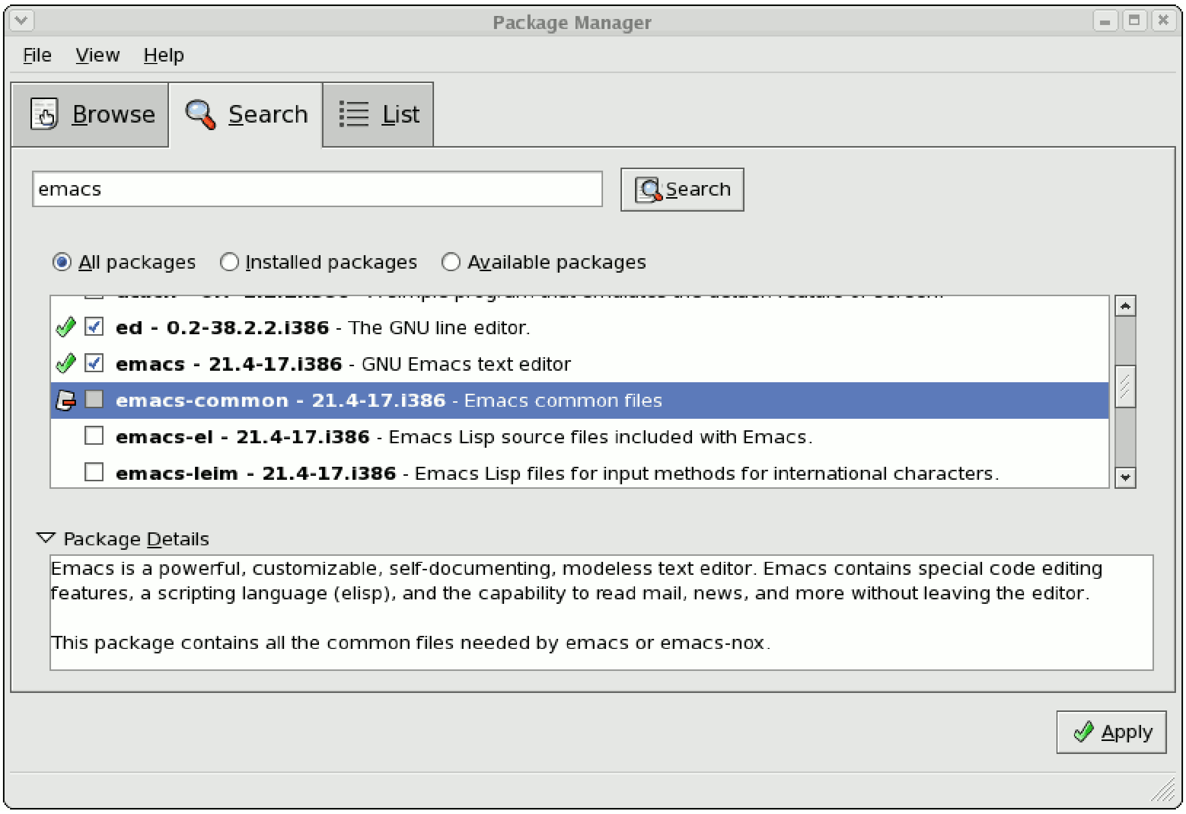

emacs

● All packages ○ Installed packages ○ Available packages

☑ **ed - 0.2-38.2.2.i386** - The GNU line editor.

☑ **emacs - 21.4-17.i386** - GNU Emacs text editor

□ **emacs-common - 21.4-17.i386** - Emacs common files

□ **emacs-el - 21.4-17.i386** - Emacs Lisp source files included with Emacs.

□ **emacs-leim - 21.4-17.i386** - Emacs Lisp files for input methods for international characters.

Emacs is a powerful, customizable, self-documenting, modeless text editor. Emacs contains special code editing features, a scripting language (elisp), and the capability to read mail, news, and more without leaving the editor.

This package contains all the common files needed by emacs or emacs-nox.

그림 12.5. Package removal

Note that if any other installed packages are dependent on the package you are removing, they will be removed as well. The Package Management Tool will notify you if there are any such dependencies. Click Details to view what packages are dependent on the one you are removing. To proceed with removing your selected package/s (along with all other dependent packages) click Continue.

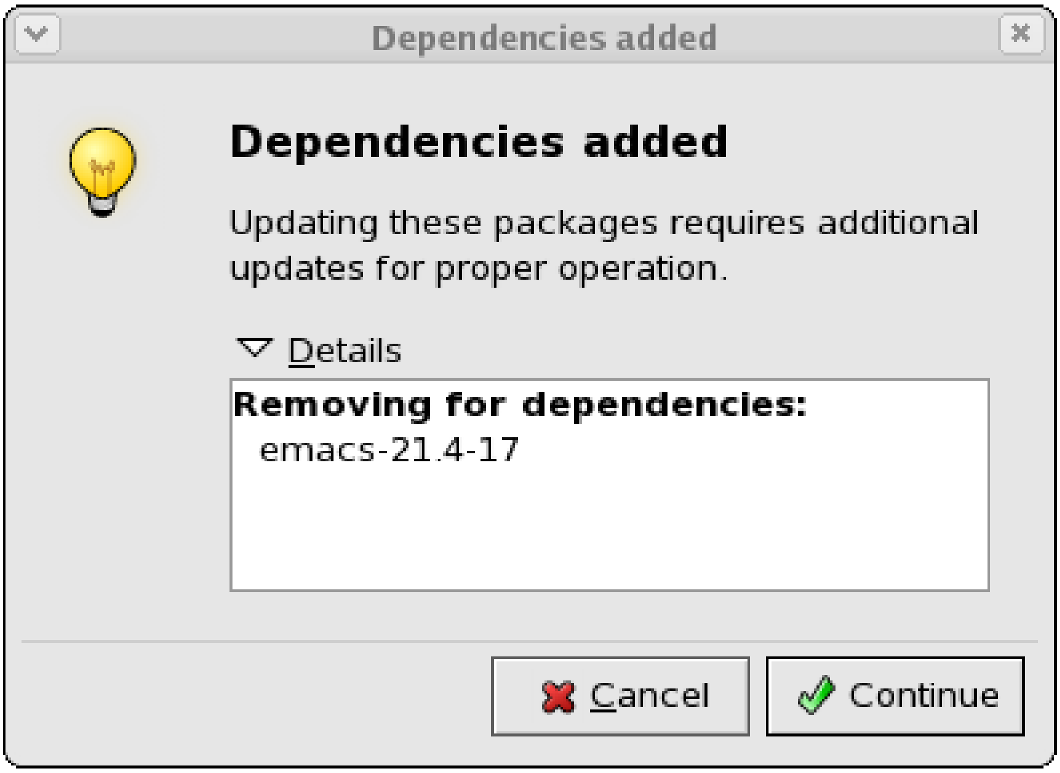

Dependencies added

Updating these packages requires additional updates for proper operation.

Removing for dependencies:
  emacs-21.4-17

그림 12.6. Package dependencies: removal

You can install and remove multiple packages by selecting packages to be installed / removed and then clicking Apply. The Package selections window displays the number of packages to be installed and removed.

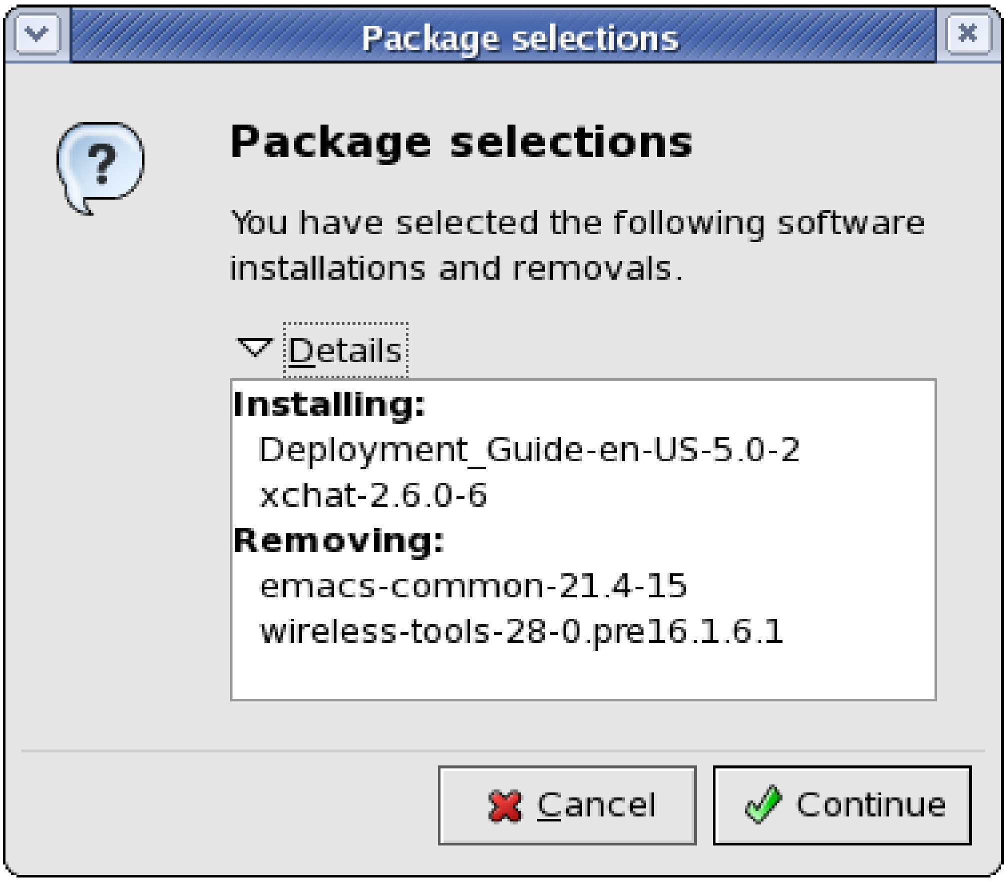

Package selections

You have selected the following software installations and removals.

Installing:
Deployment_Guide-en-US-5.0-2
xchat-2.6.0-6
Removing:
emacs-common-21.4-15
wireless-tools-28-0.pre16.1.6.1

그림 12.7. Installing and removing packages simultaneously

YUM (Yellowdog Updater Modified)

Yellowdog Update, Modified (YUM) is a package manager that was developed by Duke University to improve the installation of RPMs. yum searches numerous repositories for packages and their dependencies so they may be installed together in an effort to alleviate dependency issues. Red Hat Enterprise Linux 5.8 uses yum to fetch packages and install RPMs.

up2date is now deprecated in favor of yum (Yellowdog Updater Modified). The entire stack of tools which installs and updates software in Red Hat Enterprise Linux 5.8 is now based on yum. This includes everything, from the initial installation via Anaconda to host software management tools like pirut.

yum also allows system administrators to configure a local (i.e. available over a local network) repository to supplement packages provided by Red Hat. This is useful for user groups that use applications and packages that are not officially supported by Red Hat.

Aside from being able to supplement available packages for local users, using a local yum repository also saves bandwidth for the entire network. Further, clients that use local yum repositories do not need to be registered individually to install or update the latest packages from Red Hat Network.

13.1. Setting Up a Yum Repository

To set up a repository for Red Hat Enterprise Linux packages, follow these steps:

1. Install the createrepo package:

   ```
   ~]# yum install createrepo
   ```

2. Copy all the packages you want to provide in the repository into one directory (/mnt/local_repo for example).

3. Run createrepo on that directory (for example, createrepo /mnt/local_repo). This will create the necessary metadata for your Yum repository.

13.2. yum Commands

yum commands are typically run as yum <command> <package name/s> . By default, yum will automatically attempt to check all configured repositories to resolve all package dependencies during an installation/upgrade.

The following is a list of the most commonly-used yum commands. For a complete list of available yum commands, refer to man yum.

yum install <package name/s>
: Used to install the latest version of a package or group of packages. If no package matches the specified package name(s), they are assumed to be a shell glob, and any matches are then installed.

yum update <package name/s>
: Used to update the specified packages to the latest available version. If no package name/s are specified, then yum will attempt to update all installed packages.

  If the --obsoletes option is used (i.e. yum --obsoletes <package name/s> , yum will process obsolete packages. As such, packages that are obsoleted across updates will be removed and replaced accordingly.

yum check-update

This command allows you to determine whether any updates are available for your installed packages. yum returns a list of all package updates from all repositories if any are available.

yum remove <package name/s>

Used to remove specified packages, along with any other packages dependent on the packages being removed.

yum provides <file name>

Used to determine which packages provide a specific file or feature.

yum search <keyword>

This command is used to find any packages containing the specified keyword in the description, summary, packager and package name fields of RPMs in all repositories.

yum localinstall <absolute path to package name/s>

Used when using yum to install a package located locally in the machine.

13.3. yum Options

yum options are typically stated before specific yum commands; i.e. yum <options> <command> <package name/s> . Most of these options can be set as default using the configuration file.

The following is a list of the most commonly-used yum options. For a complete list of available yum options, refer to man yum.

-y

Answer "yes" to every question in the transaction.

-t

Sets yum to be "tolerant" of errors with regard to packages specified in the transaction. For example, if you run yum update package1 package2 and package2 is already installed, yum will continue to install package1.

--exclude=<package name>

Excludes a specific package by name or glob in a specific transaction.

13.4. Configuring yum

By default, yum is configured through /etc/yum.conf. The following is an example of a typical /etc/yum.conf file:

```
[main]
cachedir=/var/cache/yum
keepcache=0
debuglevel=2
logfile=/var/log/yum.log
distroverpkg=redhat-release
tolerant=1
exactarch=1
obsoletes=1
gpgcheck=1
plugins=1
metadata_expire=1800
[myrepo]
name=RHEL 5 $releasever - $basearch
baseurl=http://local/path/to/yum/repository/
```

```
enabled=1
```

A typical /etc/yum.conf file is made up of two types of sections: a [main] section, and a repository section. There can only be one [main] section, but you can specify multiple repositories in a single /etc/yum.conf.

13.4.1. [main] Options

The [main] section is mandatory, and there must only be one. For a complete list of options you can use in the [main] section, refer to man yum.conf.

The following is a list of the most commonly-used options in the [main] section.

cachedir

This option specifies the directory where yum should store its cache and database files. By default, the cache directory of yum is /var/cache/yum.

keepcache=<1 or 0>

Setting keepcache=1 instructs yum to keep the cache of headers and packages after a successful installation. keepcache=1 is the default.

reposdir=<absolute path to directory of .repo files>

This option allows you to specify a directory where .repo files are located. .repo files contain repository information (similar to the [repository] section of /etc/yum.conf).

yum collects all repository information from .repo files and the [repository] section of the /etc/yum.conf file to create a master list of repositories to use for each transaction. Refer to 13.4.2 절. “[repository] Options” for more information about options you can use for both the [repository] section and .repo files.

If reposdir is not set, yum uses the default directory /etc/yum.repos.d.

gpgcheck=<1 or 0>

This disables/enables GPG signature checking on packages on all repositories, including local package installation. The default is gpgcheck=0, which disables GPG checking.

If this option is set in the [main] section of the /etc/yum.conf file, it sets the GPG checking rule for all repositories. However, you can also set this on individual repositories instead; i.e., you can enable GPG checking on one repository while disabling it on another.

assumeyes=<1 or 0>

This determines whether or not yum should prompt for confirmation of critical actions. The default if assumeyes=0, which means yum will prompt you for confirmation.

If assumeyes=1 is set, yum behaves in the same way that the command line option -y does.

tolerant=<1 or 0>

When enabled (tolerant=1), yum will be tolerant of errors on the command line with regard to packages. This is similar to the yum command line option -t.

The default value for this is tolerant=0 (not tolerant).

exclude=<package name/s>

This option allows you to exclude packages by keyword during installation/updates. If you are specifying multiple packages, this is a space-delimited list. Shell globs using wildcards (for example, * and ?) are allowed.

retries=<number of retries>

This sets the number of times yum should attempt to retrieve a file before returning an error. Setting this to 0 makes yum retry forever. The default value is 6.

13.4.2. [repository] Options

The [repository] section of the /etc/yum.conf file contains information about a repository yum can use to find packages during package installation, updating and dependency resolution. A repository entry takes the following form:

```
[repository ID]
name=repository name
baseurl=url, file or ftp://path to repository
```

You can also specify repository information in a separate .repo files (for example, rhel5.repo). The format of repository information placed in .repo files is identical with the [repository] of /etc/yum.conf.

.repo files are typically placed in /etc/yum.repos.d, unless you specify a different repository path in the [main] section of /etc/yum.conf with reposdir=. .repo files and the /etc/yum.conf file can contain multiple repository entries.

Each repository entry consists of the following mandatory parts:

[repository ID]

The repository ID is a unique, one-word string that serves as a repository identifier.

name=repository name

This is a human-readable string describing the repository.

baseurl=http, file or ftp://path

This is a URL to the directory where the repodatadirectory of a repository is located. If the repository is local to the machine, use baseurl=file://path to local repository . If the repository is located online using HTTP, use baseurl=http://link . If the repository is online and uses FTP, use baseurl=ftp://link .

If a specific online repository requires basic HTTP authentication, you can specify your username and password in the baseurl line by prepending it as username:password@link. For example, if a repository on http://www.example.com/repo/ requires a username of "user" and a password os "password", then the baseurl link can be specified as baseurl=http://user:password@www.example.com/repo/.

The following is a list of options most commonly used in repository entries. For a complete list of repository entries, refer to man yum.conf.

gpgcheck=<1 or 0>

This disables/enables GPG signature checking a specific repository. The default is gpgcheck=0, which disables GPG checking.

gpgkey=URL

This option allows you to point to a URL of the ASCII-armoured GPG key file for a repository. This option is normally used if yum needs a public key to verify a package and the required key was not imported into the RPM database.

If this option is set, yum will automatically import the key from the specified URL. You will be prompted before the key is installed unless you set assumeyes=1 (in the [main] section of /etc/yum.conf) or -y (in a yum transaction).

exclude=<package name/s>
: This option is similar to the exclude option in the [main] section of /etc/yum.conf. However, it only applies to the repository in which it is specified.

includepkgs=<package name/s>
: This option is the opposite of exclude. When this option is set on a repository, yum will only be able to see the specified packages in that repository. By default, all packages in a repository are visible to yum.

13.5. Useful yum Variables

The following is a list of variables you can use for both yum commands and yum configuration files (i.e. /etc/yum.conf and .repo files).

$releasever
: This is replaced with the package's version, as listed in distroverpkg. This defaults to the version of the redhat-release package.

$arch
: This is replaced with your system's architecture, as listed by os.uname() in Python.

$basearch
: This is replaced with your base architecture. For example, if $arch=i686 then $basearch=i386.

$YUM0-9
: This is replaced with the value of the shell environment variable of the same name. If the shell environment variable does not exist, then the configuration file variable will not be replaced.

Product Subscriptions and Entitlements

Effective asset management requires a mechanism to handle the software inventory — both the type of products and the number of systems that the software is installed on. The subscription service provides that mechanism and gives transparency into both global allocations of subscriptions for an entire organization and the specific subscriptions assigned to a single system.

Red Hat Subscription Manager works with yum to unite content delivery with subscription management. The Subscription Manager handles only the subscription-system associations. yum or other package management tools handle the actual content delivery. 13장. YUM (Yellowdog Updater Modified) describes how to use yum.

This chapter provides an overview of subscription management in Red Hat Enterprise Linux and the Red Hat Subscription Manager tools which are available.

14.1. An Overview of Managing Subscriptions and Content

Red Hat Enterprise Linux and other Red Hat products are sold through subscriptions, which make packages available and provide support for a set number of systems. Subscription management clarifies the relationships between local systems and available software resources because it gives a view into where software subscriptions are assigned, apart from installing the packages.

14.1.1. The Purpose of Subscription Management

New government and industry regulations are setting new mandates for businesses to track how their infrastructure assets are used. These changes include legislation like Sarbanes-Oxley in the United States, standards like Payment Card Industry Data Security Standard (PCI-DSS), or accreditation like SAS-70. Software inventory maintenance is increasingly important to meet accounting and governmental standards.

That means that there is increasing pressure on IT administrators to have an accurate, current accounting of the software used on their systems. Generally, this is called software license management; with Red Hat's subscription model, this is subscription management.

An IT infrastructure tries to maintain parity between the products that are installed and the licensese or subscriptions that those products require. For example, if an IT environment has four servers running Red Hat Enterprise Linux, then that environment must have four active subscriptions for Red Hat Enterprise Linux. If a new server is added to the infrastructure or one of the subscriptions expired, then the infrastructure would have more installed products than it has subscriptions.

Effective subscription management helps organizations achieve four primary goals:

- Maintain regulatory compliance. One of the key responsibilities of administrators is software compliance in conformance with legal or industry requirements. Subscription management helps track both subscription assignments and contract expiration, which helps administrators manage both systems and software inventories in accordance to their regulatory requirements.

- Simplify IT audits. Having a central and clear inventory of both current subscriptions and current systems, IT administrators can monitor and report on their infrastructure better.

- Get better performance by doing better at assigning subscriptions. The subscription service maintains dual inventories of available product subscriptions and registered server systems, with clear associations between subscriptions and systems. This makes it easier for IT administrators to assign relevant subscriptions to systems, because they have a view of what is in the inventory and what the system is currently subscribed to.

- Lower costs and streamline procurement. While under-subscribing systems can run afoul of regulations, over- subscribing systems can cause a significant impact on IT budgets. Subscription management helps subscriptions be assigned most efficiently, so costs could actually be lowered.

With Red Hat's commitment to free and open software, subscription management is focused on delivering tools that help IT administrators monitor their software/systems inventory for their own benefit. Subscription management does not enforce or restrict access to products.

Important

Most Red Hat products are licensed under a GNU General Public License (GPL), which allows free use of the software or code; this is a different license than the Red Hat license agreement. A Red Hat license provides access to Red Hat services, like the Customer Portal and content delivery network.

The Red Hat subscription requires that, as long as there is any active subscription for a product, then every system which uses the Red Hat product must have an active subscription assigned to it. Otherwise, the subscription is violated. See http://www.redhat.com/subscriptions/ and http://www.redhat.com/rhel/renew/faqs/#6 for more information on Red Hat's subscription model and terms.

14.1.2. Defining Subscriptions, Entitlements, and Products

The basis of everything is a subscription. A subscription contains both the products that are available, the support levels, and the quantities, or number of servers, that the product can be installed on.

Subscriptions are managed though the Certificate-based Red Hat Network service, which ties into the subscription service and content delivery network (CDN).

The subscription service maintains a complete list of subscriptions for an organization, identified by a unique ID (called a pool ID). A system is registered, or added, to the subscription service to allow it to manage the subscriptions for that system. Like the subscription, the system is also added to the subscription service inventory and is assigned a unique ID within the service. The subscriptions and system entries, together, comprise the inventory.

A system allocates one of the quantities of a product in a subscription to itself. When a subscription is consumed, it is an entitlement. (An entitlement is roughly analogous to a user license, in that it grants all of the rights to that product to that system. Unlike a user license, an entitlement does not grant the right to use the software; with the subscription model, an entitlement grants the ability to download the packages and receive updates.) Because the available quantity in a subscription lowers once a system subscribes to it, the system consumes the subscription.

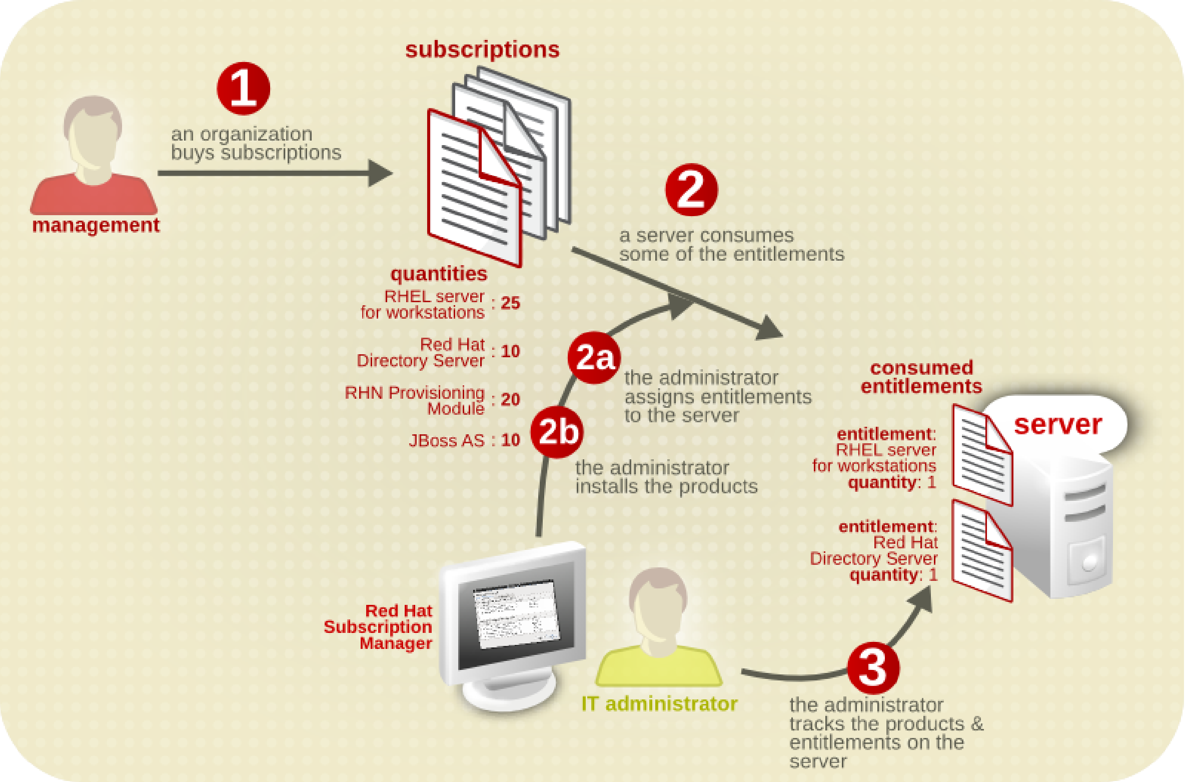
subscriptions
1
an organization buys subscriptions
management
2
a server consumes some of the entitlements
quantities
RHEL server for workstations : 25
Red Hat Directory Server : 10
RHN Provisioning Module : 20
JBoss AS : 10
2a
the administrator assigns entitlements to the server
2b
the administrator installs the products
consumed entitlements
server
entitlement: RHEL server for workstations quantity: 1
entitlement: Red Hat Directory Server quantity: 1
Red Hat Subscription Manager
IT administrator
3
the administrator tracks the products & entitlements on the server

그림 14.1. Subscription Lifecycle, Illustrated

The repository where the product software is located is organized according to the product. Each product group within the repository may contain the primary software packages and then any required dependencies or associated packages. Altogether, the product and its associated packages are called a content set. (A content set for a product even includes other versions of the product.) When a subscription grants access to a product, it includes access to all of the associated packages in that content set.

A single subscription can have multiple products, and each system can have multiple different subscriptions, depending on how many entitlement certificates are loaded on the machine.

Any number of products, for any number of different architectures, can be contained in a single subscription. The subscription options that are visible to a consumer are filtered, by default, according to whether the architecture for the product matches the architecture of the system. This is compatibility. Depending on compatible subscriptions makes sure that subscriptions are allocated efficiently, only to systems which can actually use the products.

Some subscriptions define some element count on the consumer, like the number of sockets on the machine, the number of virtual guests on a host, or the number of clients in a domain. Multiple subscriptions can be combined together to cover the counts on the consumer. For example, if there is a four socket server, two subscriptions for "RHEL Server for Two Sockets" can be consumed by the system to cover the socket count. Combining multiple subscriptions to cover the system count is called stacking.

The subscription tools can display even incompatible entitlements. Alternatively, the architecture definition for the system can be overridden by defining custom system facts for the subscription tools to use.

It is important to distinguish between subscribing to a product and installing a product. A subscription is essentially a statement of whatever products an organization has purchased. The act of subscribing to a subscription means that a system is allowed to install the product with a valid certificate, but subscribing does not actually perform any installation or updates. In the reverse, a product can also be installed apart from any entitlements for the system; the system does not require a valid certificate to install a product. Certificate-based Red Hat Network and the content delivery network harmonize with content delivery and installation by using yum plug-ins that come with the Subscription Manager tools.

14.1.3. Subscription Management Tools

Subscriptions are managed on the local system through GUI and CLI tools called Red Hat Subscription Manager. The Subscription Manager tracks and displays what entitlements are available to the local system and what entitlements have been consumed by the local system. The Subscription Manager works as a conduit back to the subscription service to synchronize changes like available product quantities or subscription expiration and renewals.

> **Note**
>
> The Red Hat Subscription Manager tools are always run as root because of the nature of the changes to the system. However, Red Hat Subscription Manager connects to the subscription service as a user account for the Customer Service Portal.

The Subscription Manager handles both registration and subscriptions for a system. The Subscription Manager is part of the firstboot process for configuring content and updates, but the system can be registered at any time through the Red Hat Subscription Manager GUI or CLI. New subscriptions, new products, and updates can be viewed and applied to a system through the Red Hat Subscription Manager tools.

The different Subscription Manager clients are covered in 14.2절. "Using Red Hat Subscription Manager Tools" .

14.1.4. Subscription and Content Architecture

Content includes new downloads, ISOs, updates, and errata, anything that can be installed on a system.

Subscription management helps to clarify and to define the relationships between local server infrastructure and the content delivery systems. Subscription management and content delivery are tightly associated. Entitlements (assigned subscriptions) identify what a system is allowed to install and update. In other words, entitlements define access to content. The content delivery system actually provides the software packages.

There are three parties that are involved in subscriptions and content:

- The subscription service
- The content delivery network
- The system which uses the content

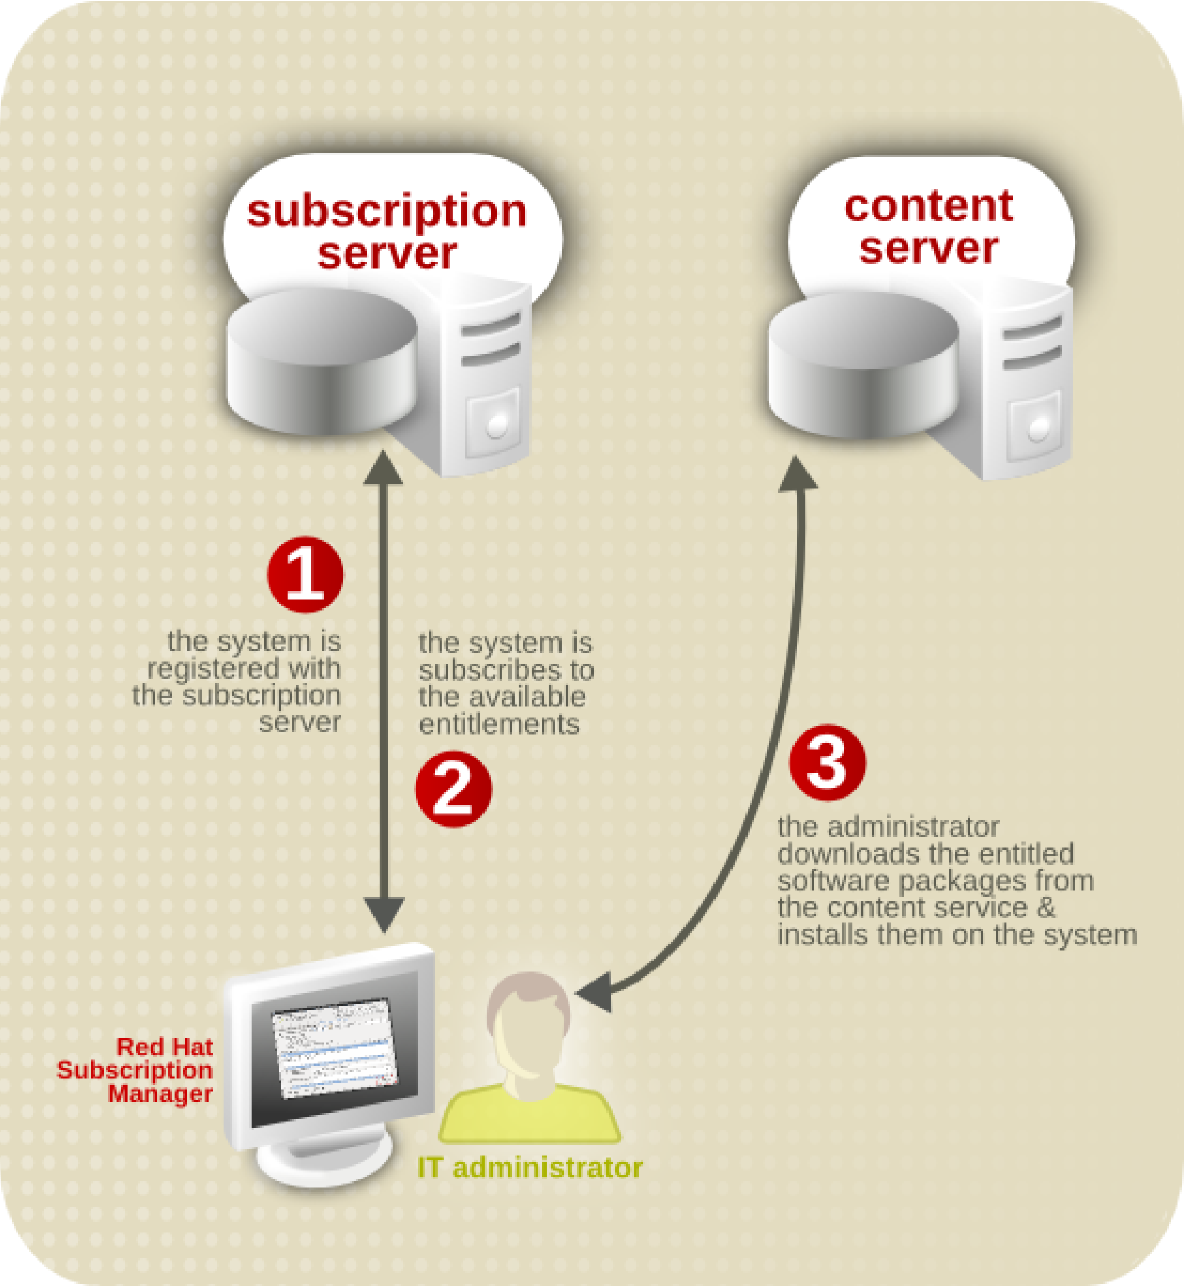
subscription server
content server
1
the system is registered with the subscription server
the system is subscribes to the available entitlements
2
3
the administrator downloads the entitled software packages from the content service & installs them on the system
Red Hat Subscription Manager
IT administrator

그림 14.2. Relationship Among Systems, the Subscription Service, and Content Delivery Network

The subscription service handles the system registration (verifying that the system is allowed to access the content). It also supplies the system with information on what products are available and handles a central list of entitlements and remaining quantities for the entire organization.

The content delivery network is responsible for delivering the content to the system when requested. The content server is configured in the Red Hat Subscription Manager configuration and then tied into the system's yum service through the Red Hat Subscription Manager yum plug-in.

Both the subscription service and the content server used by a system's Red Hat Subscription Manager tools can be customized. The default settings use the public subscription service and content delivery network, but either one can be changed to use organization-specific services.

Note

Systems have the option of using the older Red Hat Network and Satellite 5.x systems to deliver content. These content delivery mechanisms bypass the subscription service in Certificate-based Red Hat Network, so there is no entitlement management. This is allowed for legacy infrastructures, but Red Hat strongly recommends registering new systems with Certificate-based Red Hat Network.

14.1.5. Advanced Content Management: Extended Update Support

Sometimes software product installations are straightforward — you want to install a Red Hat Enterprise Linux server, so you install Red Hat Enterprise Linux. However, products can have dependencies with each other (product B is only worthwhile if product A is also installed) or products can interact with each other to provide extended functionality. There are two categories of these kinds of product interactions:

- Dependencies, where one product requires or relies on another product directly
- Modifiers, where a product provides enhanced functionality or services for existing products

Dependencies are common and can be handled directly when processing content through tools like yum.

Modifiers can be more subtle. A modifier subscription extends another entitlement and provides different repository access and support than the product entitlement alone.

If the system is subscribed to that product entitlement or combination of products, then the modifier subscription brings an enhanced content set for that product. The content set can include additional new products, new functionality, or extended service and support, depending on the product being modified.

One simple example of a modifier is extended update support (EUS), which extends support for a minor release of Red Hat Enterprise Linux from six months to 24 months. An EUS subscription provides an enhanced support path, rather than a new product. EUS works only in conjunction with another product, to extend its support profile; it does not stand alone.

Red Hat Enterprise Linux Add-ons and EUS Subscriptions

Red Hat Enterprise Linux add-ons have access to EUS streams as long as the underlying Red Hat Enterprise Linux product has an EUS subscription. For example, if an administrator has a Red Hat Enterprise Linux 2 Socket subscription, a File System subscription, and a Red Hat Enterprise Linux 2 Socket EUS subscription, then the system can access both non-EUS and EUS content for both the Red Hat Enterprise Linux server and the File System product.

14.1.6. Certificate-based Red Hat Network versus RHN Classic

During the firstboot process, there are two options given for the content server: (Certificate-based) Red Hat Network and RHN Classic. These systems are mutually exclusive, but they both handle software content and updates as well as subscriptions and system inventory.

> **Important**
>
> This entire chapter deals with entitlement and subscription management through Certificate-based Red Hat Network with the subscription service tools. This is the recommended content/subscription system for Red Hat Enterprise Linux 5.7 and later systems.

In 5.7 and later versions, entitlements and subscriptions are defined by available and installed products. However, in older versions of Red Hat Enterprise Linux, subscriptions were defined by channel access. These are two different approaches to content and entitlement access. Red Hat Network uses the product-based subscription model, while RHN Classic uses the channel-based model.

Certificate-based Red Hat Network is focused on two things:

- Subscription management
- Content delivery

Certificate-based Red Hat Network integrates the Customer Portal, content delivery network, and subscription service (subscription management). It uses simple local tools (the Red Hat Subscription Manager client) to view and assign subscriptions for the installed products and to manage subscriptions as they expire.

Since the client tools for subscription management (the focus of Certificate-based Red Hat Network) are only available in Red Hat Enterprise Linux 5.7 systems and later, Certificate-based Red Hat Network can only be utilized by 5.7 and later systems.

RHN Classic uses the traditional channel entitlement model, which provides a global view of content access but does not provide insight into system-level subscription uses. Along with content and global subscription management, RHN Classic also provides some systems management functions:

- Kickstarting systems
- Managing configuration files
- Running scripts
- Taking system snapshots

Satellite 5.x systems use a channel-based model similar to RHN Classic.

While RHN Classic has an expanded systems management feature set, RHN Classic does not provide the system-level view into installed and subscribed products that the enhanced Red Hat Network and subscription service do. RHN Classic is provided for older Red Hat Enterprise Linux systems (Red Hat Enterprise Linux 4.x, Red Hat Enterprise Linux 5.x, and Satellite 5.x) to migrate systems over to Red Hat Enterprise Linux 5.7 and later versions.

When a system is registered with RHN Classic, then the Red Hat Subscription Manager shows an error that the system is already registered and cannot be managed by the Subscription Manager tools.

Likewise, similar errors are returned in the RHN Classic tools if a system is registered with Red Hat Network and the subscription service.

The two subscription services are mutually exclusive, with separate inventories and using separate client tools. Both the RHN Classic and Red Hat Subscription Manager tools correctly identify which service a system is registered with. When a system is registered with RHN Classic, then the Red Hat Subscription Manager shows an error that the system is already registered and cannot be managed by the Subscription Manager tools. Likewise, similar errors are returned in the RHN Classic tools if a system is registered with Red Hat Network and the subscription service.

For information on migrating from RHN Classic to Certificate-based Red Hat Network, see 14.5절. “Migrating Systems from RHN Classic to Certificate-based Red Hat Network” .

14.2. Using Red Hat Subscription Manager Tools

The Red Hat Subscription Manager tool set encompasses three different tools:

- A GUI-based local client to manage the local machine
- A CLI client for advanced users and administrators to manage a local machine (and which can be tied into other applications and actions, like kickstarting machines)
- A web-based client for organizational, multi-system views of the subscriptions and inventoried resources

All of these tools, both local clients and the web-based tools, allow administrators to perform three major tasks directly related to managing subscriptions: registering machines, assigning subscriptions to systems, and updating the certificates required for authentication. Some minor operations, like updating system facts, are available to help display and track what subscriptions are available.

> **Note**
>
> Both the Red Hat Subscription Manager GUI and CLI must be run as root.

14.2.1. Launching Red Hat Subscription Manager

Red Hat Subscription Manager is listed as one of the administrative tools in the Applications > System Tools menu in the top management bar.

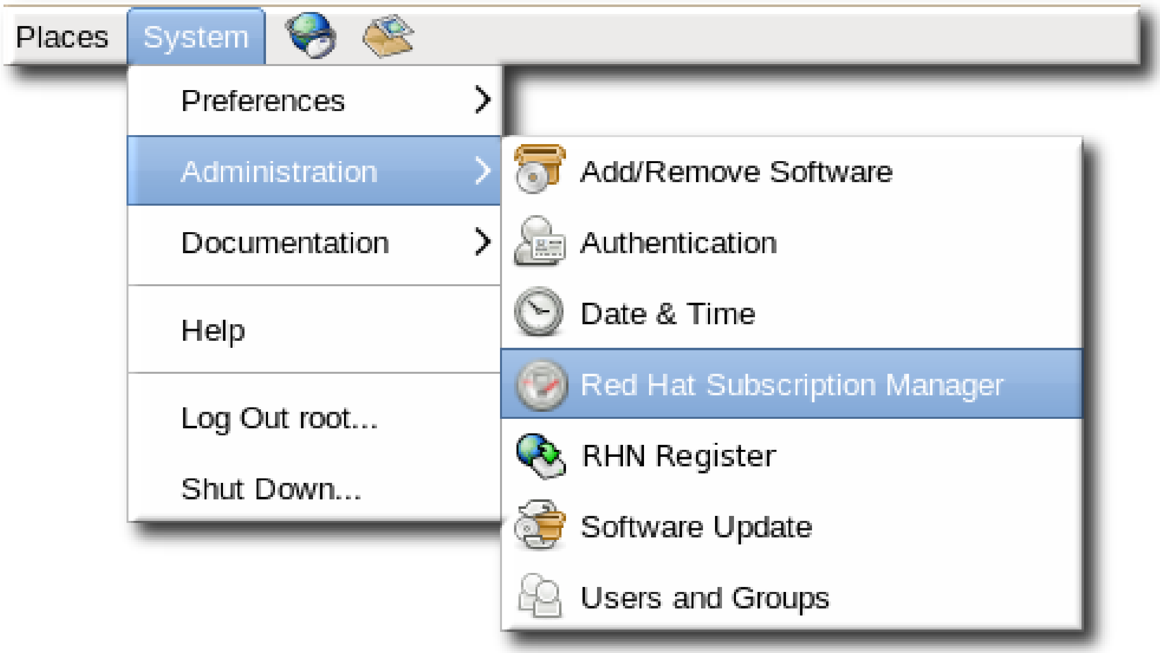

그림 14.3. Red Hat Subscription Manager Menu Option

Alternatively, the Red Hat Subscription Manager GUI can be opened from the command line with a single command:

```
[root@server1 ~]# subscription-manager-gui
```

The Red Hat Subscription Manager UI has a single window with tabbed sections that offer quick views into the current state of the system, showing installed products, subscriptions for the system, and available subscriptions the system has access to. These tabs also allow administrators to manage subscriptions by subscribing and unsubscribing the system.

The Red Hat Subscription Manager has three tabs which manage products and subscriptions:

- The My Subscriptions tab shows all of the current entitlements that the system is subscribed to.
- The All Available Subscriptions tab shows all of the subscriptions that are available to the system. The default displays only entitlements that are compatible with the hardware, but these can be filtered to show entitlements corresponding to other installed programs, only subscriptions that have not been installed, and subscriptions based on date.
- The My Installed Software tab shows the currently installed products on the system, along with their subscription status. This does not allow administrators to install software, only to view installed software.

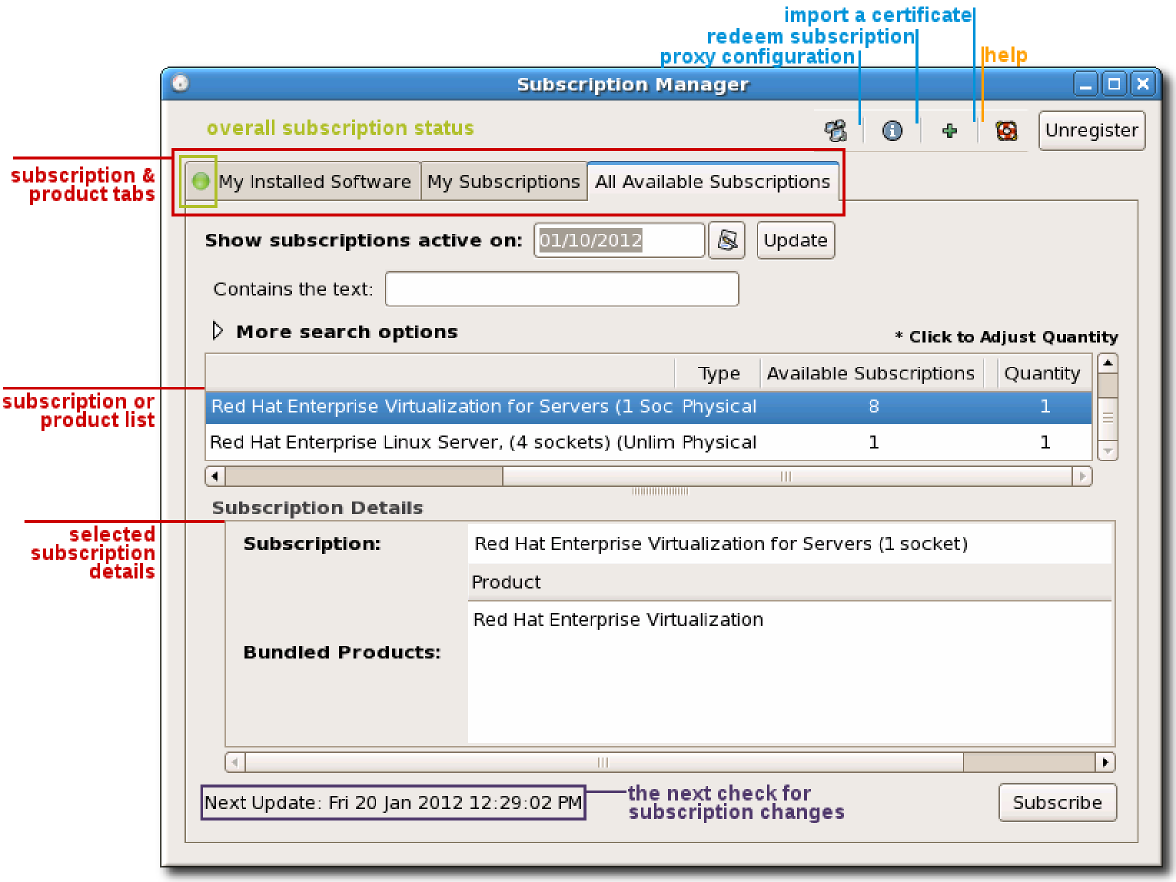
import a certificate
redeem subscription
proxy configuration
help
Subscription Manager
overall subscription status
Unregister
subscription & product tabs
My Installed Software
My Subscriptions
All Available Subscriptions
Show subscriptions active on:
01/10/2012
Update
Contains the text:
More search options
* Click to Adjust Quantity
subscription or product list
selected subscription details
Subscription Details
the next check for subscription changes
Subscribe

	Type	Available Subscriptions	Quantity
Red Hat Enterprise Virtualization for Servers (1 Soc	Physical	8	1
Red Hat Enterprise Linux Server, (4 sockets) (Unlim	Physical	1	1

Subscription: Red Hat Enterprise Virtualization for Servers (1 socket)

Bundled Products: Product
Red Hat Enterprise Virtualization

Next Update: Fri 20 Jan 2012 12:29:02 PM

그림 14.4. Red Hat Subscription Manager Main Screen

The series of icons in the top right corner of the box are used to perform system-related maintenance tasks like changing the proxy connection information and viewing system facts.

14.2.2. About subscription-manager

Any of the operations that can be performed through the Red Hat Subscription Manager UI can also be performed by running the subscription-manager tool. This tool has the following format:

```
[root@server1 ~]# subscription-manager command [options]
```

Each command has its own set of options that are used with it. The subscription-manager help and manpage have more information.

표 14.1. subscription-manager Commands

Command	Description
register	Registers or identifies a new system to the subscription service.
unregister	Unregisters a machine, which strips its subscriptions and removes the machine from the subscription service.
subscribe	Allocates a specific subscription to the machine.
redeem	Autosubscribes a machine to a pre-specified subscription that was purchased from a vendor, based on its hardware and BIOS information.
refresh	Pulls the latest entitlement data from the server. Normally, the system polls the entitlement server at a set interval (4 hours by

Command	Description
	default) to check for any changes in the available subscriptions. The refresh command checks with the entitlement server immediately, outside the normal interval.
unsubscribe	Removes a specific subscription or all subscriptions from the machine.
list	Lists all of the subscriptions that are compatible with a machine, either subscriptions that are actually consumed by the machine or unused subscriptions that are available to the machine.
identity	Handles the identity certificate and registration ID for a system. This command can be used to return the current UUID or generate a new identity certificate.
facts	Lists the system information, like the release version, number of CPUs, and other architecture information.
clean	Removes all of the subscription and identity data from the local system, without affecting the consumer information in the subscription service. Any of the subscriptions consumed by the system are still consumed and are not available for other systems to use. The clean command is useful in cases where the local entitlement information is corrupted or lost somehow, and the system will be reregistered using the register --consumerid=EXISTING_ID command.
orgs, repos, environments	Lists all of the configured organizations, environments, and content repositories that are available to the given user account or system. These commands are used to view information in a multi-org infrastructure. They are not used to configure the local machine or multi-org infrastructure.

14.2.3. Looking at RHN Subscription Management

The ultimate goal of entitlement management is to allow administrators to identify the relationship between their systems and the subscriptions used by those systems. This can be done from two different perspectives: from the perspective of the local system looking externally to potential subscriptions and from the perspective of the organization, looking down at the total infrastructure of systems and all subscriptions.

The Red Hat Subscription Manager GUI and CLI are both local clients which manage only the local machine. These tools are somewhat limited in their view; they only disclose information (such as available entitlements) from the perspective of that one system, so expired and depleted subscriptions or subscriptions for other architectures are not displayed.

RHN Subscription Management in the Customer Portal is a global tool which is intended to give complete, organization-wide views into subscriptions and systems. It shows all subscriptions and all consumers for the entire organization. RHN Subscription Management can perform many of the tasks of the local tools, like registering consumers, assigning subscriptions, and viewing system facts and UUID. It can also manage the subscriptions themselves, such as viewing contract information and renewing subscriptions — a task not possible in the local clients.

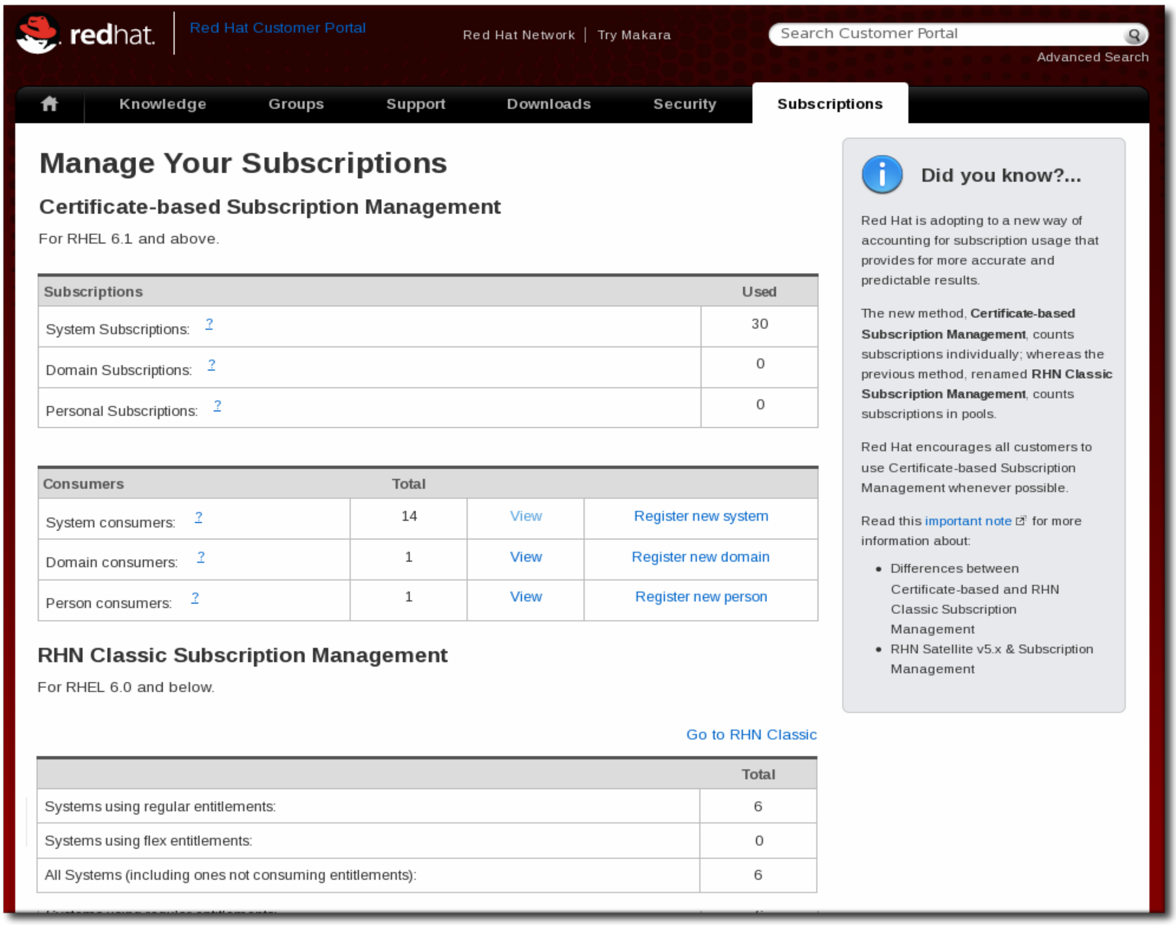

Manage Your Subscriptions

Certificate-based Subscription Management

For RHEL 6.1 and above.

Subscriptions	Used
System Subscriptions: ?	30
Domain Subscriptions: ?	0
Personal Subscriptions: ?	0

Consumers	Total		
System consumers: ?	14	View	Register new system
Domain consumers: ?	1	View	Register new domain
Person consumers: ?	1	View	Register new person

RHN Classic Subscription Management

For RHEL 6.0 and below.

Go to RHN Classic

	Total
Systems using regular entitlements:	6
Systems using flex entitlements:	0
All Systems (including ones not consuming entitlements):	6

Did you know?...

Red Hat is adopting to a new way of accounting for subscription usage that provides for more accurate and predictable results.

The new method, **Certificate-based Subscription Management**, counts subscriptions individually; whereas the previous method, renamed **RHN Classic Subscription Management**, counts subscriptions in pools.

Red Hat encourages all customers to use Certificate-based Subscription Management whenever possible.

Read this important note for more information about:

- Differences between Certificate-based and RHN Classic Subscription Management
- RHN Satellite v5.x & Subscription Management

그림 14.5. RHN Subscription Management in the Customer Portal

Note

RHN Subscription Management gives a global view of all consumers, of all types, for an organization, which is crucial for planning and effectively assigning subscriptions. However, it does not provide any insight into what products are installed on a system and whether subscriptions are assigned for those products. To track the validity of installed products, you must use the local Subscription Manager tools.

RHN Subscription Management also provides a view of systems and subscriptions managed under RHN Classic and provides access to the RHN Classic web tools.

All of the subscriptions for an entire organization — the subscriptions that have been purchased and the systems to which they have been allocated — are viewable through the account pages at https://access.redhat.com/. Additional information about RHN Subscription Management is available with the portal documentation at https://access.redhat.com/knowledge/docs/Red_Hat_Customer_Portal/.

14.2.4. Looking at Subscription Asset Manager

The simplest model for assigning subscriptions and delivering content is for local systems to connect directly to Red Hat's hosted subscription and content network. However, for large environments, highly-secure environments, and many other situations, that hosted arrangement is not feasible.

In that case, a block of subscriptions can be allocated to a distributor application. That distributor connects to Red Hat's infrastructure, and then it manages all of the systems and consumers at its local site. This has performance benefits by lowering bandwidth, and it offers significant management benefits to administrators by allowing local and flexible control over subscription management.

The Subscription Asset Manager application is a distributor. It is available as an additional layered application as part of a Red Hat Enterprise Linux subscription.

Subscription Asset Manager provides a local site not only to view subscriptions and systems for an infrastructure (as with the Customer Portal) but also to manage all of those systems. Subscription Asset Manager has three major functional areas:

- Works with the client machine's Subscription Manager to manage subscriptions and content. In that way, it is a centralized, global, web-based Subscription Manager.
- Helps manage the subscriptions themselves. It receives a subscription manifest from Red Hat Network. The manifest allocates that Subscription Asset Manager service a subset of all of an organization's subscriptions. From there, the Subscription Asset Manager locally assigns subscriptions to individual systems and can create activation keys.
- Works as a real-time proxy between the local system assets and the Red Hat content delivery network.

Subscription Asset Manager handles both client-side, local system management and backend subscription management. This allows Subscription Asset Manager to provide more in-depth information on the status of products and certificates through tools like its dashboard and activity reports.

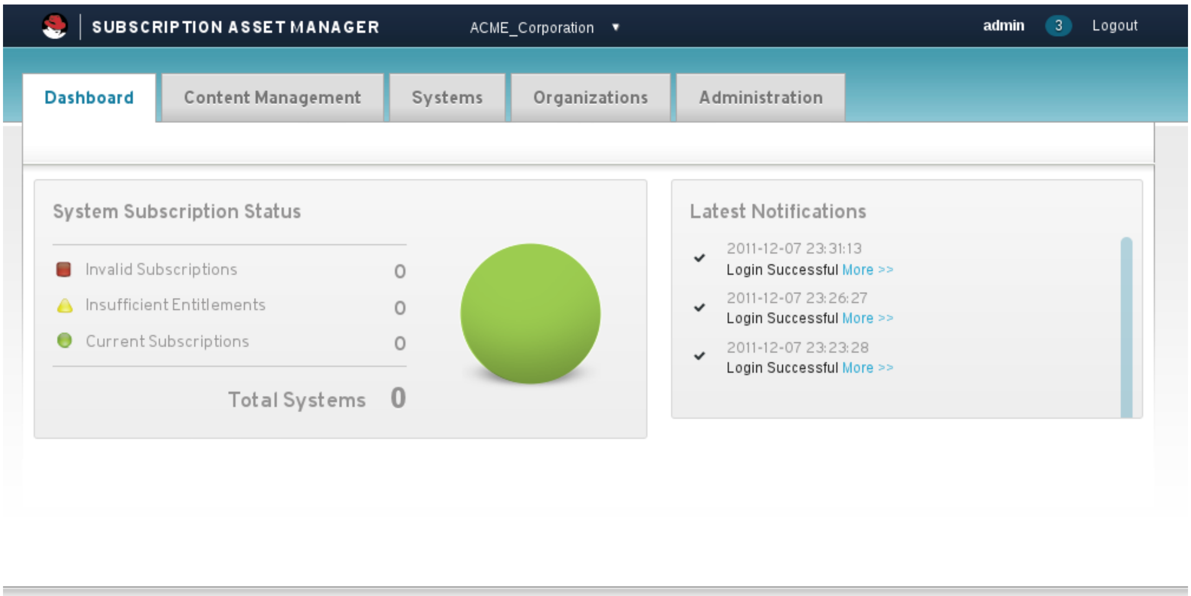

System Subscription Status

🟥 Invalid Subscriptions	0
⚠ Insufficient Entitlements	0
● Current Subscriptions	0
Total Systems	0

Latest Notifications

✓ 2011-12-07 23:31:13
Login Successful More >>

✓ 2011-12-07 23:26:27
Login Successful More >>

✓ 2011-12-07 23:23:28
Login Successful More >>

그림 14.6. Subscription Asset Manager Dashboard

Because of the insight Subscription Asset Manager has into the local server assets, it can be used to define multi-tenant organizations. Multi-tentant organizations allow completely separate silos of assets (organizations). Organizations can then be subdivided into environments; since a system can belong to multiple environments, it is possible to organize systems into overlapping circles according to the real-world infrastructure. This is covered more in 14.3.1절. "Local Subscription Services, Local Content Providers, and Multi-Tenant Organizations".

Subscription Asset Manager is available with Red Hat Enterprise Linux, but it must be installed and configured before it can be used to manage assets.

For more information on configuring and using Subscription Asset Manager, see the documentation at http://docs.redhat.com/docs/en-US/Red_Hat_Subscription_Asset_Manager/1.0/html/Installation_Guide/index.html.

14.3. Managing Special Deployment Scenarios

There are different types of consumers and different ways of organizing consumers. Subscription Manager, and the underlying concepts of a subscription server and content provider, are flexible enough to accommodate special types of consumers and different infrastructure setups.

In particular, there are three deployment scenarios that are common in IT environments:

- Multi-tenant organizations.

  The most basic subscription/content service scenario has a consumer connecting directly to Red Hat's hosted services and receiving content and subscription updates directly from Red Hat.

  That flat model is simple, but it does not accurately describe many enterprise environments, which are divided across disparate organizational units and even subunits. It also does not account for network performance or security issues which may require a company to have subscription and content information maintained locally.

  For this scenario, Red Hat supports local subscription/content infrastructures through distributors like Subscription Asset Manager. These distributor applications can subdivide both subscription and content into organizations and subordinate environments.

- Virtual and physical machines.

  There are separate entitlements for virtual and physical machines. Subscription Manager can detect and manage subscriptions appropriately depending on the type of machine.

- Server domains.

  In some cases, a group of machines may act in concert to perform a certain function, like a mail domain or cluster. In those situations, subscriptions may apply to the domain as a group, as opposed to applying to any one machine.

All of these special deployment scenarios are described in more detail in the following sections.

14.3.1. Local Subscription Services, Local Content Providers, and Multi-Tenant Organizations

As 14.1.4절. "Subscription and Content Architecture" outlines, the subscription service, content repository, and client tools and inventory all work together to define the entitlements structure for a customer. The way that these elements are organized depends on a lot of factors, like who is maintaining the individual services, how systems in the inventory are group, and how user access to the different services is controlled.

The most simplistic structure is the hosted structure. The content and subscription services are hosted by Red Hat, and all systems within the inventory are contained in one monolithic group. User access is defined only by Red Hat Customer Portal account access.

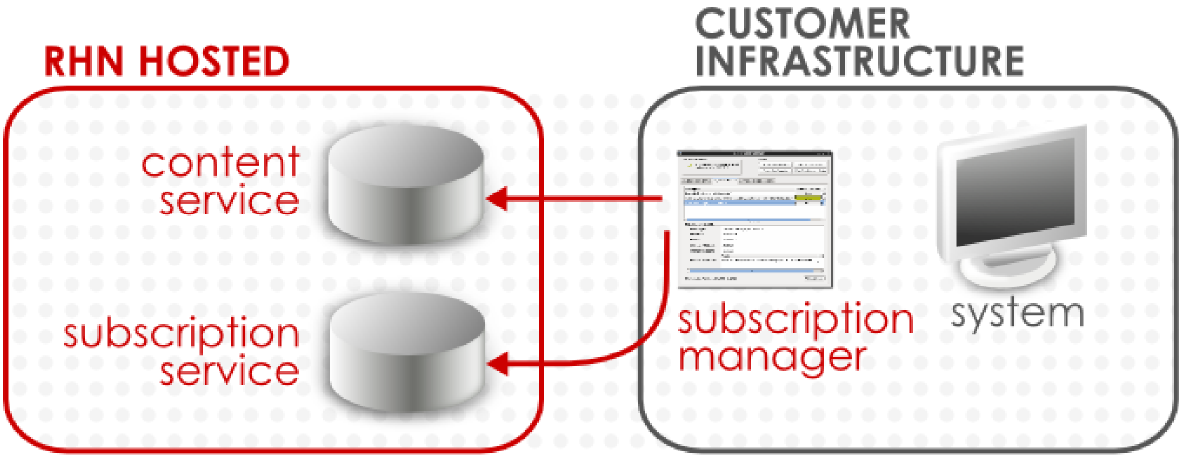

그림 14.7. Hosted Structure

The next configuration allows a customer to have its own, local subscription and content services. This allocates a block of subscriptions to that service, the distributor, and then the distributor interacts directly with local systems.

One distributor scenario is to have the distributor function as a subscription service, while still using Red Hat's hosted content delivery network. This is common when using Subscription Asset Manager, which can define subscription allocation to local systems based on organization and environment.

Additionally, user accessss can be defined locally, within the Subscription Asset Manager configuration. Subscription Asset Manager can define independent groups, called organizations. Systems belong to those organizations, and users are granted access to those organizations. Systems and users in one organization are essentially invisible to systems and users in other organizations.

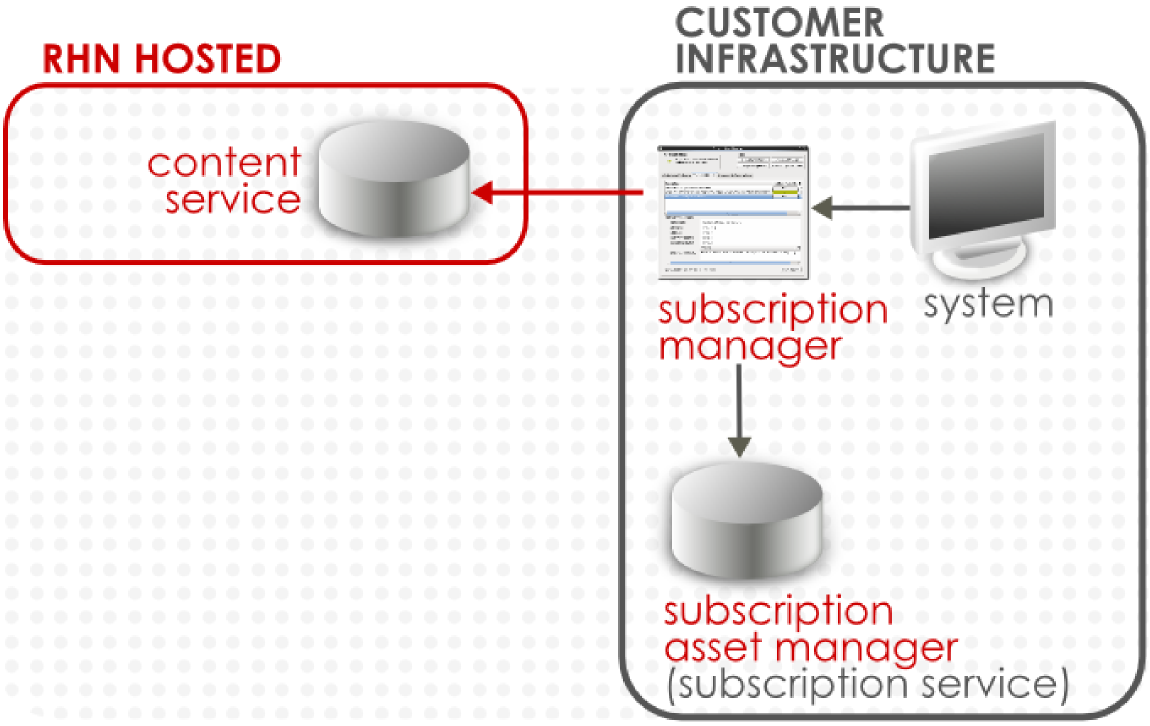

그림 14.8. Hosted Content/Local Subscriptions Structure

The last style of infrastructure is almost entirely local, with a distributor that provides locally-hosted content providers and an integrated local subscription service.

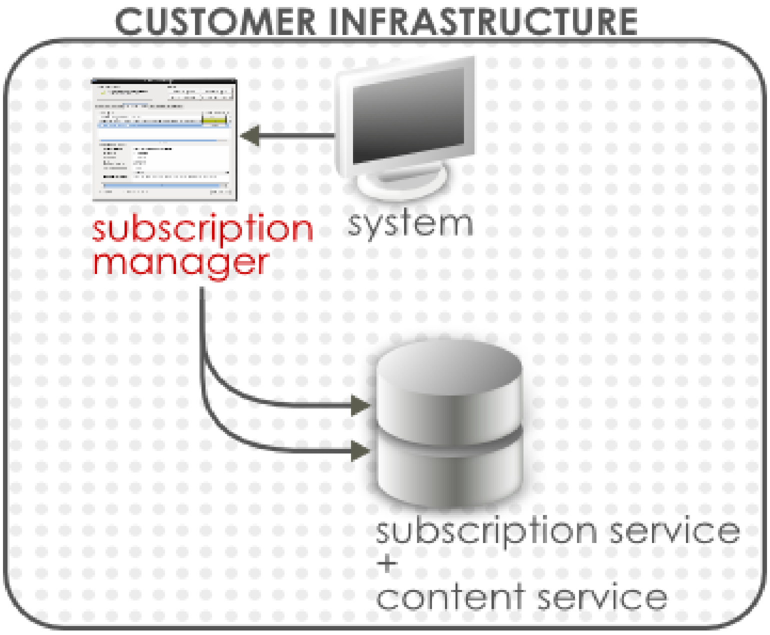

그림 14.9. Local Subscriptions and Local Content Provider Structure

This allows the most control over how systems are grouped within the subscriptions/content service. A customer's main account can be divided into separate and independent organizations. These organizations can use different content providers, can have different subscriptions allocated to them, and can have different users assigned to them with levels of access set per organization. Access control in this scenario is controlled entirely locally. The local distributor, not the remote Red Hat Customer Portal, processes user authentication requests and applies local access control policies.

A system is assigned to one organization. It is identified with that organization.

Within an organization, there can be different environments which define access to product versions and content sets. There can be overlap between environments, with a system belonging to multiple environments.

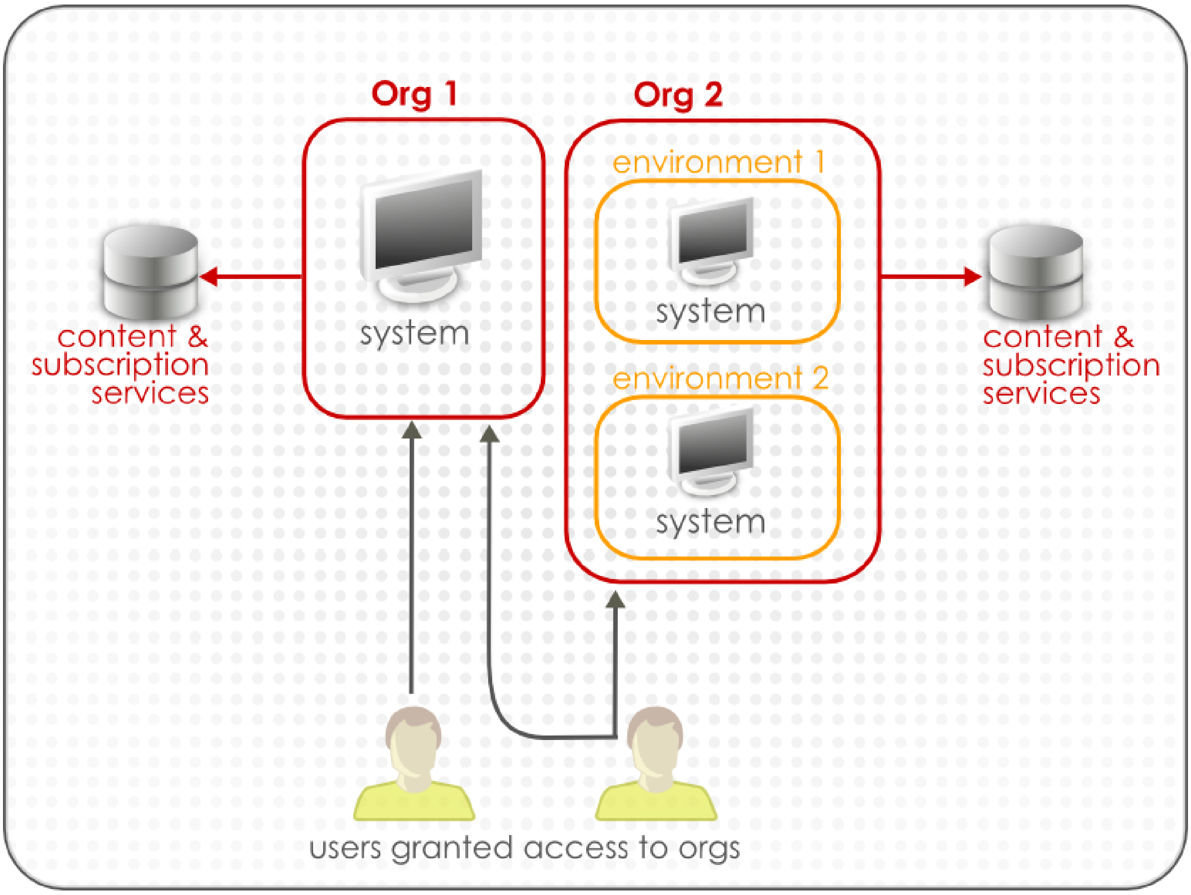
Org 1
Org 2
environment 1
environment 2
system
system
system
content & subscription services
content & subscription services
users granted access to orgs

그림 14.10. Multi-Org

When there is only one organization — such as a hosted environment (where the single organization is implicit) — then the systems all default to use that one organization. When there are multiple organizations, then the organization for a system to use must be defined for that system. This affects register operations, where the system is registered to the subscription service and then joined to the organization. It also affects other operations tangentially. It may affect subscribe operations because it affects repository availability and subscription allocations, and it affects redeem operations (activation of existing subscriptions) because subscriptions must be redeemed from the organization which issued the subscription.

For more information on configuring and managing organizations, environments, and content repositories, see the Subscription Asset Manager documentation[1].

14.3.2. Virtual Guests and Hosts

When the Red Hat Subscription Manager process checks the system facts, it attempts to identify whether the system is a physical machine or a virtual guest. The Subscription Manager can detect guests for several different virtualization services, including:

- KVM
- Xen
- HyperV

[1] http://docs.redhat.com/docs/en-US/Red_Hat_Subscription_Asset_Manager/1.0/html/Installation_Guide/index.html

- VMWare ESX

Subscription Manager records a unique identifier called a guest ID as one of the system facts for a virtual guest. A special process, virt-who, checks virtual processes and then relays that information to Subscription Manager and any configured subscription service (Certificate-based Red Hat Network or a local Subscription Asset Manager). Each guest machine on a host is assigned a guest ID, and that guest ID is both associated with the host and used to generate the identity certificate for the guest when it is registered.

Some Red Hat Enterprise Linux variants are specifically planned for virtual hosts and guests. The corresponding subscriptions are divided into a certain quantity of physical hosts and then a quantity of allowed guests. Red Hat Enterprise Linux add-ons may even be inherited, so that when a host machine is subscribed to that entitlement, all of its guests are automatically included in that subscription. (Red Hat layered products usually do not draw any distinction between virtual and physical systems; the same type of subscription is used for both.) If the system is a guest, then virtual entitlements are listed with the available subscriptions. If no more virtual entitlements are available, then the subscription service will apply physical entitlements.

Virtual and physical subscriptions are identified in the Type column.

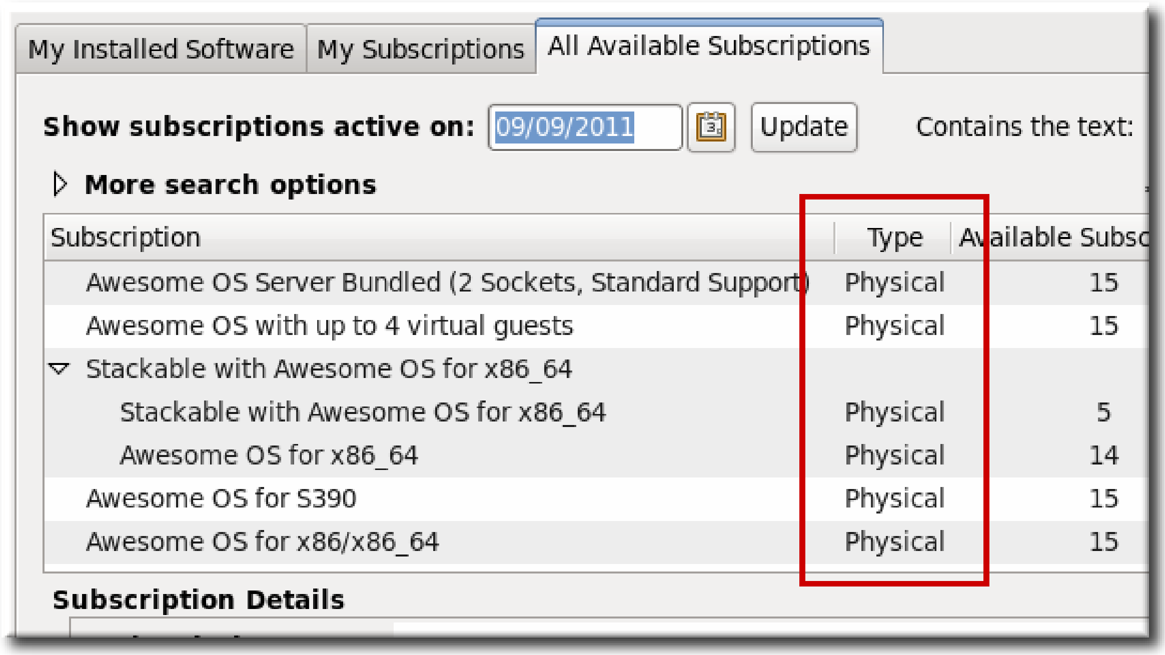

Show subscriptions active on: 09/09/2011

Subscription	Type	Available Subsc
Awesome OS Server Bundled (2 Sockets, Standard Support)	Physical	15
Awesome OS with up to 4 virtual guests	Physical	15
Stackable with Awesome OS for x86_64		
Stackable with Awesome OS for x86_64	Physical	5
Awesome OS for x86_64	Physical	14
Awesome OS for S390	Physical	15
Awesome OS for x86/x86_64	Physical	15

그림 14.11. Virtual and Physical Subscription

참고

The distinction of being a physical machine versus virtual machine matters only in the priority of how entitlements are consumed.

Virtual guests are registered to the subscription service inventory as regular systems and subscribe to entitlements just like any other consumer.

Virtual entitlements can only be used by virtual machines. Physical entitlements can be used by both physical and virtual machines. When ascertaining what subscriptions are available for autosubscription, preference is given first to virtual entitlements (which are more restrictive in the type of consumer which can use them), and then to physical entitlements.

14.3.3. Domains

Consumers in the subscription service inventory are identified by type. Most consumers will have a type of system, meaning that each individual server subscribes to its own entitlements for its own use. There is another type of consumer, though, which is available for server groups, the domain type. domain-based entitlements are not allocated to a single system; they are distributed across the group of servers to govern the behavior of that group of servers. (That server group is called a domain.)

There are two things to keep in mind about domain entitlements:

- Each member of the domain is still registered to the subscription service as a system consumer and added to the inventory individually.

- The domain entitlements apply to the behavior of the entire server group, not to any one system.

The domain entitlement simply governs the behavior of the domain. A domain entitlement is not limited to a specific type of behavior. Domain entitlements can describe a variety of types of behavior, such as storage quotas or the maximum number of messages to process per day. The entire domain is bound to the subscriptions when one of the domain servers subscribes to the domain entitlements using the Red Hat Subscription Manager tools, and the entitlement certificate is replicated between the domain servers.

14.4. Registering, Unregistering, and Reregistering a System

Entitlements are managed by organizing and maintaining the systems which use entitlement subscriptions. The entitlements and subscriptions are managed by Red Hat through the subscription service. A system is recognized to the subscription service by being registered with the service. The subscription service assigns the system (called a consumer) a unique ID (essentially as an inventory number) and issues that system an identifying certificate (with the UUID in the certificate subject name) to identify that system.

Whenever a subscription is purchased by an organization, the consumer can subscribe to that subscription. This means that a portion of the subscription is allocated to that consumer ID; when the consumer contacts the content delivery network and downloads the software, the licenses have been already assigned to the system. The system has valid certificates for its subscriptions.

Systems can be registered with a subscription service during the firstboot process or as part of the kickstart setup (both described in the Installation Guide). Systems can also be registered after they have been configured or removed from the subscription service inventory (unregistered) if they will no longer be managed within that entitlement system.

14.4.1. Registering Consumers in the Hosted Environment

For infrastructures which use Red Hat's hosted subscription and content delivery network, all that is required to register the system is the username and password of the Red Hat Network account.

1. Launch Subscription Manager. For example:

```
[root@server ~]# subscription-manager-gui
```

2. If the system is not already registered, then there will be a Register button at the top of the window in the Tools area.

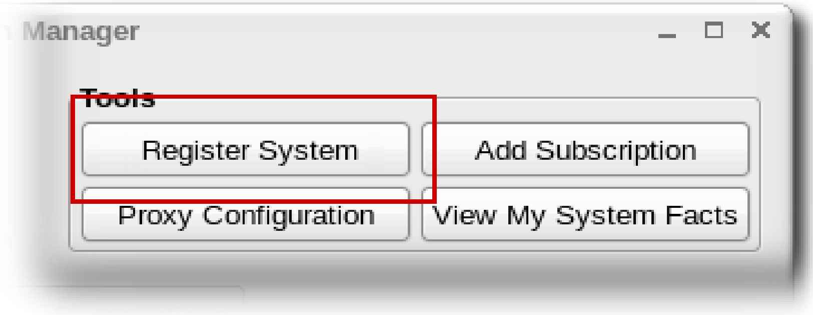

3. Enter the username and password of the user account on the subscription service; this is the account used to access the Customer Portal.

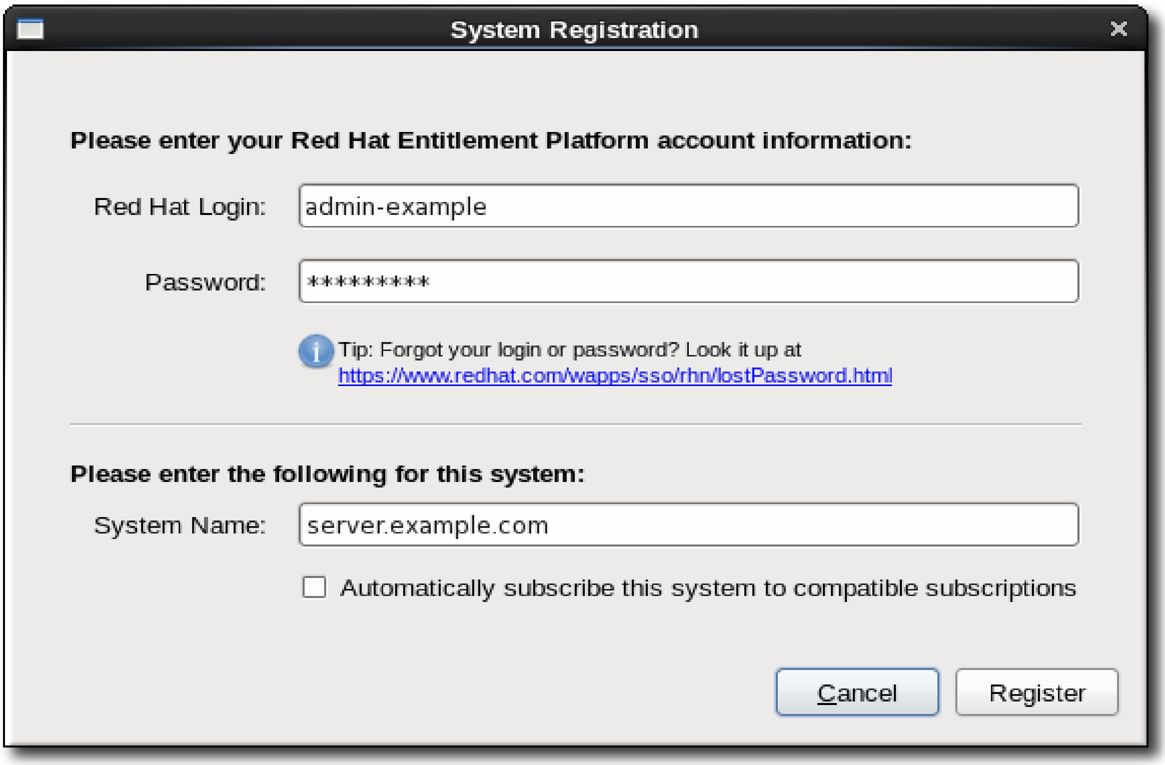

Please enter your Red Hat Entitlement Platform account information:

Red Hat Login: admin-example

Password: *********

Tip: Forgot your login or password? Look it up at https://www.redhat.com/wapps/sso/rhn/lostPassword.html

Please enter the following for this system:

System Name: server.example.com

☐ Automatically subscribe this system to compatible subscriptions

4. Optionally, select the Automatically subscribe... checkbox, so that the system is subscribed to the best matched subscription when it is registered. Otherwise, the system must be subscribed manually, as in 14.6절. "Handling Subscriptions" .

14.4.2. Registering Consumers to a Local Distributor (Organization)

Infrastructures which manage their own local content repository and subscription service are distributors.

A distributor application has a defined organization. This organization is essentially a group definition, and systems must be assigned to that group as part of the registration process. This allows there to be multiple, discrete organizations or tenants within the infrastructure.

When a system is registered using the Subscription Manager GUI, Subscription Manager automatically scans the local subscription and content service to see what organizations are configured.

참고

A system can only be registered to one organization, and that registration cannot be altered.

A system can be registered to multiple environments, if environments are configured. As with the organization assignment, environments are selected during registration and cannot be altered after registering.

1. Update the local Subscription Manager configuration to point to the distributor server rather than Red Hat's hosted services.

   This is described in more detail in 14.12.1절. "Configuring Subscription Manager to Work with Subscription Asset Manager" . In general, there are three settings that must be changed:

   - The subscription server's hostname
   - The content provider's hostname and port (8088)
   - The CA certificate for the subscription service

2. Launch Subscription Manager. For example:

```
[root@server ~]# subscription-manager-gui
```

3. Click the Register button at the top of the window in the Tools area.

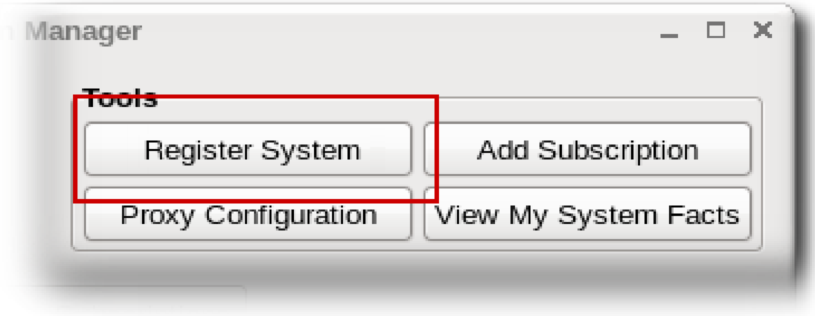

4. Enter the username and password of the user account on the subscription service; this is the account used to access the Customer Portal.

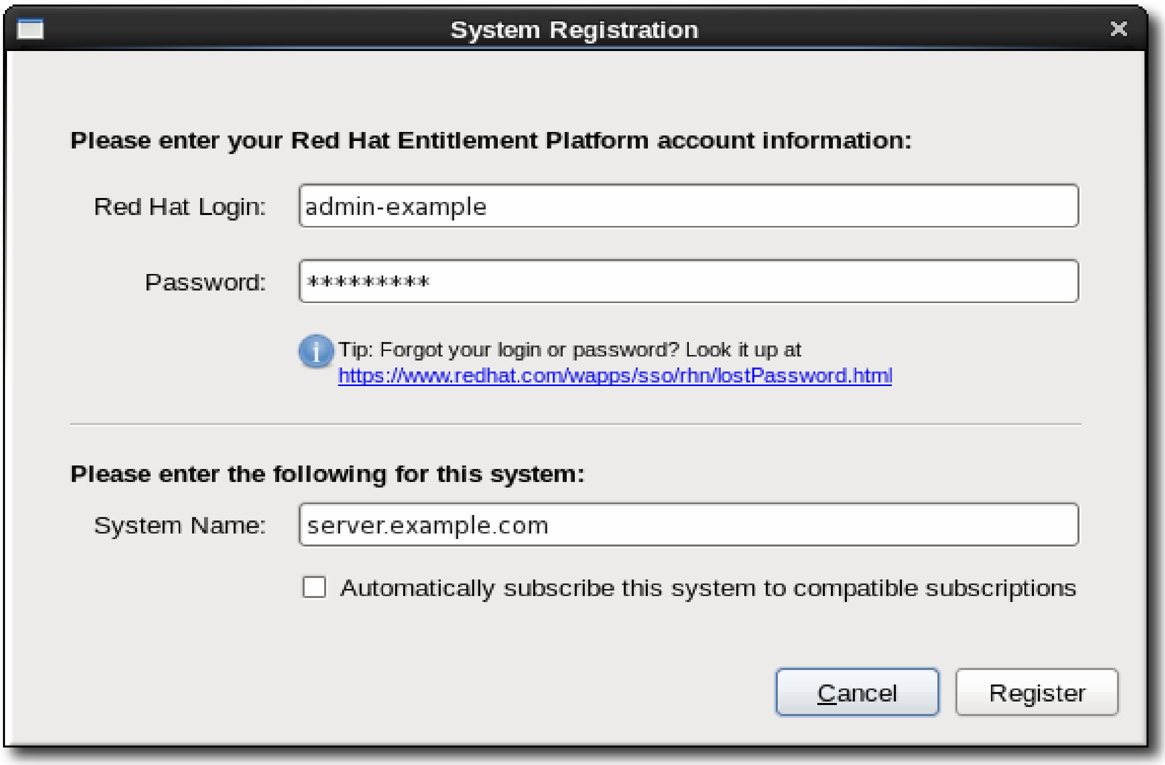

System Registration

Please enter your Red Hat Entitlement Platform account information:

Red Hat Login: admin-example

Password: *********

Tip: Forgot your login or password? Look it up at https://www.redhat.com/wapps/sso/rhn/lostPassword.html

Please enter the following for this system:

System Name: server.example.com

□ Automatically subscribe this system to compatible subscriptions

5. Subscription Manager scans the network for available organizations.

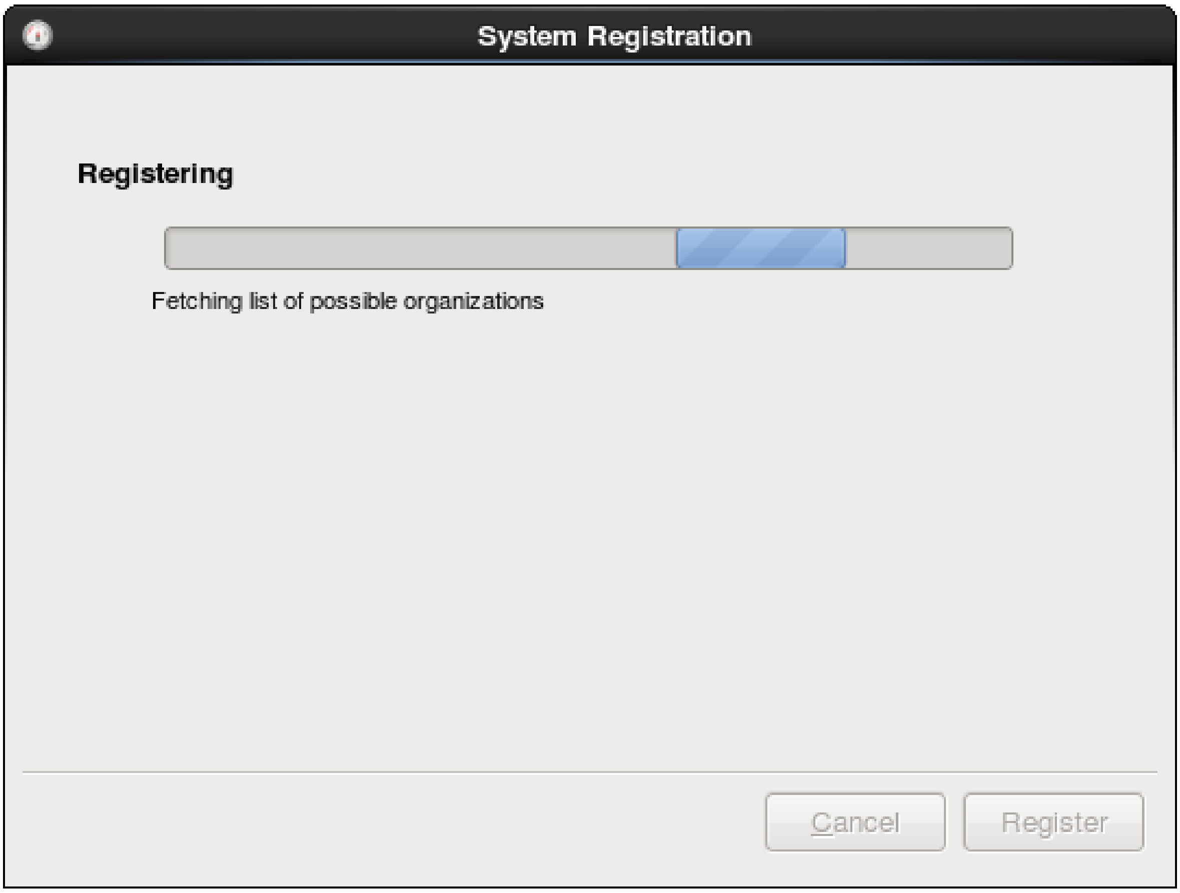

System Registration

Registering

Fetching list of possible organizations

When the configured organizations are detected, Subscription Manager prompts for the organization for the system to join. It is only possible to register with one organization.

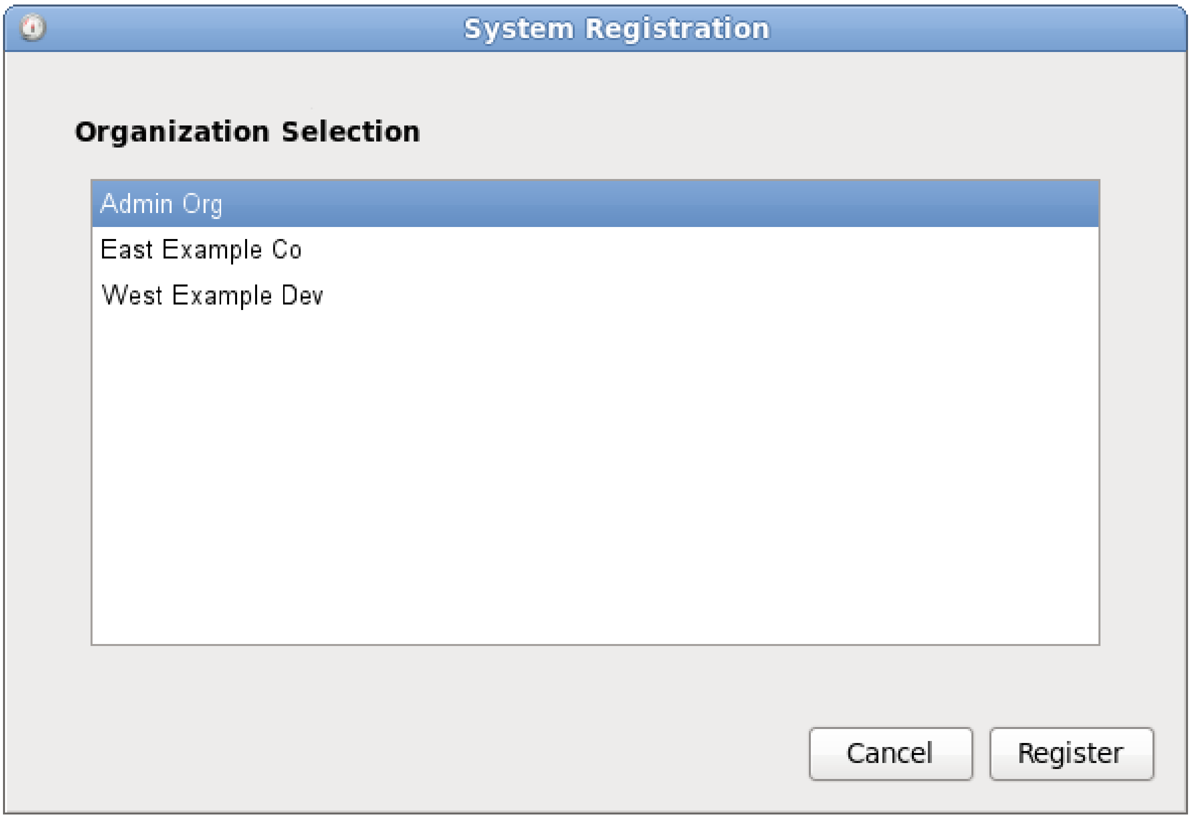

Organization Selection

Admin Org
East Example Co
West Example Dev

6. If the selected organization has multiple environments available, then the Subscription Manager will detect them and provide a list. It is possible to join multiple environments. Use the Ctrl key to select multiple environments from the list.

   If no environment is selected, then Subscription Manager uses the default environment for the organization.

   > **참고**
   >
   > It is only possible to join an environment during registration. The environments cannot be changed after registration.

7. Optionally, select the Automatically subscribe... checkbox, so that the system is subscribed to the best matched subscription when it is registered. Otherwise, the system must be subscribed manually, as in 14.6절. "Handling Subscriptions" .

14.4.3. Registering an Offline Consumer

Some systems may not have Internet connectivity, but administrators still want to assign and track the subscriptions for that system. This can be done by manually registering the system, rather than depending on Subscription Manager to perform the registration. This has two major steps, first to create an entry on the subscriptions service and then to configure the system.

1. Open the Subscriptions tab in the Customer Portal, and select the Overview item under the Certificate-based Management area.

2. In the Utilization area, click the Register a consumer link to create the new inventory entry.

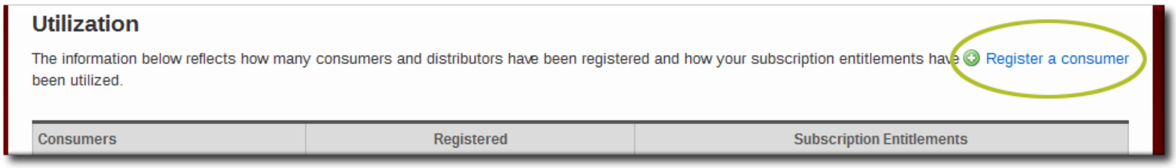

Utilization

The information below reflects how many consumers and distributors have been registered and how your subscription entitlements have been utilized.

Consumers	Registered	Subscription Entitlements

3. Fill in the required information for the new consumer type. A system requires information about the architecture and hardware in order to ascertain what subscriptions are available to that system.

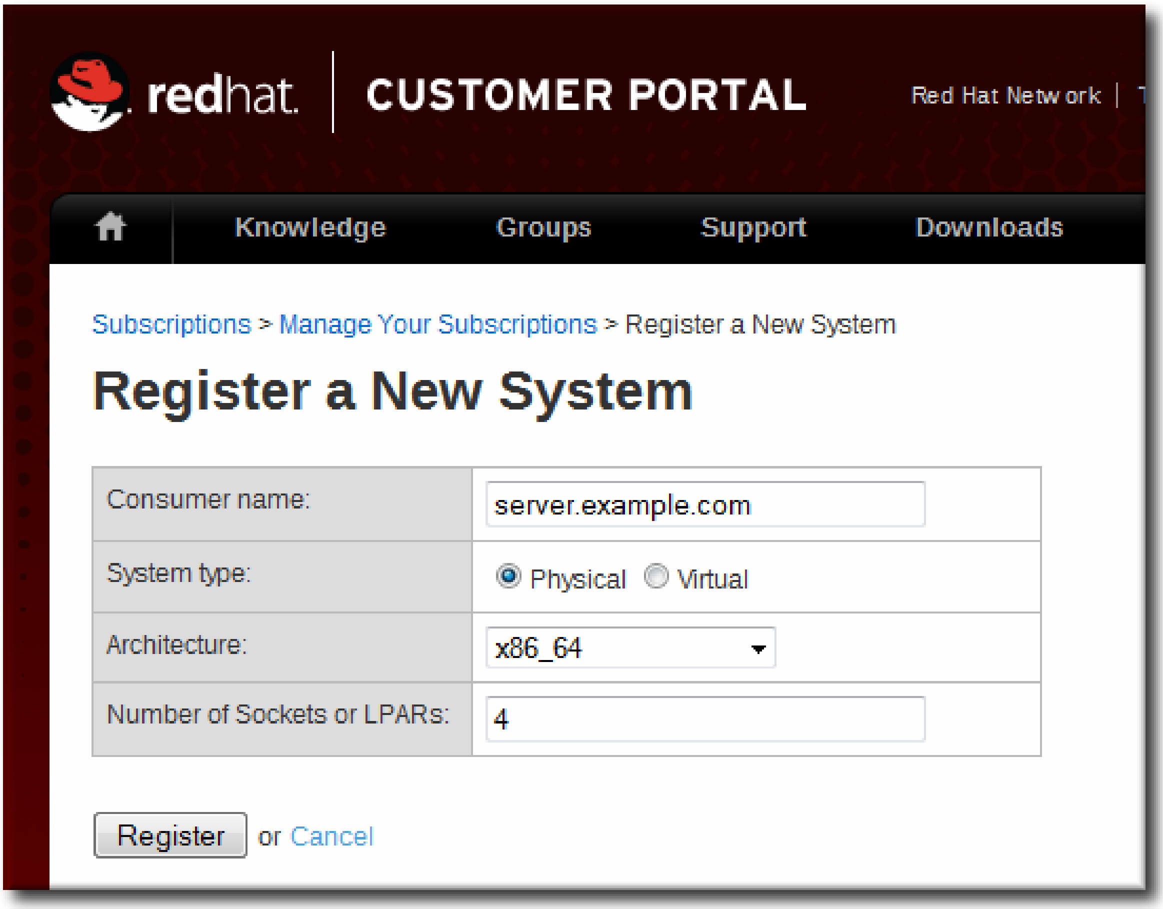

Subscriptions > Manage Your Subscriptions > Register a New System

Register a New System

Consumer name:	server.example.com
System type:	● Physical ○ Virtual
Architecture:	x86_64
Number of Sockets or LPARs:	4

4. Once the system is created, assign the appropriate subscriptions to that system.

   a. Open the Available Subscriptions tab.

   b. Click the checkboxes by all of the subscriptions to assign, and then click the Add button.

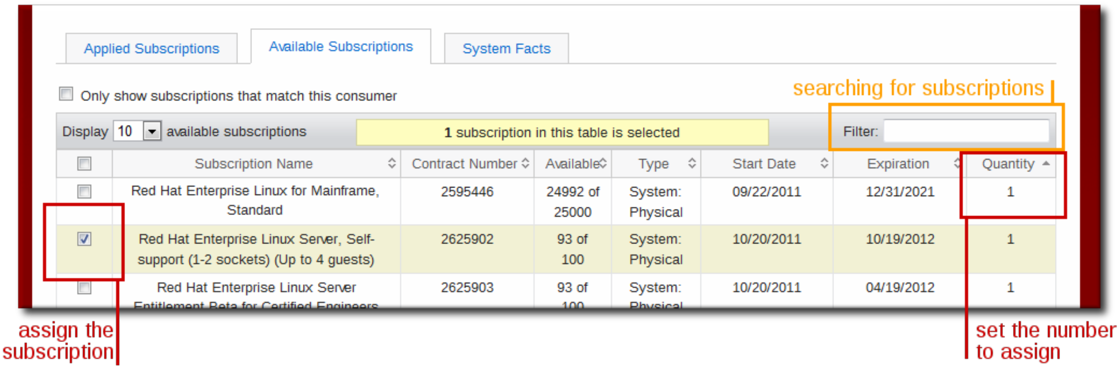

□ Only show subscriptions that match this consumer

Display 10 available subscriptions | **1** subscription in this table is selected | Filter:

□	Subscription Name	Contract Number	Available	Type	Start Date	Expiration	Quantity
□	Red Hat Enterprise Linux for Mainframe, Standard	2595446	24992 of 25000	System: Physical	09/22/2011	12/31/2021	1
☑	Red Hat Enterprise Linux Server, Self-support (1-2 sockets) (Up to 4 guests)	2625902	93 of 100	System: Physical	10/20/2011	10/19/2012	1
□	Red Hat Enterprise Linux Server Entitlement Beta for Certified Engineers	2625903	93 of 100	System: Physical	10/20/2011	04/19/2012	1

5. Once the subscriptions are added, open the Applied Subscriptions tab.

6. Click the Download All Certificates button. This exports all of the entitlements certificates, for each product, to a single .zip file. Save the file to some kind of portable media, like a flash drive.

7. Optionally, click the Download Identity Certificate button. This saves the identity certificate for the registered consumer and could be used by the consumer to connect to the subscription service. If the consumer will permanently be offline, then this is not necessary, but if the consumer could ever be brought onto the network, then this is useful.

8. Copy the entitlements certificates from the media device over to the consumer.

9. If all entitlement certificates were downloaded in an archive file, then there are multiple archives in the downloaded certificates.zip file. Unzip the directories until the PEM files for the entitlement certificates are available.

10. Import the entitlement certificates. This can be done using the Import Certificates button in the Subscription Manager GUI or using the import command. For example:

```
# subscription-manager import --certificate=/tmp/export/entitlement_certificates/596576341785244687.pem --certificate=/
tmp/export/entitlement_certificates/3195996649750311162.pem
Successfully imported certificate 596576341785244687.pem
Successfully imported certificate 3195996649750311162.pem
```

11. If you downloaded an identity certificate, copy the cert.pem file directly into the /etc/pki/consumer directory. For example:

```
cp /tmp/downloads/cert.pem /etc/pki/consumer
```

14.4.4. Registering from the Command Line

The simplest way to register a machine is to pass the register command with the user account information required to authenticate to the Certificate-based Red Hat Network (the credentials used to access subscription service or the Customer Portal). When the system is successfully authenticated, it echoes back the newly-assigned consumer ID and the user account name which registered it.

The register options are listed in 표 14.2. “register Options” .

예 14.1. Registering a New Consumer

```
[root@server1 ~]# subscription-manager register --username admin-example --password secret

The system has been registered with id: 7d133d55-876f-4f47-83eb-0ee931cb0a97
```

In a multi-org environment, it is required that you specify which organization (essentially an independent group or unit within the main account) to join the system to. This is done by using the --org option in addition to the username and password. The given user must also have the access permissions to add systems to that organization. (See 14.12절. "Working with Subscription Asset Manager" for information about organizations and Subscription Asset Manager.)

예 14.2. Registering a New Consumer with an Organization

If there is more than one organization, then the system must be assigned to one specific organization:

```
[root@server1 ~]# subscription-manager register --username admin-example --password secret --org="IT Department"

The system has been registered with id: 7d133d55-876f-4f47-83eb-0ee931cb0a97
```

Organizations can be subdivided into environments, which define access to content based on repositories, product versions, and content sets. While a consumer can only belong to a single organization, it can be assigned to multiple environments within that organization. If no environment is given, the subscription service uses the default environment. See 14.12절. "Working with Subscription Asset Manager" for information about organizations and Subscription Asset Manager.

A system can only be added to an environment during registration.

```
[root@server1 ~]# subscription-manager register --username admin-example --password secret --org="IT Department" --
environment=Dev1,ITall
```

참고

If the system is in a multi-org environment and no organization is given, the register command returns a Remote Server error.

The register command has an option, --autosubscribe, which allows the system to be registered to the subscription service and immediately subscribed to the subscription which best matches its architecture in a single step.

예 14.3. Automatically Subscribing While Registering

```
[root@server1 ~]# subscription-manager register --username admin-example --password secret --autosubscribe
```

예 14.4. Applying Subscriptions During Registration

When using the command-line tools to register the system, the --activationkey option can pass the activation key to apply existing, already-assigned certificates along with the other registration information. The activation keys for multiple subscriptions are set in a comma-separated list.

With an activation key, it is not necessary to give a username and password because the authentication is implicit in the activation key.

In hosted or single organization environments, it is not necessary to specify an organization with the --org option, but in multi-org environments, the --org option is required. The organization is not defined as part of the activation key. See 14.12절. "Working with Subscription Asset Manager" for information about activation keys and Subscription Asset Manager.

For example:

```
# subscription-manager register --activationkey=1234abcd --org="IT Dept"
```

표 14.2. register Options

Options	Description	Required
--username=name	Gives the content server user account name.	Required
--password=password	Gives the password for the user account.	Required
--org=name	Gives the organization to which to join the system.	Required, except for hosted environments
--environment=name	Registers the consumer to an environment within an organization.	Optional
--name=machine_name	Sets the name of the consumer (machine) to register. This defaults to be the same as the hostname.	Optional
--autosubscribe	Automatically subscribes this system to the best-matched compatible subscription. This is good for automated setup operations, since the system can be configured in a single step.	Optional
--activationkey=key	Applies existing subscriptions as part of the registration process. The subscriptions are pre-assigned by a vendor or by a systems administrator using Subscription Asset Manager.	Optional
--force	Registers the system even if it is already registered. Normally, any register operations will fail if the machine is already registered.	Optional

14.4.5. Unregistering

The only thing required to unregister a machine is to run the unregister command. This removes the system's entry from the subscription service, unsubscribes it from any subscriptions, and, locally, deletes its identity and entitlement certificates.

In the Red Hat Subscription Manager GUI, there is an Unregister button at the top of the window in the Tools area.

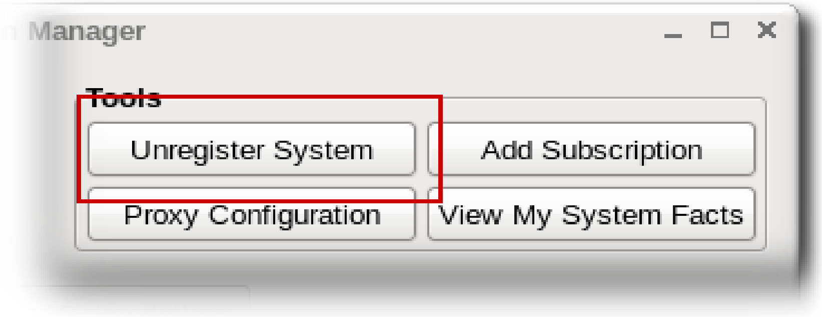

From the command line, this requires only the unregister command.

예 14.5. Unregistering a Consumer

```
[root@server1 ~]# subscription-manager unregister
```

14.4.6. Restoring a Registration

There are times when the local registration and subscription information could be lost or corrupted. There could be a hardware failure or system crash. Or other IT considerations may require that a system be moved to a different machine. Whatever the reason, the local subscription configuration is lost.

A system can be registered against an existing system entry in the Red Hat subscription service, which essentially restores or reregisters that consumer. The reregister operation uses the original consumer ID with the registration request, so that all of the previous subscriptions associated with the consumer entry are restored along with the registration.

Reregistering a system uses the register command. This command passes the original UUID for a system to issue a request to the subscription service to receive a new certificate using the same UUID. This essentially renews its previous registration.

예 14.6. Registering a System Against an Existing Identity Certificate

The register command uses the original ID to identify itself to the subscription service and restore its previous subscriptions.

```
[root@server1 ~]# subscription-manager register --username admin-example --password secret --
consumerid=7d133d55-876f-4f47-83eb-0ee931cb0a97
```

표 14.3. register Options to Reregister the System

Options	Description	Required
--consumerid	Gives the consumer UUID used by an existing consumer. The system's consumer entry must exist in the Red Hat subscription service for the reregister operation to succeed.	Required
--username=name	Gives the content server user account name.	Optional
--password=password	Gives the password for the user account.	Optional

14.5. Migrating Systems from RHN Classic to Certificate-based Red Hat Network

As described in 14.1.6절. "Certificate-based Red Hat Network versus RHN Classic" and https://access.redhat.com/kb/docs/DOC-45987, there are differences in how RHN Classic and Certificate-based Red Hat Network define and manage subscriptions.

As part of migration, the RHN Classic channels are mapped to Certificate-based Red Hat Network X.509 product certificates for every installed product. Subscription Manager can use those certificates to subscribe or autosubscribe the system to the appropriate subscriptions once it is registered.

Migration tools are available to transition system registration from RHN Classic to Certificate-based Red Hat Network and then re-apply its previous subscriptions. Product certificates in general are described in 14.15.3절. "The Structure of Product Certificates" .

There are two migration paths supported:

- From being registered with RHN Classic Hosted to being registered with Certificate-based Red Hat Network, using rhn-migrate-classic-to-rhsm
- From a disconnected (offline) system using RHN Classic-style channels to using Certificate-based Red Hat Network X.509 certificates for installed products, using install-num-migrate-to-rhsm

Important

There is no migration path from a Satellite system to Certificate-based Red Hat Network.

14.5.1. Installing the Migration Tools

The migration tools are contained in the subscription-manager-migration package. An additional package, subscription-manager-migration-data, is required to map the RHN Classic channels to Certificate-based Red Hat Network product certificates.

1. The migration tools and data are in supplementary channels. If necessary, enable the supplementary repositories, as described in 14.9절. "Working with Subscription yum Repos" .

2. Install the migration tool packages.

```
[root@server ~]# yum install subscription-manager-migration subscription-manager-migration-data
```

14.5.2. Migrating from RHN Classic to Certificate-based Red Hat Network

A system which was registered against the hosted subscription service, RHN Classic, can be migrated to Certificate-based Red Hat Network using the rhn-migrate-classic-to-rhsm script.

The general action is that it unregisters the system from RHN Classic, registers it with Certificate-based Red Hat Network, and opens Subscription Manager (either GUI or CLI) to assign subscriptions.

The rhn-migrate-classic-to-rhsm script has this syntax:

```
rhn-migrate-classic-to-rhsm [--force|--cli-only|--help|--no-auto]
```

After running migration, the system facts list what script was used for migration and what the previous system ID was.

```
[root@server ~]# subscription-manager facts --list | grep migr
migration.classic_system_id: 09876
migration.migrated_from: rhn_hosted_classic
```

This makes it easy to track the migration process for systems within the infrastructure.

예 14.7. Basic RHN Classic to Certificate-based Red Hat Network Migration

Simply running the rhn-migrate-classic-to-rhsm tool migrates the system profile and then opens the Subscription Manager GUI so that administrators can assign subscriptions to the system.

While administrators only have to run the command, the script itself runs through a series of steps to migrate the account.

```
[root@server ~]# rhn-migrate-classic-to-rhsm
RHN Username: jsmith@example.com
Password:
```

The script prompts for the username and password to use to connect to Red Hat Network. It uses these credentials to authenticate to both Red Hat Network Classic and Certificatebased Red Hat Network, to verify the account settings.

Once the account is verified, the script creates a channel list for the system.

```
Retrieving existing RHN classic subscription information ...
+----------------------------------+
System is currently subscribed to:
+----------------------------------+
rhel-i386-client-5
```

Each discovered channel is then mapped to a corresponding product certificate (14.5.5절. "Looking at Channel and Certificate Mappings"). Not every product has a product certificate, so not every channel may have a map. Only the products with a channel have a corresponding certificate map.

The matching certificates are copied into the /etc/pki/product directory.

```
List of channels for which certs are being copied
```

```
rhel-i386-client-5

Product Certificates copied successfully to /etc/pki/product !!
```

Then, the script unregisters the system from RHN Classic.

```
Preparing to unregister system from RHN classic ...
System successfully unregistered from RHN Classic.
```

Then, it registers the system with Certificate-based Red Hat Network.

```
Attempting to register system to Certificate-based RHN ...
The system has been registered with id: abcd1234
System server.example.com successfully registered to Certificate-based RHN.

Launching the GUI tool to manually subscribe the system ...
```

The last step opens the Subscription Manager GUI to the All Available Subscriptions tab so that the administrator can manually assign the subscriptions to the system.

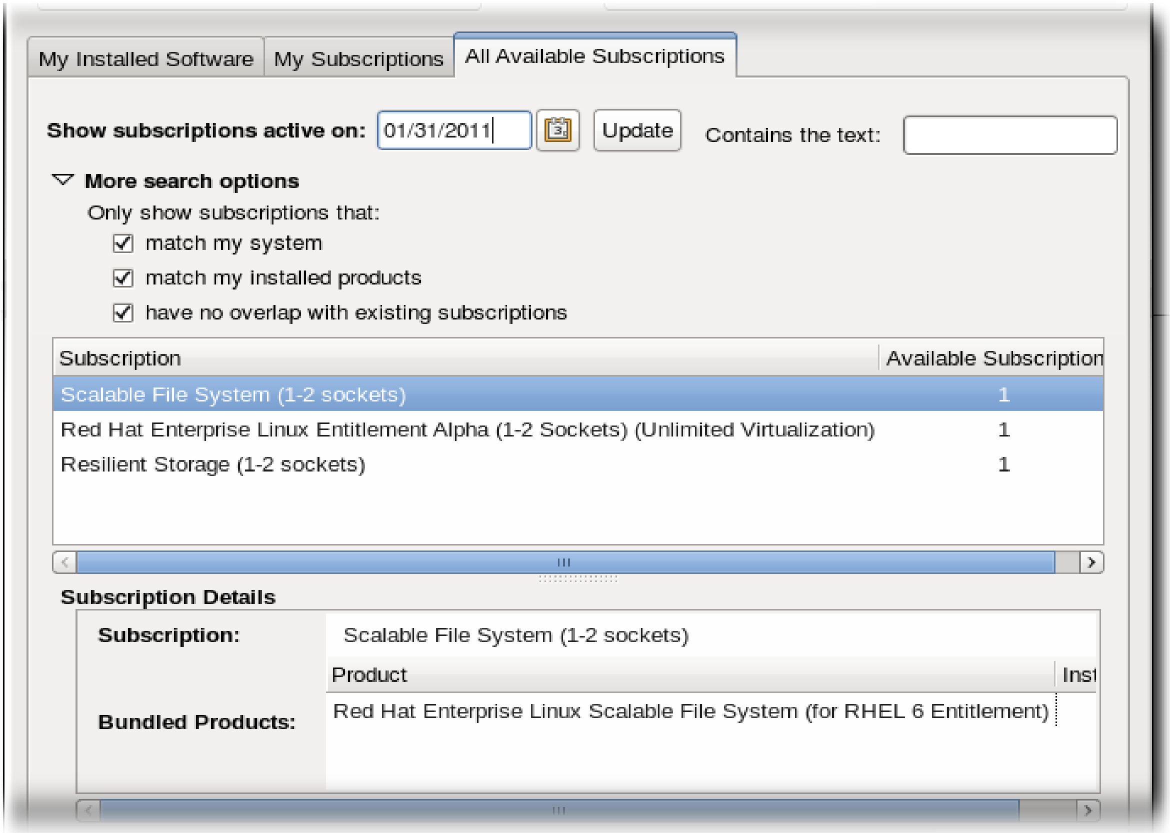

Show subscriptions active on: 01/31/2011

Contains the text:

Only show subscriptions that:

☑ match my system

☑ match my installed products

☑ have no overlap with existing subscriptions

Subscription	Available Subscription
Scalable File System (1-2 sockets)	1
Red Hat Enterprise Linux Entitlement Alpha (1-2 Sockets) (Unlimited Virtualization)	1
Resilient Storage (1-2 sockets)	1

Subscription Details

Subscription: Scalable File System (1-2 sockets)

Bundled Products:

Product	Inst
Red Hat Enterprise Linux Scalable File System (for RHEL 6 Entitlement)	

Alternatively, the rhn-migrate-classic-to-rhsm can automatically subscribe the system to matching subscriptions.

예 14.8. All CLI-Based Migration

The --cli-only option tells the rhn-migrate-classic-to-rhsm to register the system with the autosubscribe option, so all of the migration process occurs in the command line.

The overall process is identical to the one in 예 14.7. “Basic RHN Classic to Certificate-based Red Hat Network Migration” until the final step.

```
[root@server ~]# rhn-migrate-classic-to-rhsm --cli-only
RHN Username: jsmith@example.com
Password:

....

Attempting to auto-subscribe to appropriate subscriptions ...
Installed Product Current Status:
ProductName:          Red Hat Enterprise Linux Desktop
Status:              Subscribed

Please visit https://access.redhat.com/management/consumers/abcd1234 to view the details, and to make changes if necessary.
```

14.5.3. Unregistering from RHN Classic Only

There may be an instance where a system should be unregistered from RHN Classic but is not yet ready to be registered to Certificate-based Red Hat Network. The rhn-migrate-classic-to-rhsm tool can be used simply to unregister a system from RHN Classic. This still copies over the product certificates for the classic channels to configure the system in the style of certificate-based subscriptions, but it does not register the machine with the subscription service.

To unregister the system only, use the --no-auto option.

```
[root@server ~]# rhn-migrate-classic-to-rhsm --no-auto
RHN Username: jsmith@example.com
Password:

Retrieving existing RHN classic subscription information ...
+----------------------------------+
System is currently subscribed to:
+----------------------------------+
rhel-i386-client-5

List of channels for which certs are being copied
rhel-i386-client-5

Product Certificates copied successfully to /etc/pki/product !!

Preparing to unregister system from RHN classic ...
System successfully unregistered from RHN Classic.
```

Because there are product certificates, Subscription Manager will show a red, invalid status for the system and issue notifications until the system is registered and subscriptions applied.

14.5.4. Migrating a Disconnected System

Some systems may never be connected to an external network or may be prevented from accessing Red Hat Network or a Satellite system. These systems still require valid subscriptions and product certificates, though.

The rhn-migrate-classic-to-rhsm script uses the information in /etc/sysconfig/rhn/systemid to get the previous registration information and map channels to certificates. If a system is disconnected, it may not have a systemid file.

Most systems, even ones never registered with RHN Classic, do have an installation number. When Red Hat software is purchased through a vendor, the purchased software is identified in an installation number or subscription number (described in https://access.redhat.com/kb/docs/DOC-15408) in the /etc/sysconfig/rhn/install-num file.

The installation number is in essence a code which contains all of the information about the products and versions purchased for the system. For example, this installation number shows that it is valid for RHEL Client and RHEL Workstation channels.

```
[root@server ~]# python /usr/lib/python2.4/site-packages/instnum.py da3122afdb7edd23
Product: RHEL Client
Type: Installer Only
Options: Eval FullProd Workstation
Allowed CPU Sockets: Unlimited
Allowed Virtual Instances: Unlimited
Package Repositories: Client Workstation

key: 14299426 "da3122"
checksum: 175 "af"
options: 4416 "Eval FullProd Workstation"
socklimit: -1 "Unlimited"
virtlimit: -1 "Unlimited"
type: 2 "Installer Only"
product: 1 "client"

{"Workstation": "Workstation", "Base": "Client"}
```

For a system which is not connected to either RHN Classic or a Satellite system, the installation number can be used to transition the product information from the older channel-based subscription model to the X.509 certificate model, managed by Subscription Manager.

The install-num-migrate-to-rhsm script identifies the channels that a disconnected system is subscribed to and then copies in the appropriate product certificates. Simply run the command:

```
[root@server ~]# install-num-migrate-to-rhsm
```

The script copies in the product certificates for the channels into the /etc/pki/product directory.

Once the system is migrated, it can be registered remotely and have entitlement certificates installed as described in 14.4.3절. "Registering an Offline Consumer" .

Even though the system is not registered, the system facts display what script was used for migration.

```
[root@server ~]# subscription-manager facts --list | grep migr
migration.migrated_from: install_number
```

Because the system was not previously registered with RHN Classic, the migration facts do not include a system ID number.

14.5.5. Looking at Channel and Certificate Mappings

The subscription-manager-migration-data package contains a mapping file that maps RHN Classic channels to Certificate-based Red Hat Network product certificates. This file (/usr/share/rhsm/product/RHEL-5/channel-cert-mapping.txt) uses simple keys to map the values:

```
channel_name: product_name-hash-product_cert.pem
```

For example, this maps the Red Hat Enterprise Linux Client channel to the corresponding product certificate:

```
rhel-i386-client-workstation-5: Client-Workstation-i386-b0d4c042-6e31-45a9-bd94-ff0b82e43b1a-71.pem
```

During migration, that mapping is translated into product_cert.pem and the product certificate is copied into the /etc/pki/product directory. For the rhel-i386-client-workstation-5, this migrates to the 71.pem product certificate (the last two digits of the mapping).

However, many channels are available for legacy systems only or have not yet released an X.509 product certificate. In that case, the channel has no mapping.

```
jbappplatform-4.3.0-fp-i386-server-5-rpm: none
```

This can create a situation where not all channels are migrated over to Certificate-based Red Hat Network or where products are not fully subscribed.

14.6. Handling Subscriptions

Assigning a subscription to a system gives the system the ability to install and update any Red Hat product in that subscription. A subscription is a list of all of the products, in all variations, that were purchased at one time, and it defines both the products and the number of times that subscription can be used (the quantity of that product). The quantity is roughly the number of user licenses available. When one of those licenses is allocated to a system, that system is subscribed to the subscription.

A subscription is available to a system based on the system's architecture and other installed products. Subscriptions that are available for a platform (based on its hardware and operating system) are compatible. When the subscription is actually assigned to the machine, the subscription is consumed.

A system can be subscribed to multiple subscriptions, a single subscription, or a single product. Subscribing a system requires the ID number of the subscription or the subscription key for the product.

Unsubscribing a machine removes the entitlement to any of the products in the subscription, but the machine remains registered with the subscription service. Unsubscribing one system frees the subscription so that it can be allocated to another system.

14.6.1. Subscribing and Unsubscribing through the GUI

14.6.1.1. Subscribing to a Product

1. Launch Subscription Manager. For example:

   ```
   [root@server ~]# subscription-manager-gui
   ```

2. Open the All Available Subscriptions tab.

3. Set the filters to use to search for available entitlements and click Update. Subscriptions can be filtered by their active date and by their name. The checkboxes provide more fine-grained filtering:

   - match my system shows only subscriptions which match the system architecture.

   - match my installed products shows subscriptions which work with currently installed products on the system.

   - have no overlap with existing subscriptions excludes subscriptions with duplicate products. If a system is already subscribed to an entitlement for a specific product or if multiple entitlements supply the same product, then the subscription service filters those subscriptions and shows only the best fit.

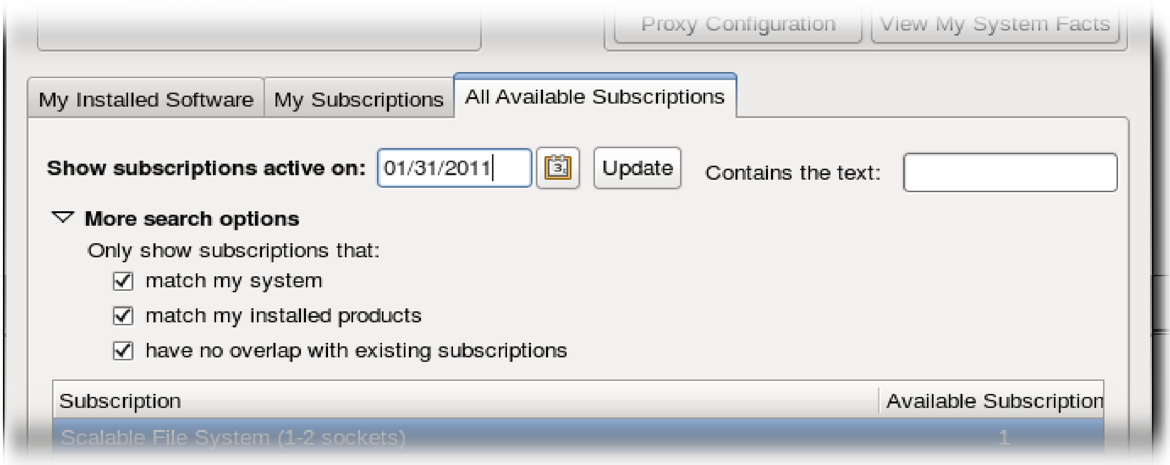

Show subscriptions active on: 01/31/2011 Contains the text:

▽ **More search options**

Only show subscriptions that:

☑ match my system

☑ match my installed products

☑ have no overlap with existing subscriptions

Subscription	Available Subscription
Scalable File System (1-2 sockets)	1

4. Select one of the available entitlements.

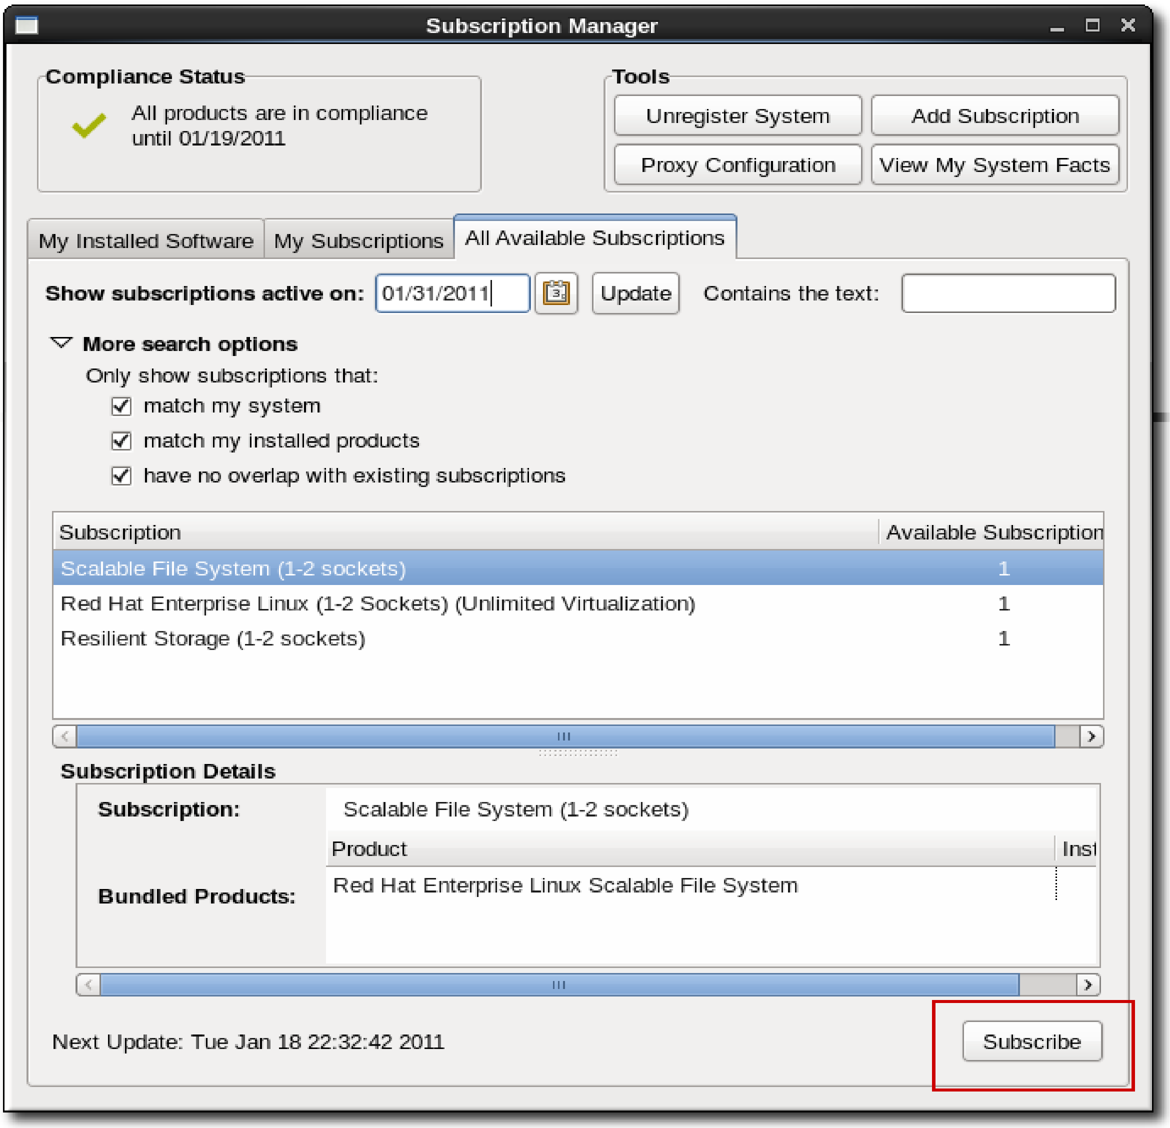

Compliance Status

✓ All products are in compliance until 01/19/2011

Show subscriptions active on: 01/31/2011 Contains the text:

▽ **More search options**

Only show subscriptions that:

☑ match my system

☑ match my installed products

☑ have no overlap with existing subscriptions

Subscription	Available Subscription
Scalable File System (1-2 sockets)	1
Red Hat Enterprise Linux (1-2 Sockets) (Unlimited Virtualization)	1
Resilient Storage (1-2 sockets)	1

Subscription Details

Subscription: Scalable File System (1-2 sockets)

Bundled Products:

Product	Inst
Red Hat Enterprise Linux Scalable File System	

Next Update: Tue Jan 18 22:32:42 2011

5. Click the Subscribe button.

14.6.1.2. Unsubscribing through the GUI

1. Launch Subscription Manager. For example:

```
[root@server ~]# subscription-manager-gui
```

2. Open the My Subscriptions tab.

   All of the active subscriptions to which the system is currently subscribed are listed. (The products available through the subscription may or may not be installed.)

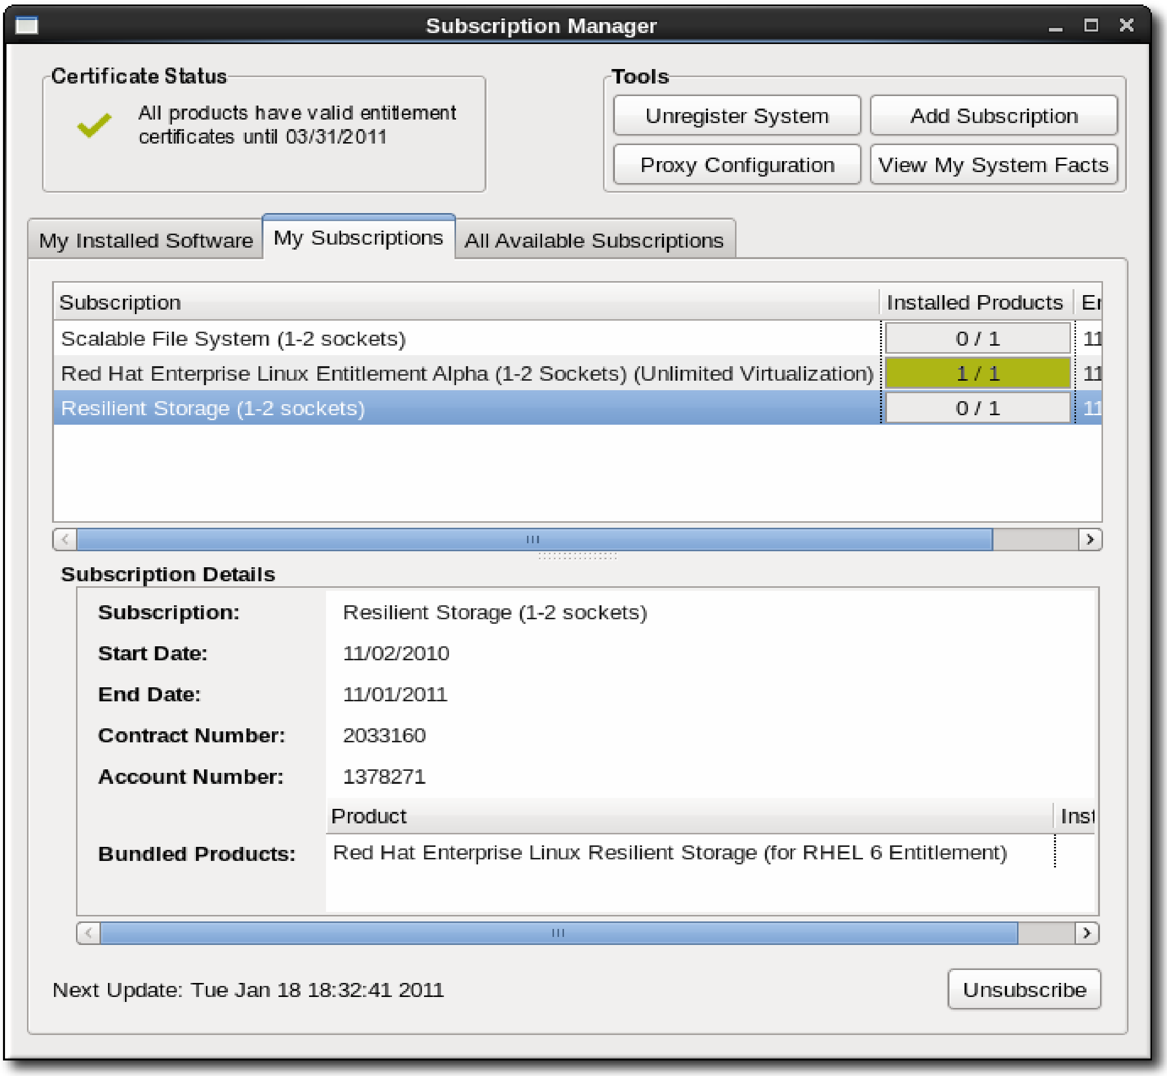

All products have valid entitlement certificates until 03/31/2011

Subscription	Installed Products	En
Scalable File System (1-2 sockets)	0 / 1	11
Red Hat Enterprise Linux Entitlement Alpha (1-2 Sockets) (Unlimited Virtualization)	1 / 1	11
Resilient Storage (1-2 sockets)	0 / 1	11

Subscription: Resilient Storage (1-2 sockets)

Start Date: 11/02/2010

End Date: 11/01/2011

Contract Number: 2033160

Account Number: 1378271

Bundled Products:

Product	Inst
Red Hat Enterprise Linux Resilient Storage (for RHEL 6 Entitlement)	

Next Update: Tue Jan 18 18:32:41 2011

3. Select the entitlement to unsubscribe.

4. Click the Unsubscribe button in the bottom right of the window.

14.6.2. Handling Subscriptions through the Command Line

14.6.2.1. Subscribing from the Command Line

Subscribing a machine through the command line requires specifying the individual product or subscription to subscribe to, using the --pool option.

```
[root@server1 ~]# subscription-manager subscribe --pool=XYZ01234567
```

The options for the subscribe command are listed in 표 14.4. “subscribe Options” .

The ID of the subscription pool for the purchased product must be specified, and this pool ID is listed with the product subscription information, from running the list command:

```
[root@server1 ~]# subscription-manager list --available

+-------------------------------------------+
    Available Subscriptions
+-------------------------------------------+


ProductName:            RHEL for Physical Servers
ProductId:              MKT-rhel-server
PoolId:                 ff8080812bc382e3012bc3845ca000cb
Quantity:               10
Expires:                2011-09-20
```

Alternatively, the system can be subscribed to the best-fitting subscriptions, as identified by the subscription service, by using the --auto option (which is analogous to the --autosubscribe option with the register command).

```
[root@server1 ~]# subscription-manager subscribe --auto
```

표 14.4. subscribe Options

Options	Description	Required
--pool=pool-id	Gives the ID for the subscription to subscribe the machine to.	Required, unless --auto is used
--auto	Automatically subscribes the system to the best-match subscription or subscriptions.	Optional
--quantity=number	Subscribes multiple counts of an entitlement to the system. This is used to cover subscriptions that define a count limit, like using two 2-socket server subscriptions to cover a 4-socket machine.	Optional

14.6.2.2. Unsubscribing from the Command Line

A system can be subscribed to multiple subscriptions and products. Similarly, the system can be unsubscribed from a single subscription or product or from every subscribed product.

Running the unsubscribe command with the --all option unsubscribes the system from every product and subscription pool it is currently subscribed to.

```
[root@server1 ~]# subscription-manager unsubscribe --all
```

It is also possible to unsubscribe from a single product. Each product has an identifying X.509 certificate installed with it. The product to unsubscribe is identified in the unsubscribe command by referencing the ID number of that X.509 certificate.

1. Get the serial number for the product certificate, if you are unsubscribing from a single product. The serial number can be obtained from the entitlement#.pem file (for example, 392729555585697907.pem) or by using the list command. For example:

```
[root@server1 ~]# subscription-manager list --consumed

+-------------------------------------------+
    Consumed Product Subscriptions
+-------------------------------------------+


ProductName:            High availability (cluster suite)
ContractNumber:         0
SerialNumber:           11287514358600162
Active:                 True
Begins:                 2010-09-18
Expires:                2011-11-18
```

2. Run the subscription-manager tool with the --serial option to specify the certificate.

```
[root@server1 ~]# subscription-manager unsubscribe --serial=11287514358600162
```

14.6.3. Stacking Subscriptions

Some subscriptions define a count which works as a restriction on the subscription. For example, counts can be set on the number of sockets or CPUs on a machine, the number of virtual guests on a host, or the number of clients in a domain.

The entire count must be covered for the system to be fully entitled. If there are four sockets on a machine, then the server subscriptions must cover four sockets, or if there are eight guests, then there must be enough to cover all eight guests.

Many subscriptions can be combined together to cover the count on the system. Two subscriptions for RHEL Server for 2-Sockets can be combined together to cover a four-socket machine. These subscriptions can be stacked.

There are some rules on what subscriptions can be stacked:

- Subscriptions can be stacked by using multiple quantities from the same subscription set.
- Subscriptions from different contracts can be stacked together.
- Only the same product subscription can be stacked. RHEL Server for 2-Sockets can be stacked with another RHEL Server for 2-Sockets subscription, but not with RHEL Server for Virtualization, even if they both cover the socket count.

Stackable entitlements are indicated in the Subscription Manager UI with an asterisk (*). In the UI, available subscriptions are grouped first by what subscriptions are compatible for stacking, and then by other available subscriptions.

To stack subscriptions in the Subscription Manager UI, simply set the Quantity field to the required quantity to cover the count.

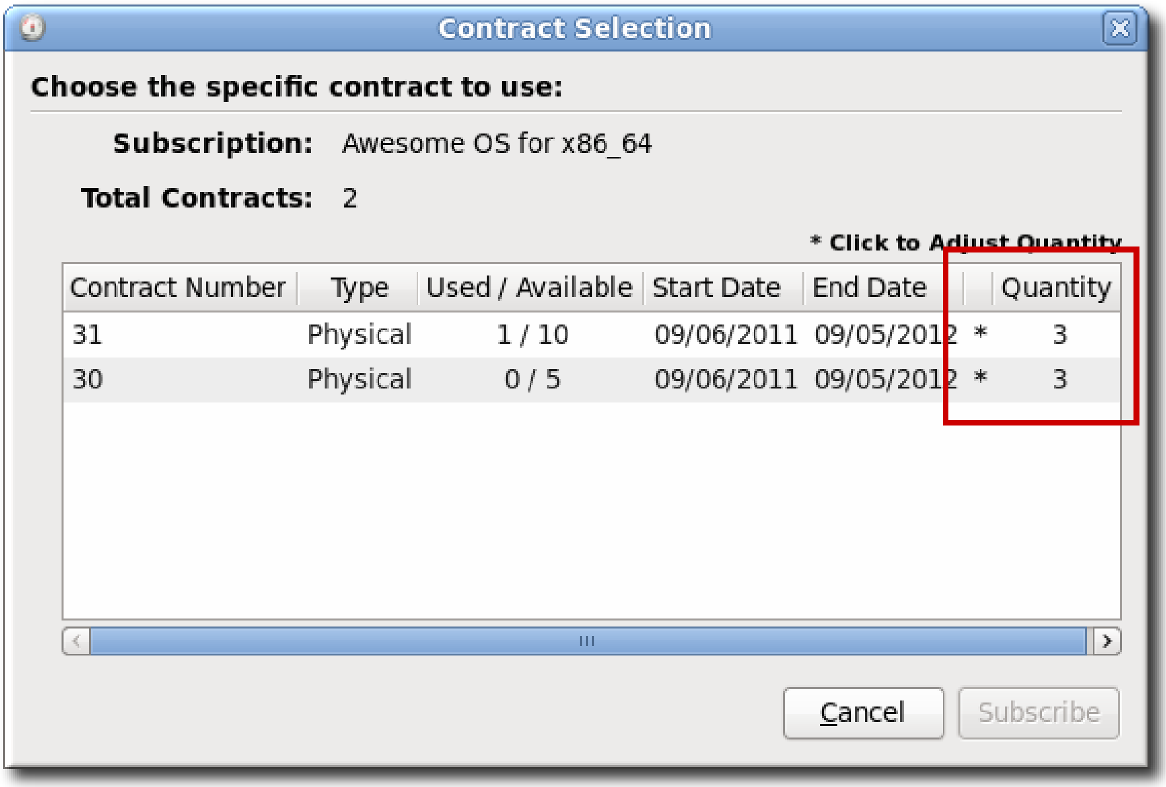

Choose the specific contract to use:

Subscription: Awesome OS for x86_64

Total Contracts: 2

* Click to Adjust Quantity

Contract Number	Type	Used / Available	Start Date	End Date		Quantity
31	Physical	1 / 10	09/06/2011	09/05/2012	*	3
30	Physical	0 / 5	09/06/2011	09/05/2012	*	3

그림 14.12. Stacking Quantities

To stack subscriptions from the command line, use the --quantity option. The quantity taken applies to the product in the --pool option:

```
[root@server1 ~]# subscription-manager subscribe --pool=XYZ01234567 --quantity=2
```

14.6.4. Manually Adding a New Subscription

In certain situations, new product subscriptions can be added by uploading the X.509 entitlements certificate directly rather than polling the subscription service. For example, consumers which are offline must have subscriptions manually added because they cannot connect to the subscription service directly.

1. Retrieve the certificate information for the consumer from the Customer Portal.

   a. Open the Subscriptions tab in the Customer Portal, and select the Overview item under the Certificate-based Management area.

   b. In the summary of consumers, click the name of the offline consumer.

   c. If necessary, assign the subscriptions to the consumer.

   d. Open the Applied Subscriptions tab.

   e. Click the Download All Certificates button. This exports all of the entitlements certificates, for each product, to a single .zip file. Save the file to some kind of portable media device, like a flash drive.

To download individual entitlement certificates, click the Download link on the row for the subscription.

2. Copy the certificates over to the consumer machine.

3. If all certificates were downloaded in an archive file, then there are multiple archives in the downloaded certificates.zip file. Unzip the directories until the PEM files for the subscription certificates are available.

4. Import the certificates.

   This can be done from the command line using the import command:

```
# subscription-manager import --certificate=/tmp/export/entitlement_certificates/596576341785244687.pem --certificate=/
tmp/export/entitlement_certificates/3195996649750311162.pem
Successfully imported certificate 596576341785244687.pem
Successfully imported certificate 3195996649750311162.pem
```

   This can also be performed through the Subscription Manager GUI:

   a. Launch Subscription Manager. For example:

```
[root@server ~]# subscription-manager-gui
```

   b. In the Tools area, click the Import Certificate button.

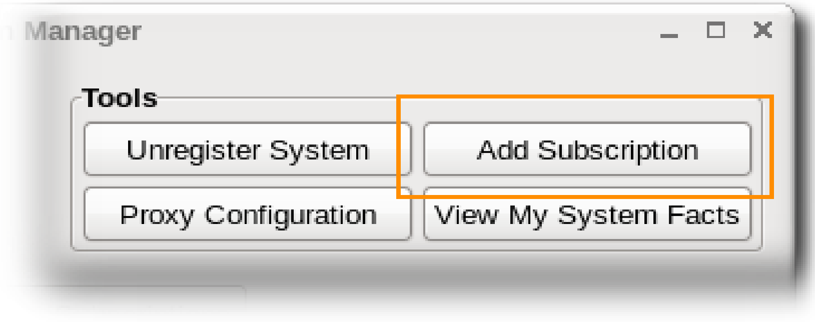

   c. Click the file folder icon at the right of the field to navigate to the .pem file of the product certificate.

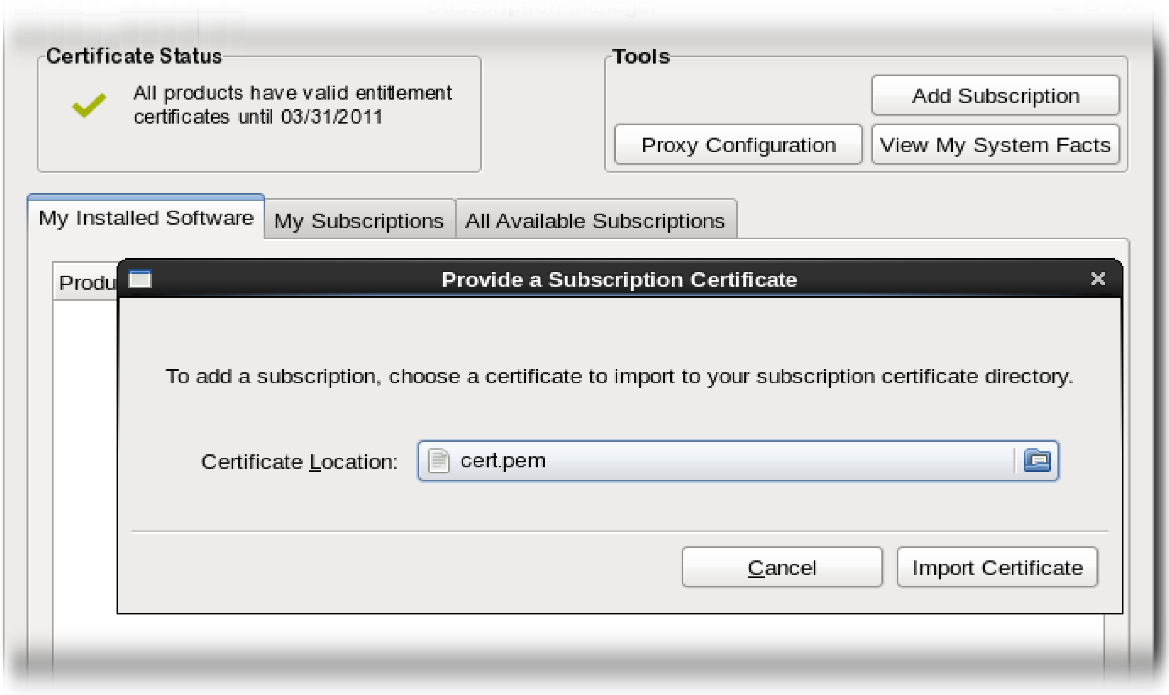

To add a subscription, choose a certificate to import to your subscription certificate directory.

Certificate Location: cert.pem

d. Click the Import Certificate button.

The consumer is then entitled for all of the subscription that were uploaded.

14.7. Redeeming Subscriptions on a Machine

Systems can be set up with pre-existing subscriptions already available to that system. For some systems which were purchased through third-party vendors, a subscription to Red Hat products is included with the purchase of the machine. Companies using the Subscription Asset Manager can allocate subscriptions to their own systems by creating activation keys which are used to claim those assigned subscriptions.

Red Hat Subscription Manager pulls information about the system hardware and the BIOS into the system facts to recognize the hardware vendor. If the vendor and BIOS information matches a certain configuration, then the subscription can be redeemed, which will allow the system to be automatically subscribed to the entitlements purchased with the machine.

This diverges from the normal subscription process by adding an extra step:

1. The machine is registered first (14.4절. "Registering, Unregistering, and Reregistering a System"). This can be done as normal or the activation keys can be submitted with command-line registrations.

2. The subscriptions are redeemed using the given activation keys.

3. The system is then subscribed to its subscriptions (14.6절. "Handling Subscriptions").

참고

Activation keys may be generated by a hardware vendor (external to your organization). Activation keys may also be generated using the Subscription Asset Manager, which is a local subscription service, described in the Subscription Asset Manager documentation[2] and 14.12절. "Working with Subscription Asset Manager" .

14.7.1. Redeeming Subscriptions through the GUI

참고

If the machine does not have any subscriptions to be redeemed, then the Redeem a Subscription button is not there.

1. Launch Subscription Manager. For example:

```
[root@server ~]# subscription-manager-gui
```

2. At the top of the main window, click the Redeem a Subscription button.

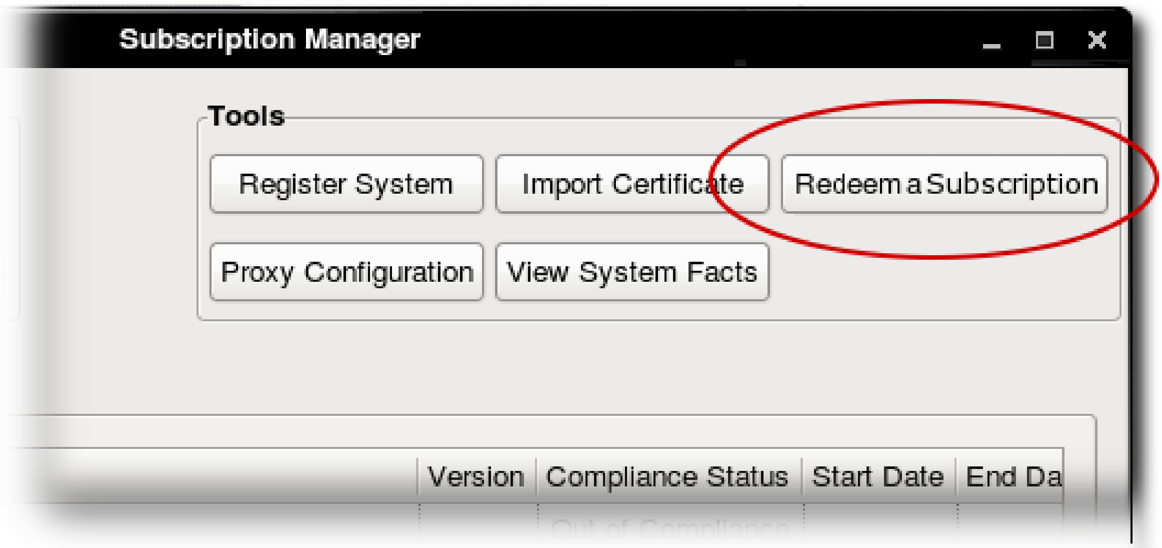

3. In the dialog window, enter the email address to send the notification to when the redemption is complete.

[2] http://docs.redhat.com/docs/en-US/Red_Hat_Subscription_Asset_Manager/1.0/html/Installation_Guide/index.html

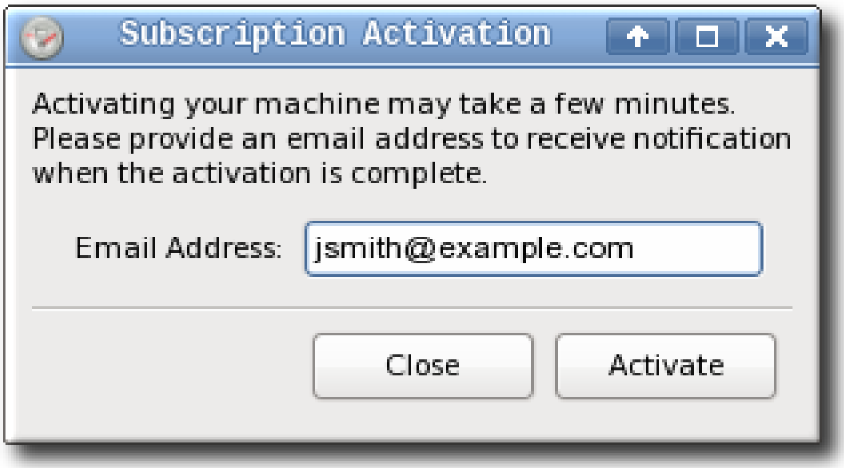

Activating your machine may take a few minutes. Please provide an email address to receive notification when the activation is complete.

Email Address: jsmith@example.com

4. Click the Redeem button.

It can take up to ten minutes for the confirmation email to arrive.

14.7.2. Redeeming Subscriptions on a Machine through the Command Line

The machine subscriptions are redeemed by running the redeem command, with an email address to send the redemption email to when the process is complete.

```
# subscription-manager redeem --email=jsmith@example.com
```

In a multi-organization environment, it is also necessary to specify the organization which issued the activation keys. For example:

```
# subscription-manager redeem --email=jsmith@example.com --org="IT Dept"
```

참고

The machine must be registered first so that the subscription service can properly identify the system and its subscriptions.

14.8. Viewing Available and Used Subscriptions

To manage subscriptions, administrators need to know both what subscriptions a system is currently consuming and what subscriptions are available to the system.

14.8.1. Viewing Subscriptions in the GUI

The Red Hat Subscription Manager tools give a more detailed view of subscriptions and entitlements than is available through the global tools of the Customer Portal. Three tabs summarize each of the subscriptions and products for the specific machine: installed products (with subscriptions), subscribed entitlements, and available subscriptions.

These summaries are always displayed in the Red Hat Subscription Manager UI.

Subscribed Entitlements

The My Subscriptions tab shows all of the current entitlements that the system is subscribed to.

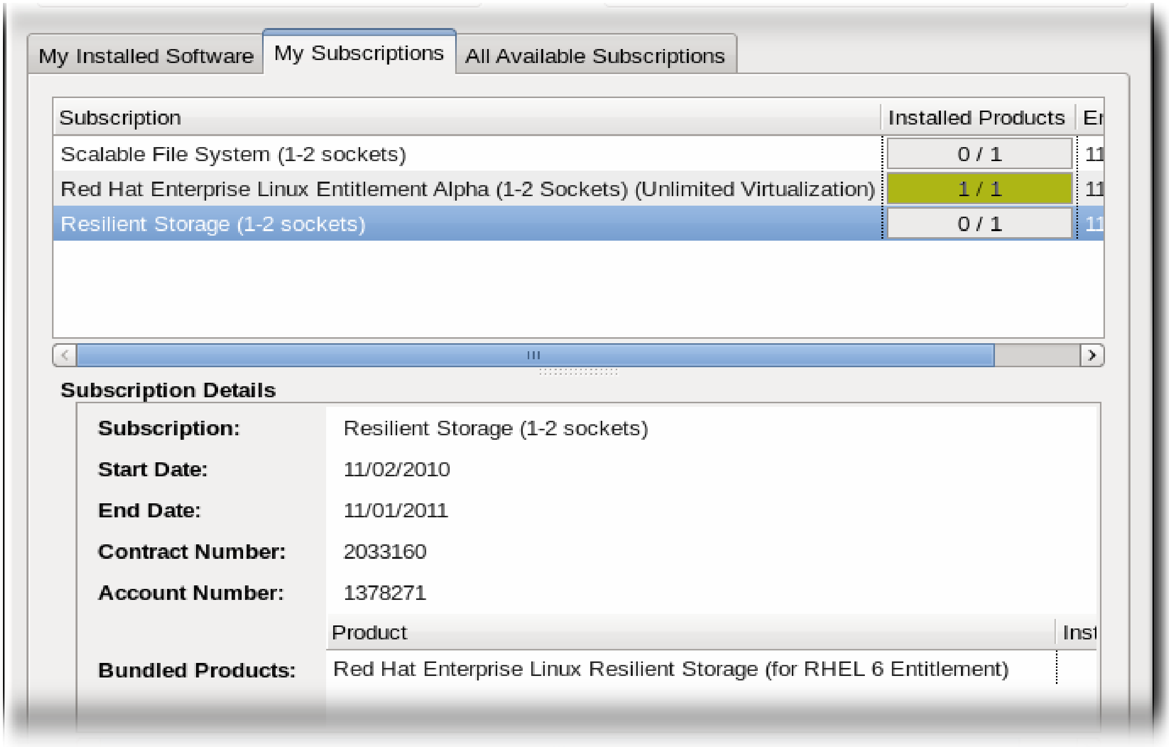

Subscription	Installed Products	Er
Scalable File System (1-2 sockets)	0 / 1	11
Red Hat Enterprise Linux Entitlement Alpha (1-2 Sockets) (Unlimited Virtualization)	1 / 1	11
Resilient Storage (1-2 sockets)	0 / 1	11

Subscription Details

Subscription: Resilient Storage (1-2 sockets)

Start Date: 11/02/2010

End Date: 11/01/2011

Contract Number: 2033160

Account Number: 1378271

Bundled Products:

Product	Inst
Red Hat Enterprise Linux Resilient Storage (for RHEL 6 Entitlement)	

그림 14.13. My Subscriptions Tab

Available Subscriptions

The All Available Subscriptions tab shows all of the subscriptions that are available to the system. The default displays only entitlements that are compatible with the hardware, but these can be filtered to show entitlements corresponding to other installed programs, only subscriptions that have not been installed, and subscriptions based on date.

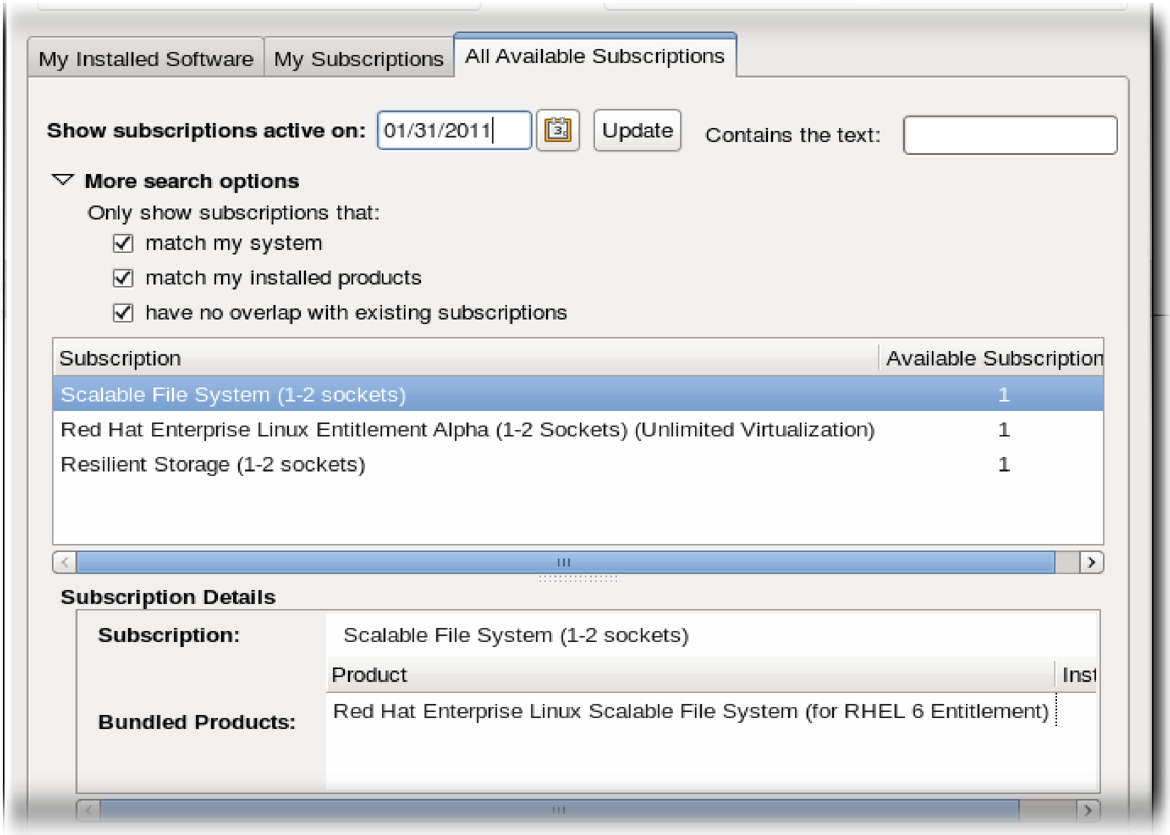

Show subscriptions active on: 01/31/2011 Update Contains the text:

More search options

Only show subscriptions that:

- ☑ match my system
- ☑ match my installed products
- ☑ have no overlap with existing subscriptions

Subscription	Available Subscription
Scalable File System (1-2 sockets)	1
Red Hat Enterprise Linux Entitlement Alpha (1-2 Sockets) (Unlimited Virtualization)	1
Resilient Storage (1-2 sockets)	1

Subscription Details

Subscription: Scalable File System (1-2 sockets)

Bundled Products:

Product	Inst
Red Hat Enterprise Linux Scalable File System (for RHEL 6 Entitlement)	

그림 14.14. All Available Subscriptions Tab

The filters dynamically search for available entitlements. Subscriptions can be filtered by their active date and by their name. The checkboxes provide more fine-grained filtering:

- match my system shows only subscriptions which match the system architecture.
- match my installed products shows subscriptions which work with currently installed products on the system.
- have no overlap with existing subscriptions excludes subscriptions with duplicate products. If a system is already subscribed to an entitlement for a specific product or if multiple entitlements supply the same product, then the subscription service filters those subscriptions and shows only the best fit.

My Installed Software

The My Installed Software tab shows the currently installed products on the system, along with their subscription status. This does not allow administrators to install software, only to view installed software.

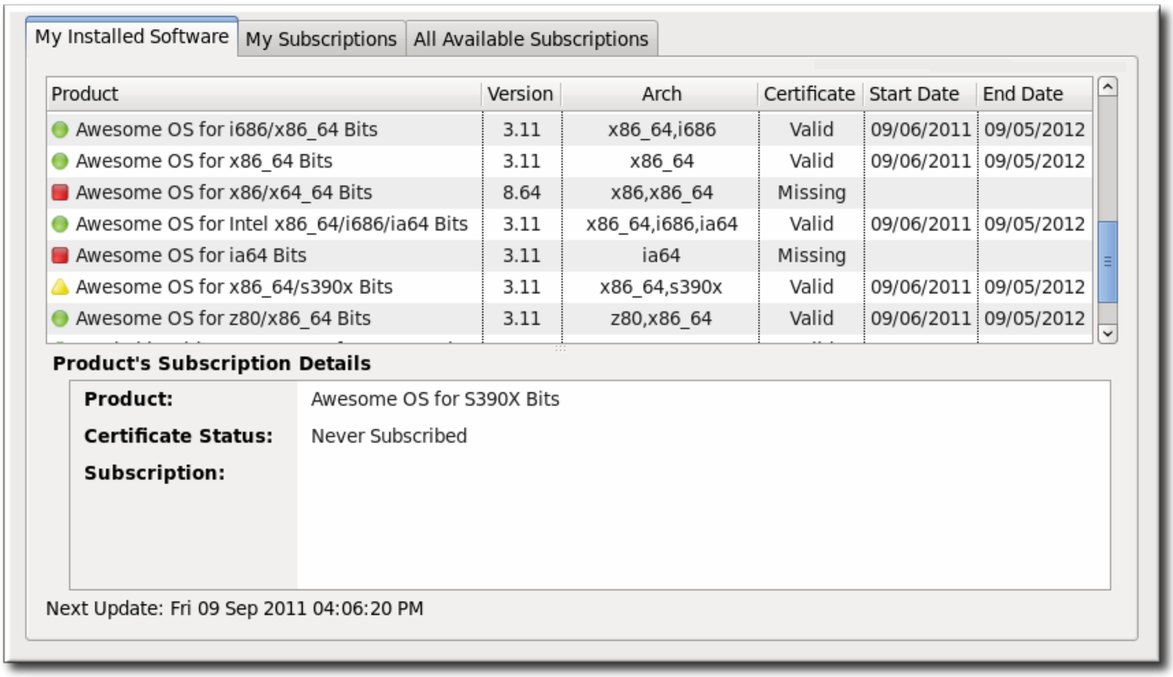

Product	Version	Arch	Certificate	Start Date	End Date
🟢 Awesome OS for i686/x86_64 Bits	3.11	x86_64,i686	Valid	09/06/2011	09/05/2012
🟢 Awesome OS for x86_64 Bits	3.11	x86_64	Valid	09/06/2011	09/05/2012
🟥 Awesome OS for x86/x64_64 Bits	8.64	x86,x86_64	Missing		
🟢 Awesome OS for Intel x86_64/i686/ia64 Bits	3.11	x86_64,i686,ia64	Valid	09/06/2011	09/05/2012
🟥 Awesome OS for ia64 Bits	3.11	ia64	Missing		
⚠ Awesome OS for x86_64/s390x Bits	3.11	x86_64,s390x	Valid	09/06/2011	09/05/2012
🟢 Awesome OS for z80/x86_64 Bits	3.11	z80,x86_64	Valid	09/06/2011	09/05/2012

Product's Subscription Details

Product: Awesome OS for S390X Bits

Certificate Status: Never Subscribed

Subscription:

Next Update: Fri 09 Sep 2011 04:06:20 PM

그림 14.15. My Installed Software Tab

14.8.2. Listing Subscriptions with the Command Line

As with the three tabs in the UI, there are several different ways to use the list command to display different areas of the subscriptions and products on the system.

표 14.5. subscription-manager list Options

Option	Description
--installed (or nothing)	Lists all of the installed and subscribed product on the system. If no option is given with list, it is the same as using the --installed argument.
--consumed	Lists all of the subscriptions allocated to the system.
--available [--all]	Using --available alone lists all of the compatible, active subscriptions for the system. Using --available --all lists all options, even ones not compatible with the system or with no more available quantities.
--ondate=YYYY-MM-DD	Shows subscriptions which are active and available on the specified date. This is only used with the --available option. If this is not used, then the command uses the current date.
--installed	Lists all of the products that are installed on the system (and whether they have a subscription) and it lists all of the product subscriptions which are assigned to the system (and whether those products are installed).

The list command shows all of the subscriptions that are currently allocated to the system by using the --consumed option.

```
[root@server1 ~]# subscription-manager list --consumed

+-------------------------------------------+
    Consumed Product Subscriptions
```

```
+-------------------------------------------+


ProductName:         Red Hat Enterprise Linux Server
ContractNumber:      1458961
SerialNumber:        171286550006020205
Active:              True
Begins:              2009-01-01
Expires:             2011-12-31
```

The list command shows all of the subscriptions that are compatible with and available to the system using the --available option. To include every subscription the organization has — both the ones that are compatible with the system and for other platforms — use the --all option with the --available. The --ondate option shows only subscriptions which are active on that date, based on their activation and expiry dates.

```
[root@server1 ~]# subscription-manager list --available --all

+-------------------------------------------+
    Available Subscriptions
+-------------------------------------------+


ProductName:            RHEL for Physical Servers
ProductId:              MKT-rhel-server
PoolId:                 ff8080812bc382e3012bc3845ca000cb
Quantity:               10
Expires:                2011-09-20


ProductName:            RHEL Workstation
ProductId:              MKT-rhel-workstation-mkt
PoolId:                 5e09a31f95885cc4
Quantity:               10
Expires:                2011-09-20

[snip]
```

The --installed option correlates the products that are actually installed on the system (and their subscription status) and the products which could be installed on the system based on the assigned subscriptions (and whether those products are installed).

```
[root@server1 ~]# subscription-manager list --installed

+-------------------------------------------+
    Installed Product Status
+-------------------------------------------+
ProductName:         Red Hat Enterprise Linux
Status:              Not Subscribed
Expires:
Subscription:
ContractNumber:
AccountNumber:


ProductName:         Awesome OS Server
Status:              Not Installed
Expires:             2012-02-20
Subscription:        54129829316535230
ContractNumber:      39
AccountNumber:       12331131231
```

14.8.3. Viewing Subscriptions Used in Both RHN Classic and Certificate-based Red Hat Network

Administrators need to have a sense of all of the subscriptions allocated for their organization, altogether, regardless of whether the system is managed in RHN Classic or Certificate-based Red Hat Network. The Customer Portal provides a way of looking at the total consumed subscriptions.

In the Subscriptions Overview page, the Subscription Utilization area at the top gives the current count for every active subscription for the entire organization, and a total count of every used subscription, regardless of whether it is used in RHN Classic or Certificate-based Red Hat Network. These numbers are updated whenever the subscription count changes in the subscription server. (The subsequent Certificate-based Red Hat Network and RHN Classic sections gives usage subcounts based on system which are registered to that specific subscription service.)

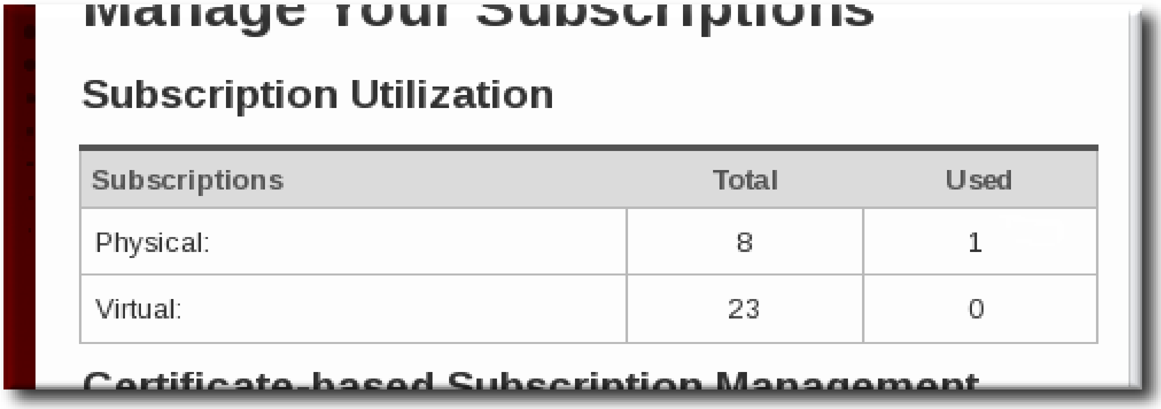

Manage Your Subscriptions

Subscription Utilization

Subscriptions	Total	Used
Physical:	8	1
Virtual:	23	0

Certificate-based Subscription Management

그림 14.16. Total Counts of Subscriptions for All Subscription Services

참고

RHN Classic is provided for legacy systems. Red Hat Enterprise Linux 5.7 and 6.1 and later systems should use Certificate-based Red Hat Network to manage subscriptions for systems.

14.9. Working with Subscription yum Repos

As 14.1.4절. "Subscription and Content Architecture" describes, Red Hat Subscription Manager works with yum. Subscription Manager has its own yum plug-ins: product-id for subscription-related information for products and subscription-manager which is used for the content repositories.

As systems are subscribed to products, the associated content repositories (identified in the entitlement certificate) are made available to the system. The content repositories are based on the product and on the content delivery network, defined in the baseurl parameter of the rhsm.conf file.

A subscription may include access to optional content channels along with the default channels. These optional channels must be enabled before the packages in them can be installed (even if the system is fully entitled to the products in those channels).

1. List all available repos for the system, including disabled repos.

```
[root@server ~]# yum repolist all
```

```
repo id                     repo name                          status
rhel-5-server               Red Hat Enterprise Linux 5Server - enabled: 1,749
rhel-5-server-beta          Red Hat Enterprise Linux 5Server Be enabled:  869
rhel-5-server-optional-rpms Red Hat Enterprise Linux 5Server Op disabled
rhel-5-server-supplementary Red Hat Enterprise Linux 5Server Su disabled
```

The optional and supplementary channels are named rhel-5-server-optional-rpms and rhel-5-server-supplementary, respectively.

2. The repositories can be enabled using the yum-config-manager command:

```
[root@server ~]# yum-config-manager --enable rhel-5-server-optional-rpms
```

Alternatively, simply specify the optional or supplementary repository when installing a package with yum. This uses the --enablerepo repo_name option. For example:

```
# yum install rubygems --enablerepo=rhel-5-server-optional-rpms
Loaded plugins: product-id, refresh-packagekit, subscription-manager
Updating Red Hat repositories.
....
```

Using yum is described in 13장. YUM (Yellowdog Updater Modified).

14.10. Responding to Subscription Notifications

The Red Hat Subscription Manager provides a series of log and UI messages that indicate any changes to the valid certificates of any installed products for a system. In the Subscription Manager GUI, the status of the system entitlements is color-coded, where green means all products are fully subscribed, yellow means that some products may not be subscribed but updates are still in effect, and red means that updates are disabled.

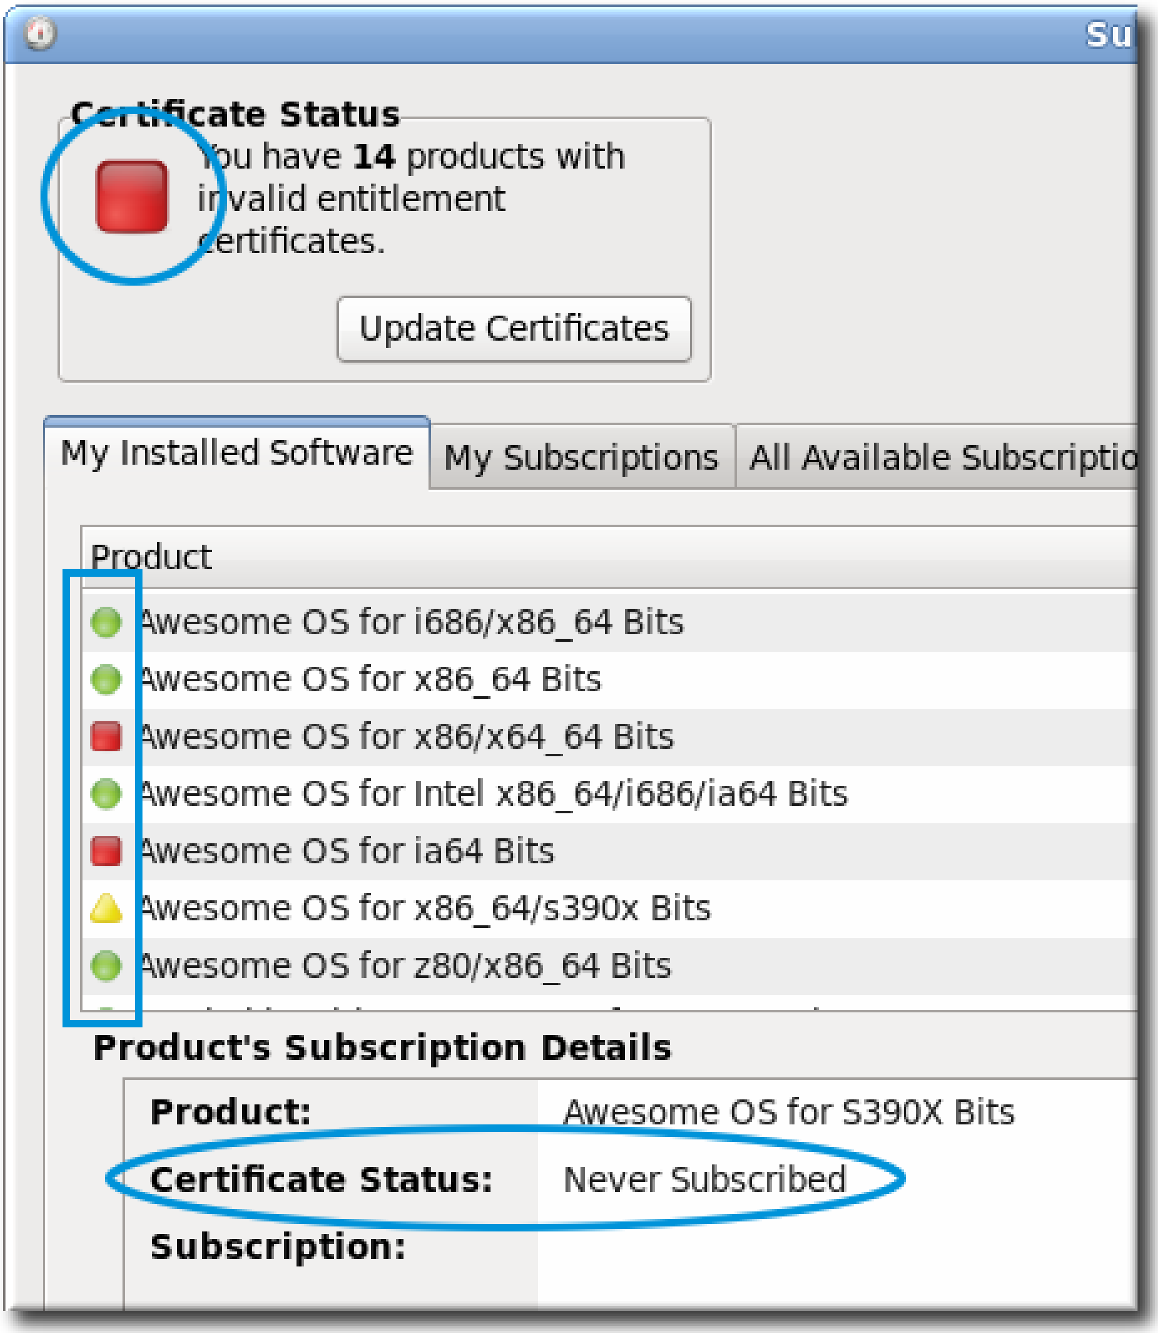

You have **14** products with invalid entitlement certificates.

- ● Awesome OS for i686/x86_64 Bits
- ● Awesome OS for x86_64 Bits
- 🟥 Awesome OS for x86/x64_64 Bits
- ● Awesome OS for Intel x86_64/i686/ia64 Bits
- 🟥 Awesome OS for ia64 Bits
- △ Awesome OS for x86_64/s390x Bits
- ● Awesome OS for z80/x86_64 Bits

Product: Awesome OS for S390X Bits

Certificate Status: Never Subscribed

Subscription:

그림 14.17. Color-Coded Status Views

The command-line tools also indicate that status of the machine. The green, yellow, and red codes translate to text status messages of subscribed, partially subscribed, and expired/not subscribed, respectively.

```
[root@server ~]# subscription-manager list
+-------------------------------------------+
    Installed Product Status
+-------------------------------------------+

ProductName:            Red Hat Enterprise Linux Server
Status: Not Subscribed
Expires:
SerialNumber:
ContractNumber:
```

```
AccountNumber:
```

Whenever there is a warning about subscription changes, a small icon appears in the top menu bar, similar to a fuel gauge.

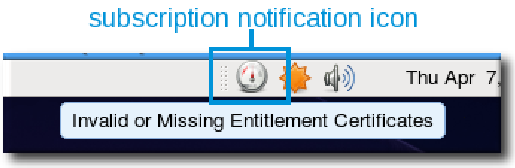
subscription notification icon
Thu Apr 7,
Invalid or Missing Entitlement Certificates

그림 14.18. Subscription Notification Icon

As any installed product nears the expiration date of the subscription, the Subscription Manager daemon will issue a warning. A similar message is given when the system has products without a valid certificate, meaning either the system is not subscribed to a subscription that entitles that product or the product is installed past the expiration of the subscription. Clicking the Manage My Subscriptions... button in the subscription notification window opens the Red Hat Subscription Manager GUI to view and update subscriptions.

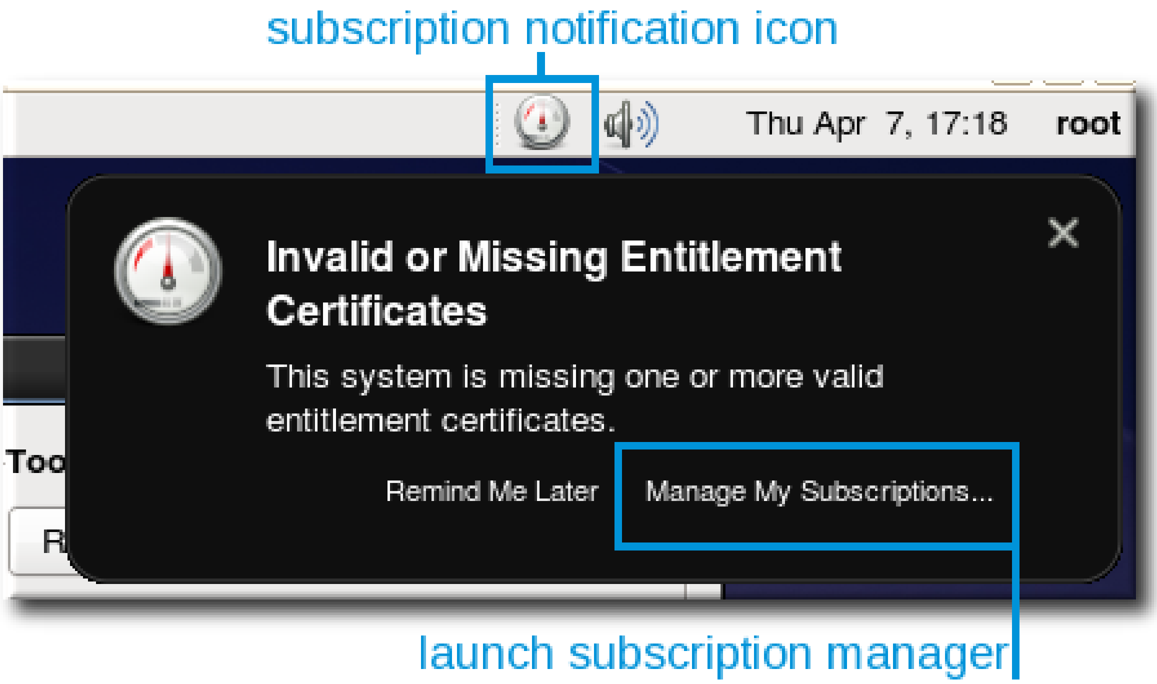
subscription notification icon
Thu Apr 7, 17:18
root
Remind Me Later
Manage My Subscriptions...
launch subscription manager

Invalid or Missing Entitlement Certificates

This system is missing one or more valid entitlement certificates.

그림 14.19. Subscription Warning Message

When the Subscription Manager UI opens, whether it was opened through a notification or just opened normally, there is an icon in the upper left corner that shows whether products lack a valid certificate. The easiest way to allocate subscriptions which match invalidated products is to click the Update Certificates button.

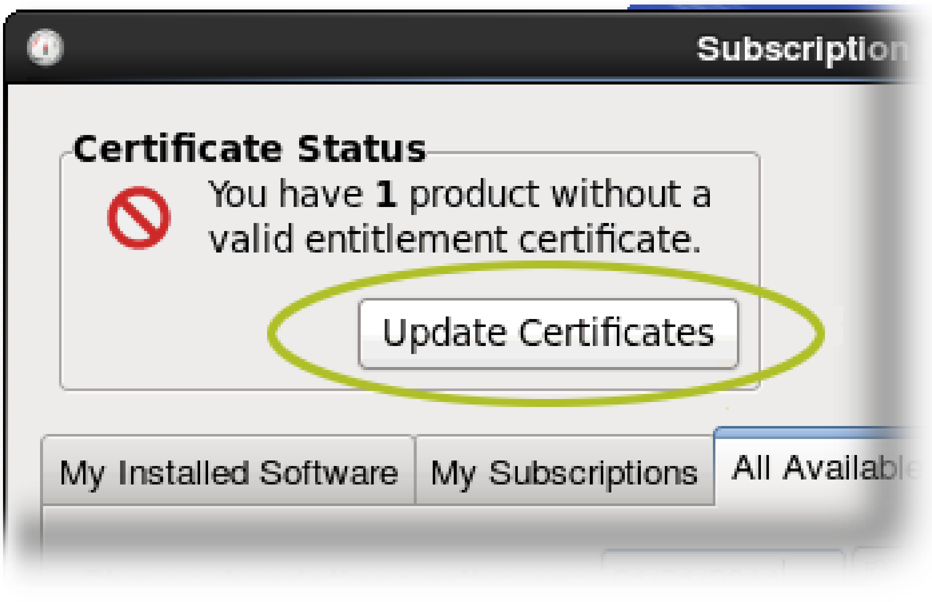

Certificate Status

You have **1** product without a valid entitlement certificate.

그림 14.20. Update Certificates Button

The Subscription Assistant dialog shows a targeted list of available subscriptions that apply to the specific products that do not have valid certificates (assuming subscriptions are available).

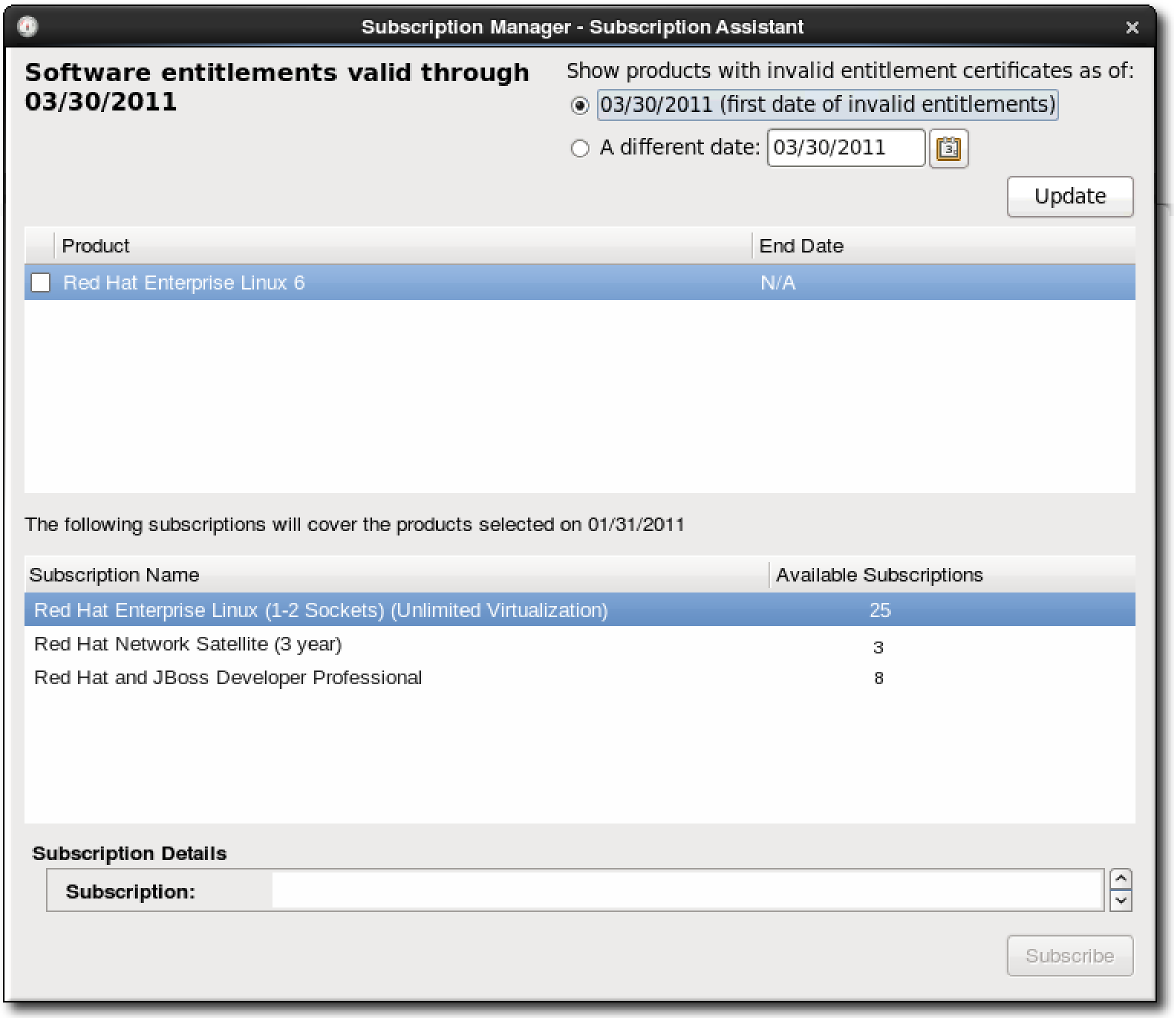

Software entitlements valid through 03/30/2011

Show products with invalid entitlement certificates as of:

● 03/30/2011 (first date of invalid entitlements)

○ A different date: 03/30/2011

Product	End Date
□ Red Hat Enterprise Linux 6	N/A

The following subscriptions will cover the products selected on 01/31/2011

Subscription Name	Available Subscriptions
Red Hat Enterprise Linux (1-2 Sockets) (Unlimited Virtualization)	25
Red Hat Network Satellite (3 year)	3
Red Hat and JBoss Developer Professional	8

Subscription Details

Subscription:

그림 14.21. Subscription Assistant

Alternatively, you can respond to entitlements notifications by managing subscriptions generally:

- The entitlements certificate can be updated or a new one can be added (14.13절. “Updating Entitlements Certificates”).
- The system can be subscribed to another subscription that contains the product (14.6절. “Handling Subscriptions”).

14.11. Changing the Healing Check Frequency

Subscription Manager can monitor all of the active entitlements for a system. Along with passively warning that a subscription is close to expiration (14.10절. “Responding to Subscription Notifications”), Subscription Manager can be configured to re-subscribe to subscriptions, automatically and actively, as one nears its expiry. This is system healing.

System healing prevents a system from having unentitled products as long as any valid subscription is available for it.

System healing is configured as part of the Subscription Manager daemon, rhsmcertd. This daemon checks the certificate validity dates daily. If a subscription is within 24 hours of expiring, then Subscription Manager will check for any available compatible subscriptions and automatically re-subscribes the system, much like auto-subscribing during registration.

참고

Healing cannot be disabled by changing the time interval. Setting the healFrequency parameter to zero means that Subscription Manager simply uses the default time setting.

1. Open the Subscription Manager configuration file:

```
# vim /etc/rhsm/rhsm.conf
```

2. In the [rhsmcertd] section, set the healFrequency parameter to the time, in minutes, to check for changed subscriptions.

```
[rhsmcertd]
certFrequency = 240
healFrequency = 1440
```

3. Restart the rhsmcertd daemon to reload the configuration.

```
# service rhsmcertd start
```

14.12. Working with Subscription Asset Manager

Subscription Asset Manager works with the local Subscription Manager tools, but the local Subscription Manager must be configured to work with the given Subscription Asset Manager service.

This section covers the procedures for setting up Subscription Manager to work with Subscription Asset Manager.

The Subscription Asset Manager documentation[3] details all the tasks for managing the infrastructure:

- Creating organizations and environments.
- Creating activation keys.
- Managing subscription manifests from Red Hat.
- Viewing notification and system reports.

14.12.1. Configuring Subscription Manager to Work with Subscription Asset Manager

Subscription Asset Manager performs two backend management functions:

- Allocate subscriptions as a subscription service
- Work as a real-time proxy for the content delivery network

That means that the local Subscription Manager client needs to be configured to use Subscription Asset Manager as its subscription service and content provider, rather than using the default Red Hat Network (hosted) configuration.

1. Obtain a copy of the CA certificate for the Subscription Asset Manager server and install it in the Subscription Manager certificate directory.

```
[root@server ~]# cd /etc/rhsm/ca
[root@server ca]# scp  sam.example.com:/etc/candlepin/certs/candlepin-ca.crt .
[root@server ca]# mv candlepin-ca.crt sam-local.pem
```

2. Update the settings in the rhsm.conf file. Several parameters need to be reset:

   - The configuration needs to point to the Subscription Asset Manager host for the subscription service host.

   ```
   hostname = sam.example.com
   ```

   - The prefix value, which sets the directory location for the subscription service, needs to be set to /sam/api.

   ```
   prefix= /sam/api
   ```

   - The content service hostname needs to be set to the Subscription Manager host, with the port set to 8088:

   ```
   baseurl= https://sam.example.com:8088
   ```

   - The CA certificate used for SSL connections needs to be set to the Subscription Asset Manager CA certificate, not Red Hat Network's global CA certificate.

   ```
   repo_ca_cert = %(ca_cert_dir)ssam_certificate.pem
   ```

[3] http://docs.redhat.com/docs/en-US/Red_Hat_Subscription_Asset_Manager/1.0/html/Installation_Guide/index.html

This can be done by editing the rhsm.conf file directly or by using the config command. For example:

```
[root@server1 ~]# subscription-manager config --server.hostname=sam.example.com --server.prefix=/sam/api --
rhsm.baseurl=https://sam.example.com:8088 --rhsm.repo_ca_cert=%(ca_cert_dir)ssam_certificate.pem
```

Changing the Subscription Manager configuration with the config command is covered in 14.14.2절. “Using the config Command” .

14.12.2. Viewing Organization Information

Infrastructures that have their own local content and subscription services, such as Subscription Asset Manager, can define groups that organize their systems. The primary division is organizations, which create independent units. The systems and users in one organization are invisible to the systems and users in another organization. Organizations can be subdivided into environments, which provide associations with content repositories and allowed products, versions, and content sets. A system can belong to multiple environments.

This is described in 14.3.1절. “Local Subscription Services, Local Content Providers, and Multi-Tenant Organizations” .

Organizations, environments, and repositories are created and configured in the distributor application, such as Subscription Asset Manager. However, the organization structure for a system or for a user account can be viewed using the Subscription Manager command-line tools. The orgs, environments, and repos commands list the organization, environment, and repository information for the system, depending on the organization and environments it belongs to.

The orgs lists the friendly name of the organization, such as Dev East, and then the key or ID for the organization which is used when registering consumers.

```
[root@server1 ~]# subscription-manager orgs --username=jsmith --password=secret
+-------------------------------------------+
          admin Organizations
+-------------------------------------------+

OrgName:          Admin Owner
OrgKey:          admin

OrgName:          Dev East
OrgKey:          deveast

OrgName:          Dev West
OrgKey:          devwest
```

The environments lists whatever environments are configured for the given organization which are assigned to that system. The organization may have other environments available, but they are only listed if the system belongs to them.

```
[root@server1 ~]# subscription-manager environments --username=jsmith --password=secret --org=admin
+-------------------------------------------+
          Environments
+-------------------------------------------+

Name:                      Locker
Description:               None

Name:                      Dev
Description:
```

```
Name:                          Prod
Description:
```

Distributor applications can defined a number of different content repositories, based on environments, physical locations, and other factors. Even when using the Red Hat content delivery network, multiple repositories are available, depending on the product. The repos command lists all of the repositories that are available to the configuration environments and organization for a system, and then shows whether those repositories are enabled for the system.

```
[root@server1 ~]# subscription-manager repos --list
+----------------------------------------------------------+
    Entitled Repositories in /etc/yum.repos.d/redhat.repo
+----------------------------------------------------------+

RepoName:                        never-enabled-content
RepoId:                         never-enabled-content
RepoUrl:                        https://content.example.com/repos/optional
Enabled:                        0


RepoName:                        always-enabled-content
RepoId:                         always-enabled-content
RepoUrl:                        https://content.example.com/repos/dev
Enabled:                        1


RepoName:                        content
RepoId:                         content-label
RepoUrl:                        https://content.example.com/repos/prod
Enabled:                        1
```

14.13. Updating Entitlements Certificates

An entitlement certificate represents a subscription that has been consumed by a given system. It includes all of the products which are included in the subscription for service and support, the subscription's start and end dates, and the number of entitlements included for each product. An entitlement certificate does not list products that are currently installed on the system; rather, it lists all products that are available to the system.

The entitlement certificate is an X.509 certificate and is stored in a base 64-encoded blob in a .pem file.

When a subscription expires or is changed, then the entitlement certificate must be updated to account for the changes. The Red Hat Subscription Manager polls the subscription service periodically to check for updated entitlement certificates; this can also be updated immediately or pulled down from the Customer Portal. The entitlement certificates are updated by revoking the previous entitlement certificate and generating a new one to replace it.

14.13.1. Updating Entitlement Certificates

1. Open the Red Hat Customer Portal.

   ```
   https://access.redhat.com/
   ```

2. Click the Subscriptions tab to open the subscriptions menu, and select the Consumers List option under Certificate-based Management.

3. Click the system name in the list of consumers.

4. Open the Applied Subscriptions tab for the list of all active, assigned subscriptions for the consumer.

5. Click the Download All Certificates button above the list of subscriptions. If there is only one subscription, then click the Download button by the certificate.

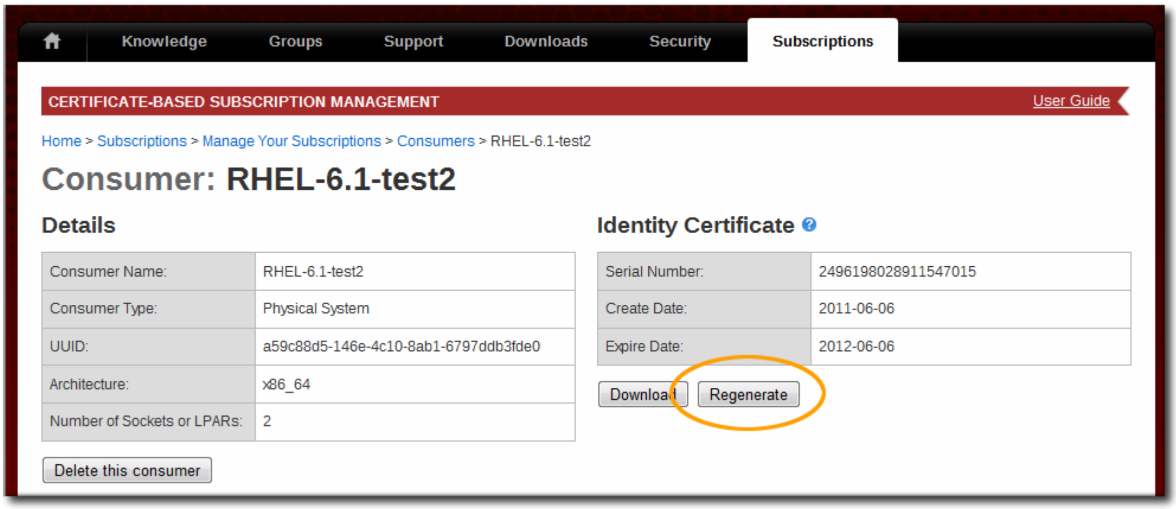

CERTIFICATE-BASED SUBSCRIPTION MANAGEMENT

Home > Subscriptions > Manage Your Subscriptions > Consumers > RHEL-6.1-test2

Consumer: RHEL-6.1-test2

Details

Consumer Name:	RHEL-6.1-test2
Consumer Type:	Physical System
UUID:	a59c88d5-146e-4c10-8ab1-6797ddb3fde0
Architecture:	x86_64
Number of Sockets or LPARs:	2

Identity Certificate

Serial Number:	2496198028911547015
Create Date:	2011-06-06
Expire Date:	2012-06-06

To retrieve an individual entitlement certificate, click the Download link in the subscription row.

6. If all entitlement certificates were downloaded in an archive file, then there are multiple archives in the downloaded certificates.zip file. Unzip the directories until the PEM files for the entitlement certificates are available.

7. Import the certificate PEM file. This can be done using the Import Certificates button in the Subscription Manager GUI or using the import command:

```
# subscription-manager import --certificate=/tmp/export/entitlement_certificates/596576341785244687.pem --certificate=/
tmp/export/entitlement_certificates/3195996649750311162.pem
Successfully imported certificate 596576341785244687.pem
Successfully imported certificate 3195996649750311162.pem
```

14.13.2. Updating Subscription Information

The refresh command updates all of the subscription information that is available to the consumer. This removes expired subscriptions and adds new subscriptions to the list. This does not subscribe the machine, but it does pull in the newest data for administrators to use.

```
[root@server1 ~]# subscription-manager refresh
```

14.14. Configuring the Subscription Service

By default, Red Hat Subscription Manager (both GUI and CLI) talk to the subscription service and the Customer Portal for their subscription services and content delivery, respectively. Red Hat Subscription Manager can be configured to use different content servers or subscription services. Other aspects of the Red Hat Subscription Manager — like the locations to look for system and product certificates or the system information used by Red Hat Subscription Manager to identify compatible entitlements — can also be customized to fit the network environment.

14.14.1. Red Hat Subscription Manager Configuration Files

The primary configuration file for Red Hat Subscription Manager, both the GUI and CLI tools, is the rhsm.conf configuration file. There are other support files that either influence the Red Hat Subscription Manager service or can help administrators better use the Subscription Manager.

14.14.1.1. All Files Used by Red Hat Subscription Manager

All of the files related to the configuration of Red Hat Subscription Manager are used by both the GUI and CLI; there is no separate configuration.

표 14.6. Red Hat Subscription Manager Files and Directories

File or Directory	Description
/etc/rhsm	The primary Red Hat Subscription Manager configuration directory.
/etc/rhsm/rhsm.conf	The Red Hat Subscription Manager configuration file. This is used by both the GUI and the CLI.
/etc/rhsm/facts	Any user-defined JSON files that override or add system facts to determine entitlement compatibility. Any facts files must end in .facts.
/var/lib/rhsm/cache/installed_products.json	A master list of installed products, which is sent by Subscription Manager to a hosted content service, such as Subscription Asset Manager.
/var/lib/rhsm/facts/facts.json	The default system facts filed, gathered by the Subscription Manager.
/var/lib/rhsm/packages/	The package profile cache (a list of installed products) which is gathered and periodically updated by the Subscription Manager.
/var/log/rhsm	The Red Hat Subscription Manager log directory.
/var/log/rhsm/rhsm.log	The log for the Red Hat Subscription Manager tools.
/var/log/rhsm/rhsmcertd.log	The log for the Red Hat Subscription Manager daemon, rhsmcertd.
/etc/pki/consumer	The directory which contains the identity certificates used by the system to identify itself to the subscription service.
/etc/pki/consumer/cert.pem	The base-64 consumer identity certificate file.
/etc/pki/consumer/key.pem	The base-64 consumer identity key file.
/etc/pki/entitlement	The directory which contains the entitlement certificates for the available subscriptions.
/etc/pki/product/product_serial#.pem	The product certificates for installed software products.
/var/run/subsys/rhsm	Runtime files for Red Hat Subscription Manager
/etc/init.d/rhsmcertd	The subscription certificate daemon.
/etc/cron.daily/rhsm-complianced and /usr/libexec/rhsm-complianced	Files to run daily checks and notifications for subscription validity.

File or Directory	Description
/etc/yum/pluginconf.d/rhsmplugin.conf	The configuration file to include the Red Hat Subscription Manager plug-in in the yum configuration.
/usr/share/rhsm	All of the Python and script files used by both Red Hat Subscription Manager tool to perform subscription tasks.
/usr/share/rhsm/gui	All of the Python script and image files used to render the Red Hat Subscription Manager GUI.

14.14.1.2. About the rhsm.conf File

The main configuration file for the Subscription Manager is rhsm.conf. This file configures several important aspects of how Red Hat Subscription Manager interacts with both entitlements and content services:

- The subscription service connection information, including the server host and port
- The content service to use, in the form of a web address
- The location of all of the different certificates used by the subscription service, including CA certificates for SSL authentication, identity certificates for the system, and entitlement and product certificates

The rhsm.conf file is divided into three sections. Two major sections define the subscription service ([server]) and content and product delivery ([rhsm]). The third section relates to the rhsmcertd daemon. Each assertion is a simple attribute= value pair. Any of the default values can be edited; all possible attributes are present and active in the default rhsm.conf file.

예 14.9. Default rhsm.conf File

```
# Red Hat Subscription Manager Configuration File:

# Unified Entitlement Platform Configuration
[server]
# Server hostname:
hostname = subscription.rhn.redhat.com

# Server prefix:
prefix = /subscription

# Server port:
port = 443

# Set to 1 to disable certificate validation:
insecure = 0

# Set the depth of certs which should be checked
# when validating a certificate
ssl_verify_depth = 3

# Server CA certificate location:
ca_cert_dir = /etc/rhsm/ca/

# an http proxy server to use
proxy_hostname =

# port for http proxy server
```

```
proxy_port =

# user name for authenticating to an http proxy, if needed
proxy_user =

# password for basic http proxy auth, if needed
proxy_password =

[rhsm]
# Content base URL:
baseurl= https://cdn.redhat.com

# Default CA cert to use when generating yum repo configs:
repo_ca_cert = %(ca_cert_dir)sredhat-uep.pem

# Where the certificates should be stored
productCertDir = /etc/pki/product
entitlementCertDir = /etc/pki/entitlement
consumerCertDir = /etc/pki/consumer

[rhsmcertd]
# Frequency of certificate refresh (in minutes):
certFrequency = 240
# Frequency of autoheal check (1440 min = 1 day):
healFrequency = 1440
```

표 14.7. rhsm.conf Parameters

Parameter	Description	Default Value
[server] Parameters		
hostname	Gives the IP address or fully-qualified domain name of the subscription service.	subscription.rhn.redhat.com
prefix	Gives the directory, in the URL, to use to connect to the subscription service.	/subscription
port	Gives the port to use to connect to the subscription service.	443
insecure	Sets whether to use a secure (0) or insecure (1) connection for connections between the Subscription Manager clients and the subscription service.	0
ssl_verify_depth	Sets how far back in the certificate chain to verify the certificate.	3
proxy_hostname	Gives the hostname of the proxy server. This is required.	
proxy_port	Gives the port of the proxy server. This is required.	
proxy_user	Gives the user account to use to access the proxy server. This may not be required,	

Parameter	Description	Default Value
	depending on the proxy server configuration.	
proxy_password	Gives the password credentials to access the proxy server. This may not be required, depending on the proxy server configuration.	
ca_cert_dir	Gives the location for the CA certificate for the CA which issued the subscription service's certificates. This allows the client to identify and trust the subscription service for authentication for establishing an SSL connection.	/etc/rhsm/ca
[rhsm] Parameters		
baseurl	Gives the full URL to access the content delivery system.	https://cdn.redhat.com
repo_ca_cert	Identifies the default CA certificate to use to set the yum repo configuration.	%(ca_cert_dir)sredhat-uep.pem
productCertDir	Sets the root directory where the product certificates are stored and can be accessed by Subscription Manager.	/etc/pki/product
consumerCertDir	Sets the directory where the identity certificate for the system is stored and can be accessed by Subscription Manager.	/etc/pki/consumer
entitlementCertDir	Sets the directory where the entitlement certificates for the system are stored and can be accessed by Subscription Manager. Each subscription has its own entitlement certificate.	/etc/pki/entitlement
[rhsmcertd] Parameters		
certFrequency	Sets the interval, in minutes, to check and update entitlement certificates used by Subscription Manager.	240
healFrequency	Sets the interval, in minutes, to check for change subscriptions and installed products and to allocate subscriptions, as necessary, to maintain subscription status for all products.	1440

Parameter	Description	Default Value
healFrequency	Sets the interval, in minutes, to check for change subscriptions and installed products and to allocate subscriptions, as necessary, to maintain subscription status for all products.	1440

14.14.2. Using the config Command

subscription-manager has a subcommand that can change the rhsm.conf configuration file. Almost all of the connection information used by Subscription Manager to access the subscription server, content server, and any proxies is set in the configuration file, as well as general configuration parameters like the frequency Subscription Manager checks for entitlements updates. There are major divisions in the rhsm.conf file, such as [server] which is used to configure the subscription server. When changing the Subscription Manager configuration, the settings are identified with the format section.parameter and then the new value. For example:

```
server.hostname=newsubscription.example.com
```

When changing the value for a parameter, the parameter is passed as an argument to the config command:

```
[root@server1 ~]# subscription-manager config --section.parameter=newValue
```

For example, to change the hostname of the subscription service:

```
[root@server1 ~]# subscription-manager config --server.hostname=subscription.example.com
```

All of the rhsm.conf file parameters are listed in 표 14.7. “rhsm.conf Parameters” . This is most commonly used to change connection settings:

- server.hostname (subscription server)
- server.proxy
- server.proxy_port
- server.proxy_user
- server.proxy_password
- rhsm.baseurl (content server)
- rhsm.certFrequency

The config command also has a --remove option. This deletes the current value for the parameter without supplying a new parameter. A blank value tells Subscription Manager to use any default values that are set for that parameter rather than a user-defined value. For example:

```
[root@server1 ~]# subscription-manager config --remove=rhsm.certFrequency

The default value for rhsm.certFrequency will now be used.
```

If a value does not have a default, then the command returns simply that the value has been removed:

```
[root@server1 ~]# subscription-manager config --remove=server.proxy

You have removed the value in section server for parameter proxy.
```

14.14.3. Using an HTTP Proxy

Some network environments may only allow external Internet access or access to content servers by going through an HTTP proxy.

14.14.3.1. Configuring an HTTP Proxy for GUI Use

Subscription Manager can be configured to use an HTTP proxy for all of its connections to the subscription service. (This is also an advanced configuration option at firstboot.) To configure the proxy:

1. Launch Subscription Manager. For example:

   ```
   [root@server ~]# subscription-manager-gui
   ```

2. Click the Proxy Configuration button at the top of the window in the Tools area.

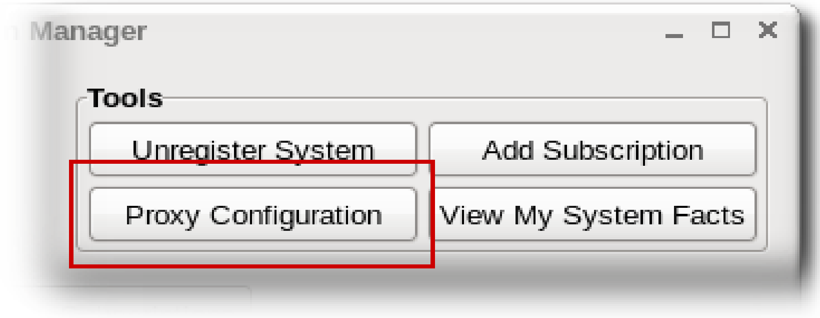

3. Check the ...Connect to Red Hat Network via an HTTP Proxy checkbox and enter the server location, in the format hostname:port.

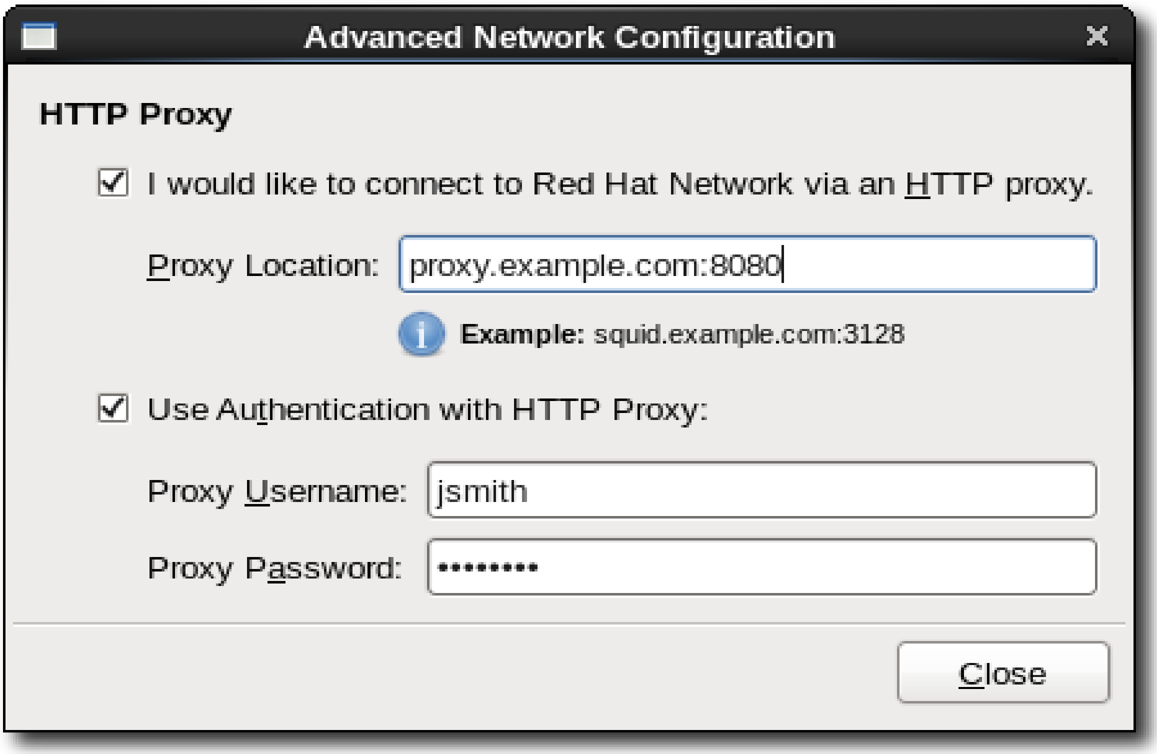

Advanced Network Configuration

HTTP Proxy

☑ I would like to connect to Red Hat Network via an HTTP proxy.

Proxy Location: proxy.example.com:8080

Example: squid.example.com:3128

☑ Use Authentication with HTTP Proxy:

Proxy Username: jsmith

Proxy Password: ••••••••

4. If the proxy requires a username/password to allow access, then also select the authentication checkbox and fill in the user credentials.

5. The configuration is automatically applied, so when the proxy is configured, simply close the window.

14.14.3.2. Configuring HTTP Proxy in the rhsm.conf File

The HTTP proxy settings can be configured in the rhsm.conf file; this is the same as configuring it in the Subscription Manager GUI. The proxy configuration is stored and used for every connection between the subscription service and the local system.

1. Open the Subscription Manager configuration file.

```
vim /etc/rhsm/rhsm.conf
```

2. Change the settings in the [server] section that relate to the HTTP proxy. All parameters are described in 표 14.7. "rhsm.conf Parameters" . There are four parameters directly related to the proxy:

   - proxy_hostname for the IP address or fully-qualified domain name of the proxy server; this is required.

Note

Leaving the proxy_hostname argument blank means that no HTTP proxy is used.

- proxy_port for the proxy server port.
- proxy_user for the user account to connect to the proxy; this may not be required, depending on the proxy server's configuration.
- proxy_password for the password for the user account to connect to the proxy; this may not be required, depending on the proxy server's configuration.

```
[server]

# an http proxy server to use
proxy_hostname = proxy.example.com

# port for http proxy server
proxy_port = 443

# user name for authenticating to an http proxy, if needed
proxy_user =

# password for basic http proxy auth, if needed
proxy_password =
```

14.14.3.3. Passing HTTP Proxy Information with subscription-manager Commands

Rather than using a permanently-configured HTTP proxy, as the GUI does, HTTP proxy information can be passed with a command invocations. The arguments listed in 표 14.8. "Proxy Arguments" are available to every command used with subscription-manager.

표 14.8. Proxy Arguments

Argument	Description	Required for a Proxy Connection?
--proxy	Gives the proxy server to connect to, in the format hostname:port.	Yes
--proxyuser	Gives the username to use to authenticate. This is only required if user authentication is required.	No
--proxypassword	Gives the password to use with the user account. This is only required if user authentication is required.	No

The proxy information can be passed with any subscription-manager operation. For example:

```
[root@server1 ~]# subscription-manager subscribe --pool=ff8080812bc382e3012bc3845ca000cb --proxy=proxy.example.com:8443
 --proxyuser=jsmith --proxypassword=secret
```

14.14.4. Changing the Subscription Server

The Subscription Manager usually connects to the subscription service, and the public server is configured in the rhsm.conf file. There may be instances when an organization is running a mirror or an internal subscription service; in those situations, the connection settings can be altered to connect to the specific servers. The subscription service connection settings are in the [server] section of the configuration file.

1. Open the Subscription Manager configuration file.

   ```
   vim /etc/rhsm/rhsm.conf
   ```

2. Change the settings in the [server] section that relate to the subscription service connection. All parameters are described in 표 14.7. “rhsm.conf Parameters” . There are three parameters directly related to the connection:

   - hostname for the IP address or fully-qualified domain name of the machine

   - prefix for the subscription service directory

   - port for the subscription service port

   ```
   [server]
   hostname=entitlements.server.example.com
   prefix=/candlepin
   port=8443
   ```

14.14.5. Configuring Red Hat Subscription Manager to Use a Local Content Provider

By default, the Subscription Manager is configured to use Red Hat's content delivery service, which is available at https://cdn.redhat.com. This can be changed to use a different external content delivery system or to use an organization-managed content system, such as Subscription Asset Manager.

1. Open the Subscription Manager configuration file.

   ```
   vim /etc/rhsm/rhsm.conf
   ```

2. Change the baseurl directive in the [rhsm] section. This is the full URL to the service.

   ```
   [rhsm]
   # Content base URL:
   baseurl= http://content.example.com/content
   ```

14.14.6. Managing Secure Connections to the Subscription Server

Red Hat Subscription Manager assumes, by default, that the subscription clients connect to the subscription service using a secure (SSL) connection. This requires that the CA certificate of the subscription service be downloaded and available locally for the client and that the appropriate connections be configured.

1. Open the Subscription Manager configuration file.

```
vim /etc/rhsm/rhsm.conf
```

2. Change the settings in the [server] section that relate to a secure connection. All parameters are described in 표 14.7. “rhsm.conf Parameters” . There are three parameters directly related to the connection:

   - insecure to set whether to use a secure (0) or insecure (1) connection

   - ca_cert_dir for the directory location for the CA certificate for authentication and verification

   - port for the subscription service port; this should be an SSL port if a secure connection is required

```
[server]
port=8443
insecure = 1
ca_cert_dir = /etc/rhsm/ca
```

   There is also an optional parameter to set how far in a certificate chain to go to validate a certificate. By default, this is three, meaning the server validates three CAs back in the issuing chain.

```
ssl_verify_depth = 3
```

14.14.7. Starting and Stopping the Subscription Service

The Red Hat Subscription Manager daemon, rhsmcertd, runs as a service on the system. The daemon, by default, starts with the system, and it can be started, stopped, or checked with the service command.

```
service rhsmcertd status
rhsmcertd (pid 13084) is running...
```

Red Hat Enterprise Linux has a tool called chkconfig which manages the automatic startup and shutdown settings for each process on the server, described in 17.5절. “chkconfig” . When a system reboots, some services can be automatically restarted. chkconfig also defines startup settings for different run levels of the server.

The Red Hat Subscription Manager service, which runs routinely to check for changes in the entitlements for an organization, can be controlled by chkconfig. By default, the Red Hat Subscription Manager daemon, rhsmcertd, is configured to run at levels 3, 4, and 5, so that the service is started automatically when the server reboots.

The run level settings can be reset using chkconfig. For example, to enable run level 2:

```
chkconfig --level 2345 rhsmcertd on
```

To remove the rhsmcertd from the start list, change the run level settings off:

```
chkconfig --level 2345 rhsmcertd off
```

Red Hat Enterprise Linux also has a GUI console that can manage the service and chkconfig settings.

1. In the main menu, select the Administration link and open the Server Settings submenu.

2. Open the Services link.

Note

The system-config-services package must be installed for the Services wizard to be available.

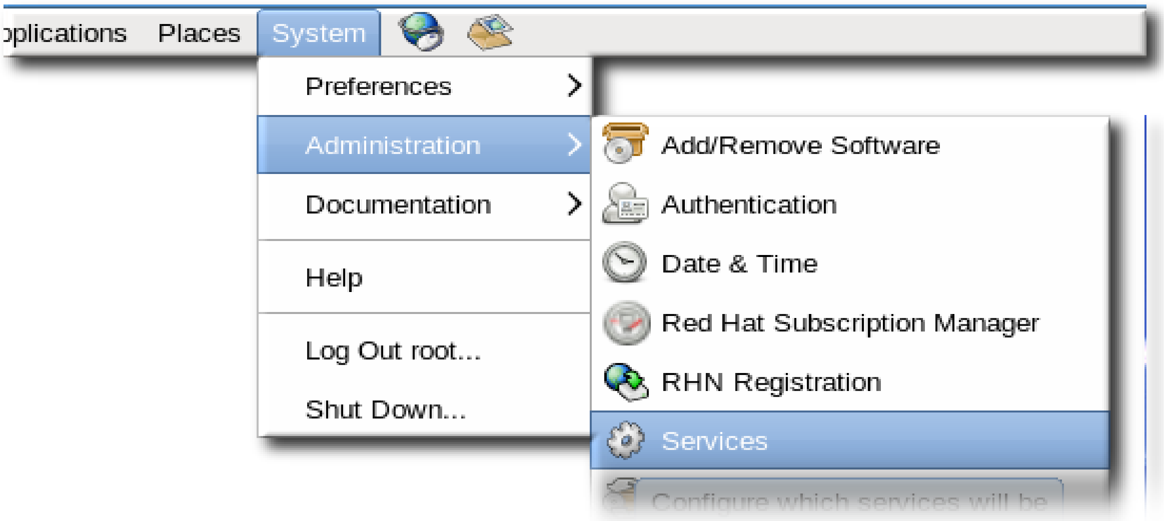

3. Scroll to the rhsmcertd item in the list of services on the left, and then edit the service as desired.

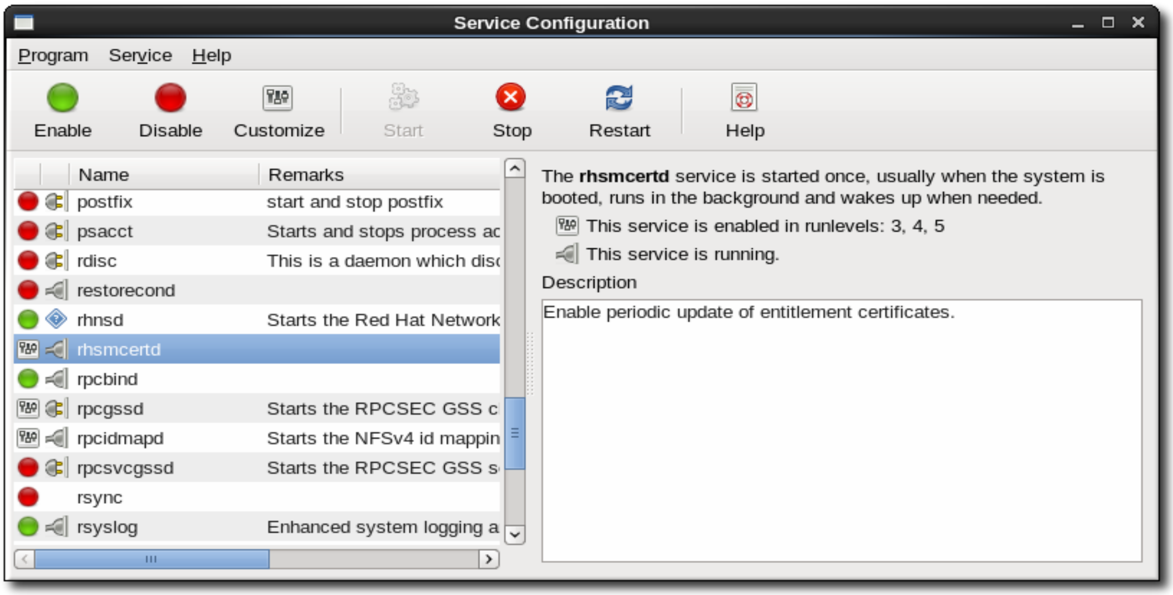

Name	Remarks
postfix	start and stop postfix
psacct	Starts and stops process ac
rdisc	This is a daemon which disc
restorecond	
rhnsd	Starts the Red Hat Network
rhsmcertd	
rpcbind	
rpcgssd	Starts the RPCSEC GSS c
rpcidmapd	Starts the NFSv4 id mappin
rpcsvcgssd	Starts the RPCSEC GSS s
rsync	
rsyslog	Enhanced system logging a

The **rhsmcertd** service is started once, usually when the system is booted, runs in the background and wakes up when needed.

This service is enabled in runlevels: 3, 4, 5

This service is running.

Description

Enable periodic update of entitlement certificates.

14.14.8. Checking Logs

There are two log files maintained for Red Hat Subscription Manager in the /var/log/rhsm directory:

- rhsm.log shows every invocation and result of running the Subscription Manager GUI or CLI
- rhsmcertd.log shows every time a new certificate is generated, which happens on a schedule defined by the certFrequency parameter in the rhsm.conf file.

The rhsm.log file contains the sequence of every Python call for every operation invoked through the Subscription Manager tools. Each entry has this format:

```
YYYY-MM-DD HH:MM:SS,process_id [MESSAGE_TYPE] call python_script response
```

The response in the log entry can be very complex, spanning multiple lines, or relatively simply, with just a status code.

Because each log entry in rhsm.log relates to the Python script or function that was called, there can be multiple log entries for a single operation.

예 14.10. rhsm.log Entry

```
2010-10-01 17:27:57,874 [INFO] _request() @connection.py:97 - status code: 200
2010-10-01 17:27:57,875 [INFO] perform() @certlib.py:132 - updated:
Total updates: 0
Found (local) serial# []
Expected (UEP) serial# []
Added (new)
  <NONE>
Deleted (rogue):
  <NONE>
Expired (not deleted):
  <NONE>
Expired (deleted):
  <NONE>
2010-10-01 17:27:57,878 [INFO] __init__() @connection.py:193 - Using certificate authentication: key = /etc/pki/consumer/
key.pem, cert = /etc/pki/consumer/cert.pem, ca = /etc/pki/CA/candlepin.pem, insecure = True
```

```
2010-10-01 17:27:57,878 [INFO] __init__() @connection.py:196 - Connection Established: host: candlepin.example.com,
 port: 443, handler: /candlepin
```

The entries in the rhsmcertd.log file are much simpler. The log only records when the rhsmcertd daemon starts or stops and every time a certificate is updated.

예 14.11. rhsmcertd.log Entry

```
Fri Oct  1 13:27:44 2010: started: interval = 240 minutes
Fri Oct  1 13:27:50 2010: certificates updated
```

14.14.9. Checking and Adding System Facts

Entitlements are available to a system based on whether the software is compatible with the system's architecture. For example, there are different products and subscriptions for 32-bit and 64-bit platforms. Red Hat Subscription Manager determines compatibility by collecting a range of facts about the system's hardware and architecture and then comparing it with all available entitlements.

The collected facts can be viewed, updated to acknowledge a hardware or configuration change, or overridden to force compatibility in the specified areas.

The system facts are very similar to the information in /etc/redhat-release or /etc/sysconfig. In both the Red Hat Subscription Manager GUI and CLI, the facts are represented as simple attribute: value pairs.

> **Tip**
>
> Updating the facts resends the information about the system to the Red Hat subscription service so that it can update the list of subscriptions which match the system architecture. Updating the facts is a very good thing to do after hardware upgrades or other important system changes.

14.14.9.1. Checking Facts from the Red Hat Subscription Manager UI

1. Launch Subscription Manager. For example:

   ```
   [root@server ~]# subscription-manager-gui
   ```

2. In the Tools at the top of the window, click the View System Facts button.

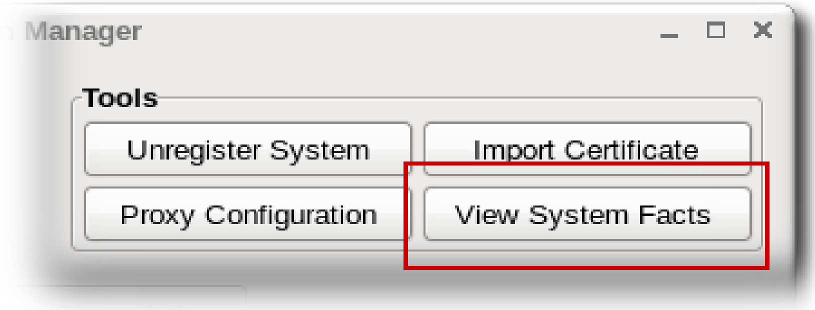

3. All of the current facts for the system are listed in the table, broken down into categories. Each category is in a closed list; to reveal all of the facts in that category, click the arrow by the category name.

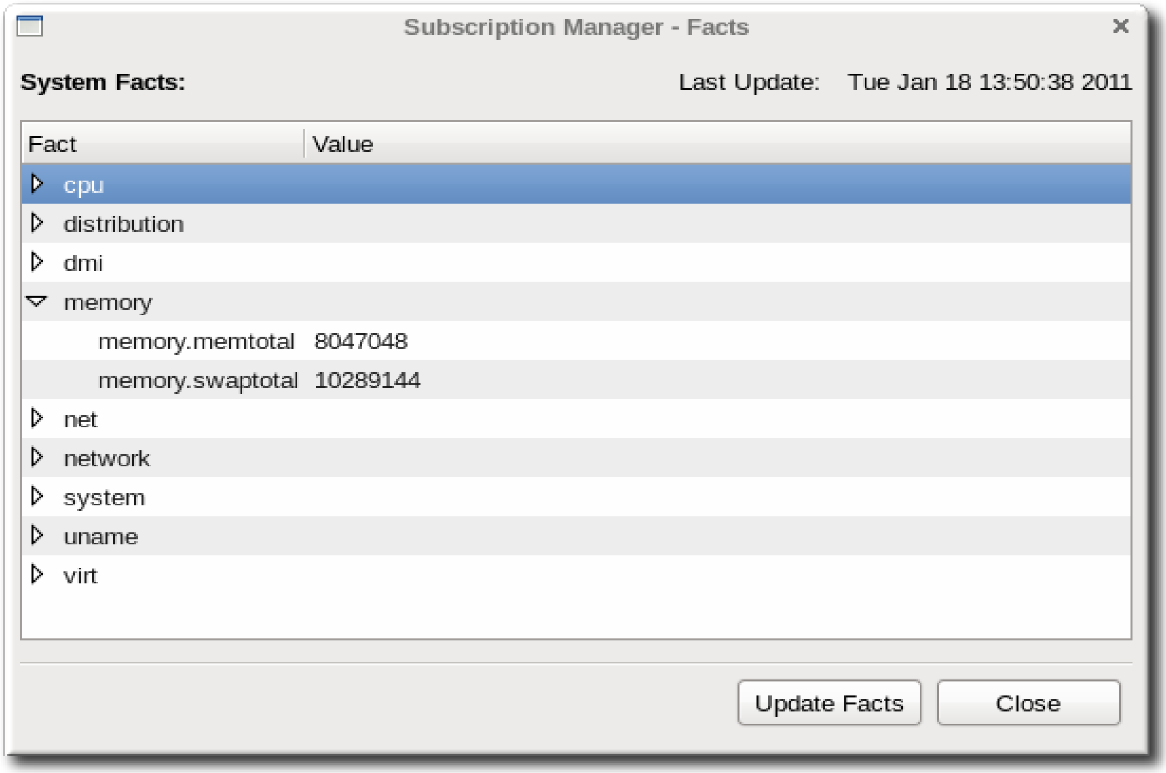

System Facts: Last Update: Tue Jan 18 13:50:38 2011

Fact	Value
▷ cpu	
▷ distribution	
▷ dmi	
▽ memory	
memory.memtotal	8047048
memory.swaptotal	10289144
▷ net	
▷ network	
▷ system	
▷ uname	
▷ virt	

To update the facts, click the Update Facts button in the bottom right of the window.

14.14.9.2. Checking Facts with subscription-manager

To simply list the facts, run the facts command with the --list option.

```
[root@server1 ~]# subscription-manager facts --list

cpu.architecture: i686
cpu.core(s)_per_socket: 4
cpu.cpu(s): 4
cpu.cpu_family: 6
cpu.cpu_mhz: 2000.010
cpu.cpu_op-mode(s): 32-bit, 64-bit
cpu.cpu_socket(s): 1
cpu.l1d_cache: 32K
cpu.l1i_cache: 32K
cpu.l2_cache: 6144K
cpu.model: 23
cpu.stepping: 6
cpu.thread(s)_per_core: 1
cpu.vendor_id: GenuineIntel
cpu.virtualization: VT-x
distribution.id: Santiago
distribution.name: Red Hat Enterprise Linux Workstation
distribution.version: 5
dmi.baseboard.manufacturer: IBM
dmi.baseboard.product_name: Server Blade
... [snip] ...
```

To update the facts after a system change, use the --update option with the facts command.

```
[root@server1 ~]# subscription-manager facts --update
```

14.14.9.3. Overriding the Default System Facts

The system facts, as collected, are stored in /var/lib/rhsm/facts/facts.json. These facts are stored as attribute: value pairs, in a comma-separated list.

```
{"fact1": "value1","fact2": "value2"}
```

The primary file is generated and maintained by the Subscription Manager service. However, these values can be overridden to force architecture or platform compatibility (and thereby widening the available compatible subscriptions) by creating additional JSON facts files and dropping them in the /etc/rhsm/facts directory. These JSON files can override existing facts or even add new facts to be used by the subscription service.

예 14.12. Example Facts Override File

```
vim /etc/rhsm/facts/my-example.facts

{"uname.machine": "x86","kernel_version": "2.6.32","physical_location": "MTV colo rack 5"}
```

14.14.10. Regenerating Identity Certificates

To regenerate the consumer's identity certificate (meaning it is revoked and replaced), use the identity command.

Although credentials are not normally required with the identity command, using the --force option will require the username and password and will cause the Subscription Manager to prompt for the credentials if they are not passed in the command. This can be helpful if the identity certificate needs to be regenerated using a different Red Hat account than the original registration.

```
[root@server1 ~]# subscription-manager identity --regenerate --force
Username: jsmith@example.com
Password:
Identity certificate has been regenerated.
```

14.14.11. Getting the System UUID

The consumer or system UUID is a unique identifier used in the inventory subscription service. This UUID can be used to re-register the system if there is some kind of corruption or for internal tracking. In the GUI (14.14.9.1절. "Checking Facts from the Red Hat Subscription Manager UI"), this is listed as one of the system facts, under the system category:

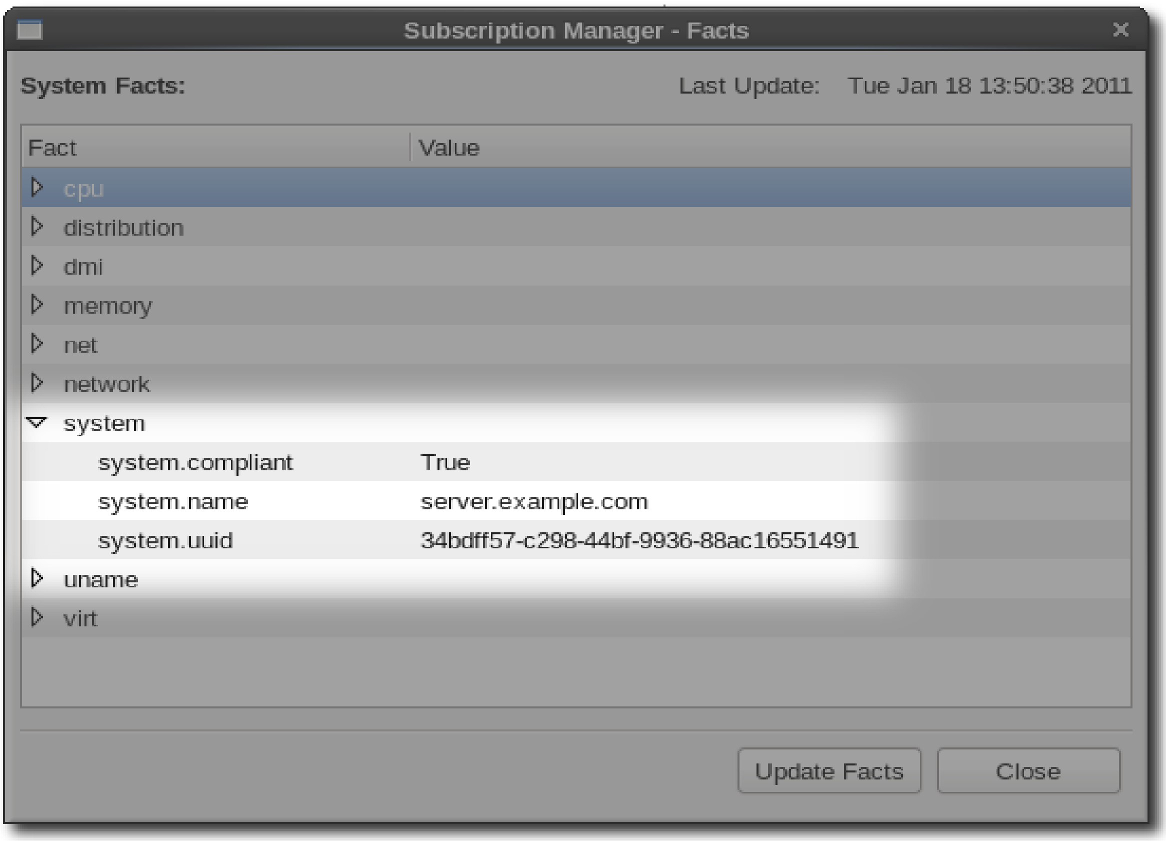

System Facts: Last Update: Tue Jan 18 13:50:38 2011

Fact	Value
▷ cpu	
▷ distribution	
▷ dmi	
▷ memory	
▷ net	
▷ network	
▽ system	
system.compliant	True
system.name	server.example.com
system.uuid	34bdff57-c298-44bf-9936-88ac16551491
▷ uname	
▷ virt	

From the command-line, use the identity command to return the current UUID. The UUID is the Current identity is value.

```
[root@server1 ~]# subscription-manager identity
Current identity is: 63701087-f625-4519-8ab2-633bb50cb261
name: server1.example.com
org name: 6340056
org id: 8a85f981302cbaf201302d89931e059a
```

14.14.12. Viewing Package Profiles

A package profile is the list of installed packages on a system (regardless of its subscription status). Red Hat Subscription Manager maintains a local list of installed packages to track the subscription status of the system. The package profile contains some general information about each package in the list:

- Package name
- Package version
- Epoch

- Publisher

This package manifest is always visible locally in the My Installed Software tab of the UI or by using the list --installed command with the command-line tools.

The Subscription Manager daemon, rhsmcertd, checks the system periodically — once when it is first registered and then when it runs a refresh operation every four hours — to get the most current list of installed products. When the system is registered and then whenever there is a change to the package list, Subscription Manager sends an updated package profile to the subscription service.

The package profile is stored in a cache file in /var/lib/rhsm/packages/.

Having an updated package profile for a system helps the subscription service identify compatible subscriptions.

14.14.13. Retrieving the Consumer ID, Registration Tokens, and Other Information

Some pieces of information are used frequently when managing entitlements using the subscription-manager script. Information like the consumer ID or subscription pool ID is pulled up and referenced automatically in the Red Hat Subscription Manager UI, but it has to be entered manually in the command line.

표 14.9. “Locations and Descriptions of Entitlement Data” lists common information that is used to manage subscriptions, the operations they are used in, and the places to find the data.

표 14.9. Locations and Descriptions of Entitlement Data

Information	Description	Operations Used In	Find It In ...

Consumer ID	A unique identifier for each system that is registered to the subscription service.	identity	The simplest method is to use the identity command to return the current UUID.
			`[root@server1 ~]# subscription-manager identity`
			`Current identity is: 63701087-f625-4519-8ab2-633bb50cb261`
			`name: consumer-1.example.com`
			`org name: 6340056`
			`org id: 8a85f981302cbaf201302d89931e059a`
			The Subject CN element of the identity certificate for the system, /etc/pki/consumer/cert.pem. The UUID can also be returned by using openssl to pretty-print the certificate.

Information	Description	Operations Used In	Find It In ...
			openssl x509 -text -in /etc/pki/consumer/cert.pem Certificate: ... snip ... Subject: CN=7d133d55 876f 4f47 83eb 0ee931cb0a97
Pool ID	An identifier for a specific set of subscriptions. This set is created when subscriptions are purchased. Whenever a system needs to subscribe to a product, it references a pool ID to identify which purchased set of subscriptions to use.	subscribe	The PoolID value given for a product when listing available subscriptions. For example: [root@server1 ~]# subscription-manager list --available +----------------------+ Available Subscriptions +----------------------+ ProductName: Red Hat Enterprise Linux, Standard (up to 2 sockets) 3 year ProductId: MCT0346F3 PoolId: ff8080812bc382e3012bc3845ca000cb Quantity: 2 Expires: 2011-02-28
Product certificate serial number	The identification used for a specific, installed product. A certificate with a unique serial number is generated when a product is installed; this serial number is used to identify that specific product installation when managing subscriptions.	unsubscribe	The SerialNumber line in the product subscription information. This can be returned by running list --consumed. [root@server1 ~]# subscription-manager list --consumed +-----------------------------+ Consumed Product Subscriptions +-----------------------------+ ProductName: High availability (cluster suite) ContractNumber: 0 SerialNumber: 11287514358600162 ....
Product ID	The internal identifier used to identify a type of product.		The ProductID value given for a product when listing available

Information	Description	Operations Used In	Find It In ...
			subscriptions. For example: [root@server1 ~]# subscription-manager list --available +---+ Available Subscriptions +---+ ProductName: RHEL for Physical Servers ProductId: MKT-rhel-server ... snip ...

14.15. About Certificates and Managing Entitlements

Part of managing subscriptions requires verifying the identity of everything involved, such as the system, the subscription service, and the available products. The subscription service uses X.509 certificates to handle the identity and authentication aspects of the subscription service. These X.509 certificates also contain the actual data about available subscriptions and installed products.

The first time a system is subscribed to a subscription, it downloads a certificate from the subscription service. The entitlement certificate contains all of the information about products that are available through that subscription. The entitlement certificate is revoked and reissued any time there is a change in the subscriptions for an organization. Once a product is actually installed on a machine, then another certificate is issued to manage the entitlements for the product on the system.

Each certificate issued and used by the Subscription Manager services is a .pem formatted file. This file format stores both keys and certificates in a base-64 blob. For example:

```
-----BEGIN CERTIFICATE-----
MIIDaTCCAtKgAwIBAgICBZYwDQYJKoZIhvcNAQEFBQAwSzEqMCgGA1UEAxMhY2Fu
ZGxlcGluMS5kZXZsYWIucGh4MS5yZWRoYXQuY29tMQswCQYDVQQGEwJVUzEQMA4G
A1UEBxMHUmFsZWlnaDAeFw0xMDEwMDYxNjMyMDVaFw0xMTEwMDYyMzU5NTlaMC8x
LTArBgNVBAMMJDQ4ODFiZDJmLTg2OGItNDM4Yy1hZjk2LThiMWQyODNkYWZmYzCC
ASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAKNyLw6+IMtjY03F7Otxj2GL
GTz5VKx1kfWY7q4OD4w+XlBHTkt+2tQV9S+4TFkUZ7XoI80LDL/BONpy/gq5c5cw
yKvjv2gjSS/pihgYNXc5zUOIfSj1vb3fHGHOkzdCcZMyWq1z0N/zaLClp/zP/pcM
og4NTAg2niNPjFYvkQ+oIl16WmQpefM0y0SY7N7oJd2T8dZjOiuLV2cVZLfwjrwG
9UpkT2J03g+n1ZA9q95ibLD5NVOdTy9+2lfRhdDViZaVoFiQXvg86qBHQ0ieENuF
a6bCvGgpTxcBuVXmsnl2+9dnMiwoDqPZp1HB6G2uNmyNe/IvkTOPFJ/ZVbtBTYUC
AwEAAaOB8zCB8DARBglghkgBhvhCAQEEBAMCBaAwCwYDVR0PBAQDAgSwMHsGA1Ud
IwR0MHKAFGiY1N2UtulxcMFy0j6gQGLTyo6CoU+kTTBLMSowKAYDVQQDEyFjYW5k
bGVwaW4xLmRldmxhYi5waHgxLnJlZGhhdC5jb20xCzAJBgNVBAYTAlVTMRAwDgYD
VQQHEwdSYWxlaWdoggkA1s54sVacN0EwHQYDVR0OBBYEFGbB5fqOzh32g4Wqrwhc
/96IupIgMBMGA1UdJQQMMAoGCCsGAQUFBwMCMB0GA1UdEQQWMBSkEjAQMQ4wDAYD
VQQDDAV4ZW9wczANBgkqhkiG9w0BAQUFAAOBgQANxHRsev4fYfnHO9kYcHo4UeK7
owN+fq92gl76iRHRnhzkPlhWL+uV2tyqGG9zJASOX+qEDOqN5sVAB4iNQTDGiUbK
z757igD2hsQ4ewv9Vq3QtnajWnfdaUZH919GgWs09Etg6ucsKwgfx1fqjSRLBbOo
lZuvBTYROOX6W2vKXw==
-----END CERTIFICATE-----
```

Tools like openssl or pk12util can be used to extract and view information from these certificates, in a pretty-print format. The product- and subscription-related information is extracted and viewable in the Red Hat Subscription Manager GUI or command-line tools.

This section describes the different certificates used by the subscription service and the entitlement information contained in those certificates. A much more detailed description of X.509 certificates and a public key infrastructure (PKI) is given in the Red Hat Certificate System documentation in chapter 1, "Introduction to Public-Key Cryptography,"[4] in the Red Hat Certificate System Deployment Guide.

표 14.10. Types of Certificates Used for Content and Entitlements

Certificate Type	Description	Default Location
Consumer Identity Certificate	Used to identify the system (consumer) to the subscription service. This contains a unique ID which is assigned to the system when it is registered to the system. The identity certificate itself is generated by the subscription service when the system is registered and then sent to the consumer.	/etc/pki/consumer
Entitlement Certificate	Contains a list of products that are available to a system to install, based on the subscriptions that the system has been subscribed to. The entitlement certificate defines the software products, the content delivery location, and validity dates. The presence of an entitlement certificate means that the system has consumed one of the quantities from the subscription.	/etc/pki/entitlement
Product Certificate	Contains the information about a product after it has been installed.	/etc/pki/ product/product_serial#.pem
CA Certificate	A certificate for the certificate authority which issued the SSL server certificate used by the subscription service. This must be installed on a system for the system to use SSL to connect to the subscription service.	/etc/rhsm/ca/candlepin-ca.pem
Satellite Certificate	An XML-formatted certificate which contains a product list. This is used by local Satellite 5.x systems, not the newer subscription service.	

[4] http://docs.redhat.com/docs/en-US/Red_Hat_Certificate_System/8.0/html/Deployment_Guide/Introduction_to_Public_Key_Cryptography.html

14.15.1. The Structure of Identity Certificates

An identity certificate is a standard SSL client certificate. This certificate is issued by the subscription service when the system registers to it. The system consumer subsequently uses this certificate to authenticate to the subscription service whenever it contacts the service after registration.

The certificate contains three important pieces of information:

- The consumer UUID, in the subject CN of the certificate
- The subscription service which the system is registered to, in the issuer field of the certificate
- The user account which registered the system, as the DirName value in the Subject Alt Name

The validity period of this certificate is associated with the time when the system was registered, not to any subscription contract periods or user account settings.

예 14.13. Identity Certificate

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 1430 (0x596)
        Signature Algorithm: sha1WithRSAEncryption
        Issuer: CN=entitlement.server.example.com, C=US, L=Raleigh
        Validity
            Not Before: Oct  6 16:32:05 2010 GMT
            Not After : Oct  6 23:59:59 2011 GMT
        Subject: CN=4881bd2f-868b-438c-af96-8b1d283daffc
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:a3:72:2f:0e:be:20:cb:63:63:4d:c5:ec:eb:71:
                    8f:61:8b:19:3c:f9:54:ac:75:91:f5:98:ee:ae:0e:
                    0f:8c:3e:5e:50:47:4e:4b:7e:da:d4:15:f5:2f:b8:
                    4c:59:14:67:b5:e8:23:cd:0b:0c:bf:c1:38:da:72:
                    fe:0a:b9:73:97:30:c8:ab:e3:bf:68:23:49:2f:e9:
                    8a:18:18:35:77:39:cd:43:88:7d:28:f5:bd:bd:df:
                    1c:61:ce:93:37:42:71:93:32:5a:ad:73:d0:df:f3:
                    68:b0:a5:a7:fc:cf:fe:97:0c:a2:0e:0d:4c:08:36:
                    9e:23:4f:8c:56:2f:91:0f:a8:22:5d:7a:5a:64:29:
                    79:f3:34:cb:44:98:ec:de:e8:25:dd:93:f1:d6:63:
                    3a:2b:8b:57:67:15:64:b7:f0:8e:bc:06:f5:4a:64:
                    4f:62:74:de:0f:a7:d5:90:3d:ab:de:62:6c:b0:f9:
                    35:53:9d:4f:2f:7e:da:57:d1:85:d0:d5:89:96:95:
                    a0:58:90:5e:f8:3c:ea:a0:47:43:48:9e:10:db:85:
                    6b:a6:c2:bc:68:29:4f:17:01:b9:55:e6:b2:79:76:
                    fb:d7:67:32:2c:28:0e:a3:d9:a7:51:c1:e8:6d:ae:
                    36:6c:8d:7b:f2:2f:91:33:8f:14:9f:d9:55:bb:41:
                    4d:85
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            Netscape Cert Type:
                SSL Client, S/MIME
            X509v3 Key Usage:
                Digital Signature, Key Encipherment, Data Encipherment
            X509v3 Authority Key Identifier:
                keyid:68:98:D4:DD:94:B6:E9:71:70:C1:72:D2:3E:A0:40:62:D3:CA:8E:82
                DirName:/CN=entitlement.server.example.com/C=US/L=Raleigh
                serial:D6:CE:78:B1:56:9C:37:41

            X509v3 Subject Key Identifier:
                66:C1:E5:FA:8E:CE:1D:F6:83:85:AA:AF:08:5C:FF:DE:88:BA:92:20
            X509v3 Extended Key Usage:
                TLS Web Client Authentication
            X509v3 Subject Alternative Name:
                DirName:/CN=admin-example
    Signature Algorithm: sha1WithRSAEncryption
        0d:c4:74:6c:7a:fe:1f:61:f9:c7:3b:d9:18:70:7a:38:51:e2:
        bb:a3:03:7e:7e:af:76:82:5e:fa:89:11:d1:9e:1c:e4:3e:58:
        56:2f:eb:95:da:dc:aa:18:6f:73:24:04:8e:5f:ea:84:0c:ea:
        8d:e6:c5:40:07:88:8d:41:30:c6:89:46:ca:cf:be:7b:8a:00:
        f6:86:c4:38:7b:0b:fd:56:ad:d0:b6:76:a3:5a:77:dd:69:46:
        47:f7:5f:46:81:6b:34:f4:4b:60:ea:e7:2c:2b:08:1f:c7:57:
        ea:8d:24:4b:05:b3:a8:95:9b:af:05:36:11:38:e5:fa:5b:6b:
        ca:5f
```

14.15.2. The Structure of Entitlement Certificates

An entitlement is analogous to an assigned software license. Entitlement certificates contain a list of available products for a system — software that the system has been granted rights to download and update. When a system is subscribed to a subscription pool, the system pulls down the entitlement certificate from the subscription service, which contains all of the information about available products.

An entitlement certificate contains a list of every potential product from every potential content source. The structure of the entitlement certificate, then, allows multiple namespaces for products, content servers, roles, orders, and systems. An entitlement certificate also contains complete information about the subscribed pool, even for products which may not be compatible with the specific system. In an entitlement certificate, the architecture and version definitions contain all of the allowed architectures and versions.

Note

The local Subscription Manager polls the subscription service routinely (every four hours by default) to check for changes in the entitlements. When a subscription is changed in some way, then the original entitlement certificate is revoked and is replaced with a new entitlement certificate.

The entitlement certificate is a *.pem file stored in the entitlement certificates directory, /etc/pki/entitlement. The name of the *.pem file is a numeric identifier that is generated by the subscription service. This ID is an inventory number that is used to associate a subscription quantity with the system in the software inventory.

The heading of the certificate contains the name of the subscription service which issued it, the validity period of the certificate (which is tied to the installation date of the product), and then the serial number of the installation of the product.

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            3c:da:6c:06:90:7f:ff
        Signature Algorithm: sha1WithRSAEncryption
        Issuer: CN=candlepin.example.com, C=US, L=City
        Validity
            Not Before: Oct  8 17:55:28 2010 GMT
            Not After : Oct  2 23:59:59 2011 GMT
        Subject: CN=8a878c912b875189012b8cfbc3f2264a
... [snip] ...
```

The key definition of the product is given in custom certificate extensions that are appended to the certificate. Each namespace defines certain information about a product, including its name, content servers which can deliver it, the format of delivery, and a GPG key to identify the release. Every individual entry is identified by a numeric object identifier (OID) with the same basic format:

```
1.3.6.1.4.1.2312.9.2.product_#.config_#:
   ..config_value
```

The 2 indicates that it is a product entry. product_# is a unique ID which identifies the specific product or variant. config_# relates to the installation information for that product, like its content server or the quantity available.

Note

Every entitlements-related extension begins with the OID base 1.3.6.1.4.1.2312.9. The subsequent numbers identify different subscription areas:

- .2. is the product-specific information
- .1. is the subscription information
- .4. contains the contract information, like its ID number and start and end dates
- .5. contains the consumer information, like the consumer ID which installed a product

A product definition contains a series of entries which configure all of the information required to identify and install the product. Each type of information has its own ID, the config_# in the OID, that is used consistently for all products. An example product is listed in 예 14.14. "Annotated Red Hat Enterprise Linux High Availability Product Extensions in an Entitlement Certificate" .

예 14.14. Annotated Red Hat Enterprise Linux High Availability Product Extensions in an Entitlement Certificate

```
             content repository type
             1.3.6.1.4.1.2312.9.2.30393.1:
                 ..yum
             product
             1.3.6.1.4.1.2312.9.2.30393.1.1:
                 .HRed Hat Enterprise Linux High Availability (for RHEL Entitlement) (RPMs)
             channel name
             1.3.6.1.4.1.2312.9.2.30393.1.2:
                 .Dred-hat-enterprise-linux-high-availability-for-rhel-entitlement-rpms
             vendor
             1.3.6.1.4.1.2312.9.2.30393.1.5:
                 ..Red Hat
             download URL
             1.3.6.1.4.1.2312.9.2.30393.1.6:
                 .Q/content/dist/rhel/entitlement/releases/$releasever/$basearch/highavailability/os
             key download URL
             1.3.6.1.4.1.2312.9.2.30393.1.7:
                 .2file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
```

```
            flex quantity
            1.3.6.1.4.1.2312.9.2.30393.1.4:
                ..0
            quantity
            1.3.6.1.4.1.2312.9.2.30393.1.3:
                ..25
            repo enabled setting
            1.3.6.1.4.1.2312.9.2.30393.1.8:
                ..1
```

14.15.3. The Structure of Product Certificates

The products that are installed on a system through the subscriptions assigned to a system are identified by X.509 certificates. When an available product is installed, the subscription service generates a product certificate, which contains the information about the product contract and the specific installation.

Structurally, entitlement certificates and product certificates are very similar, because they both provide much of the same information about products. The main difference is that a product certificate contains information about a single product that has been installed, so no other subscription information (like other available products or other product versions) is included in a product certificate the way that it is in an entitlement certificate.

A product certificate contains a single product namespace (meaning, a single product definition) which shows only what is actually installed on the system. The architecture and version definitions in a product certificate reflect the architecture and version of the product that is actually installed.

The product certificate is a *.pem file stored in the entitlement certificates directory, /etc/pki/product/product_serial#.pem. The name of the *.pem file is a numeric identifier that is generated by the subscription service. As with entitlement tracking, the generated ID is an inventory number, used to track installed products and associate them with systems within the subscription service.

14.15.4. Anatomy of Satellite Certificates

Important

Satellite certificates are used by Satellite 5.x deployments. They are not used on Red Hat Enterprise Linux 5.7 or by the subscription service.

Every system has to have a secure, authoritative way to identify what subscriptions are available. For Satellite 5.x systems, this identification is done through a digitally-signed XML document that lists the products and quantities that a customer has purchased.

As with entitlement certificates, a Satellite certificate contains the information about the subscription that was purchased, including the total number of systems that can be registered against that subscription and its start and end dates.

There are two types of subscriptions:

- System entitlements are subscriptions for services that can be performed, such as monitoring, provisioning, and virtualization.

- Channel entitlements, or content entitlements, provide access to the different software product download channels on Red Hat Network. These include Red Hat Enterprise Linux add-ons like Supplementary and FastTrack and layered products like Red Hat Directory Server.

Both types can be included in a single Satellite certificate.

A system entitlement and the metadata for an entitlement are both configured similarly in the certificate:

```
<rhn-cert-field name="configuration_area">value</rhn-cert-field>
```

The name argument identifies what entity is being configured. This can be the organization which ordered the subscription (name="owner"), the start and end dates for the entitlement (name="issued" and name="expires"), or the entitlement itself. A system entitlement uses the name argument to set the service being entitled; every content entitlement is set as a name="channel-family" type, with the specific product identified in an additional family argument.

The first section of the Satellite certificate is the metadata. The metadata identifies the organization which purchased it and the start and end dates of the entitlement. The field being set is in the name argument, while the value is between the tags. The last lines of the certificate also set metadata for the subscription, including the version of the Satellite and the signature that signs the XML document (and allows the XML file to be used as a certificate).

```

  <rhn-cert-field name="product">RHN-SATELLITE-001</rhn-cert-field>
  <rhn-cert-field name="owner">Example Corp</rhn-cert-field>
  <rhn-cert-field name="issued">2009-04-07 10:18:33</rhn-cert-field>
  <rhn-cert-field name="expires">2009-11-25 00:00:00</rhn-cert-field>

... [snip] ...

  <rhn-cert-field name="satellite-version">5.3</rhn-cert-field>
  <rhn-cert-field name="generation">2</rhn-cert-field>
  <rhn-cert-signature>
-----BEGIN PGP SIGNATURE-----
Version: Crypt::OpenPGP 1.03

iQBGBAARAwAGBQJJ22C+AAoJEJ5ynaAAAAkyyZ0An18+4hK5Ozt4HWieFvahsTnF
aPcaAJ0e5neOfdDZRLOgDE+Tp/Im3Hc3Rg==
=gqP7
-----END PGP SIGNATURE-----
</rhn-cert-signature>
```

The name="slot" field lists how many total systems are allowed to use this Satellite certificate to receive content. It is a global quantity.

```
  <rhn-cert-field name="slots">119</rhn-cert-field>
```

The system entitlements are set by identifying the service type in the name argument and then setting the quantity as the value within the tags.

```
  <rhn-cert-field name="provisioning-slots">117</rhn-cert-field>
  <rhn-cert-field name="monitoring-slots">20</rhn-cert-field>
  <rhn-cert-field name="virtualization_host">67</rhn-cert-field>
```

The content entitlements can include any combination of products, including base Red Hat Enterprise Linux subscriptions, variations of Red Hat Enterprise Linux, Red Hat Enterprise Linux add-ons, and general software products. General Red Hat Enterprise Linux server subscriptions are listed in the rhel-server family, while a specific Virtualization Server subscription provides an additional rhel-server-vt family.

```
<rhn-cert-field name="channel-families" quantity="95" family="rhel-server"/>
<rhn-cert-field name="channel-families" quantity="67" family="rhel-server-vt"/>
```

Add-ons and products for Red Hat Enterprise Linux systems (but not necessarily operating system products) are also in a rhel-* family, because that refers to the platform the product is supported on. In this example, Red Hat Directory Server is in the rhel-rhdirserv family.

```
<rhn-cert-field name="channel-families" quantity="3" family="rhel-rhdirserv"/>
```

Most subscriptions will also include a subscription tool set to manage and enable within clients features such as provisioning or configuration management when registered to RHN Classic or Satellite 5.x.

```
<rhn-cert-field name="channel-families" quantity="212" family="rhn-tools"/>
```

부 III. 네트워크 관련 설정

네트워크를 설정하는 방법에 관하여 설명한 후에, 원격 로그인을 허용하는 방법, 네트워크를 통한 파일 및 디렉토리를 공유하는 방법, 웹 서버를 설정하는 방법과 같은 네트워킹과 관련된 주제에 관하여 논의하겠습니다.

네트워크 인터페이스

Red Hat Enterprise Linux 에서, 모든 네트워크 통신은 설정된 소프트웨어 인터페이스와 시스템에 연결된 물리적 네트워크 장치 사이에서 발생합니다.

The configuration files for network interfaces are located in the /etc/sysconfig/network-scripts/ directory. The scripts used to activate and deactivate these network interfaces are also located here. Although the number and type of interface files can differ from system to system, there are three categories of files that exist in this directory:

1. Interface configuration files

2. Interface control scripts

3. Network function files

이러한 범주에 있는 파일이 함께 작동하여 여러 네트워크 장치를 활성화시킵니다.

이 장에서는 이러한 파일 사이의 관계와 이러한 파일들이 어떻게 사용되는지에 대해 살펴봅니다.

15.1. 네트워크 설정 파일

인터페이스 설정 파일을 살펴보기에 앞서, 우선 네트워크 설정에서 사용되는 기본 설정 파일을 항목별로 나열해 보겠습니다. 이러한 파일의 기능을 이해하여 네트워크 스택을 설정하는 것은 Red Hat Enterprise Linux 시스템을 사용자 설정할 때 도움이 됩니다.

기본 네트워크 설정 파일은 다음과 같습니다:

/etc/hosts

The main purpose of this file is to resolve hostnames that cannot be resolved any other way. It can also be used to resolve hostnames on small networks with no DNS server. Regardless of the type of network the computer is on, this file should contain a line specifying the IP address of the loopback device (127.0.0.1) as localhost.localdomain. For more information, refer to the hosts man page.

/etc/resolv.conf

This file specifies the IP addresses of DNS servers and the search domain. Unless configured to do otherwise, the network initialization scripts populate this file. For more information about this file, refer to the resolv.conf man page.

/etc/sysconfig/network

This file specifies routing and host information for all network interfaces. For more information about this file and the directives it accepts, refer to 30.1.21절. “/etc/sysconfig/network”.

/etc/sysconfig/network-scripts/ifcfg-<interface-name>

For each network interface, there is a corresponding interface configuration script. Each of these files provide information specific to a particular network interface. Refer to 15.2절. “인터페이스 설정 파일” for more information on this type of file and the directives it accepts.

주의

The /etc/sysconfig/networking/ directory is used by the Network Administration Tool (system-config-network) and its contents should not be edited manually. Using only one method for network configuration is strongly encouraged, due to the risk of configuration deletion.

For more information about configuring network interfaces using the Network Administration Tool, refer to 16장. 네트워크 설정

15.2. 인터페이스 설정 파일

Interface configuration files control the software interfaces for individual network devices. As the system boots, it uses these files to determine what interfaces to bring up and how to configure them. These files are usually named ifcfg-<name> , where <name> refers to the name of the device that the configuration file controls.

15.2.1. 이더넷 인터페이스

One of the most common interface files is ifcfg-eth0, which controls the first Ethernet network interface card or NIC in the system. In a system with multiple NICs, there are multiple ifcfg-eth<X> files (where <X> is a unique number corresponding to a specific interface). Because each device has its own configuration file, an administrator can control how each interface functions individually.

The following is a sample ifcfg-eth0 file for a system using a fixed IP address:

```
DEVICE=eth0
BOOTPROTO=none
ONBOOT=yes
NETWORK=10.0.1.0
NETMASK=255.255.255.0
IPADDR=10.0.1.27
USERCTL=no
```

The values required in an interface configuration file can change based on other values. For example, the ifcfg-eth0 file for an interface using DHCP looks different because IP information is provided by the DHCP server:

```
DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
```

The Network Administration Tool (system-config-network) is an easy way to make changes to the various network interface configuration files (refer to 16장. 네트워크 설정 for detailed instructions on using this tool).

However, it is also possible to manually edit the configuration files for a given network interface.

Below is a listing of the configurable parameters in an Ethernet interface configuration file:

BONDING_OPTS=<parameters>

sets the configuration parameters for the bonding device, and is used in /etc/sysconfig/network-scripts/ifcfg-bond<N> (see 15.2.3절. “채널 결합 인터페이스”). These parameters are identical to those used for bonding devices in /sys/class/net/<bonding device>/bonding, and the module parameters for the bonding driver as described in bonding Module Directives.

This configuration method is used so that multiple bonding devices can have different configurations. If you use BONDING_OPTS in ifcfg-<name> , do not use /etc/modprobe.conf to specify options for the bonding device.

BOOTPROTO=<protocol>

where <protocol> is one of the following:

- none — No boot-time protocol should be used.
- bootp — The BOOTP protocol should be used.
- dhcp — The DHCP protocol should be used.

BROADCAST=<address>

where <address> is the broadcast address. This directive is deprecated, as the value is calculated automatically with ipcalc.

DEVICE=<name>

where <name> is the name of the physical device (except for dynamically-allocated PPP devices where it is the logical name).

DHCP_HOSTNAME=<name>

where <name> is a short hostname to be sent to the DHCP server. Use this option only if the DHCP server requires the client to specify a hostname before receiving an IP address.

DNS{1,2}=<address>

where <address> is a name server address to be placed in /etc/resolv.conf if the PEERDNS directive is set to yes.

ETHTOOL_OPTS=<options>

where <options> are any device-specific options supported by ethtool. For example, if you wanted to force 100Mb, full duplex:

```
ETHTOOL_OPTS="autoneg off speed 100 duplex full"
```

Instead of a custom initscript, use ETHTOOL_OPTS to set the interface speed and duplex settings. Custom initscripts run outside of the network init script lead to unpredictable results during a post-boot network service restart.

> 참고
>
> Changing speed or duplex settings almost always requires disabling autonegotiation with the autoneg off option. This needs to be stated first, as the option entries are order-dependent.

GATEWAY=<address>

where <address> is the IP address of the network router or gateway device (if any).

HOTPLUG=<answer>
where <answer> is one of the following:

- yes — This device should be activated when it is hot-plugged (this is the default option).

- no — This device should not be activated when it is hot-plugged.

The HOTPLUG=no option can be used to prevent a channel bonding interface from being activated when a bonding kernel module is loaded.

Refer to 15.2.3절. "채널 결합 인터페이스" for more about channel bonding interfaces.

HWADDR=<MAC-address>
where <MAC-address> is the hardware address of the Ethernet device in the form AA:BB:CC:DD:EE:FF. This directive must be used in machines containing more than one NIC to ensure that the interfaces are assigned the correct device names regardless of the configured load order for each NIC's module. This directive should not be used in conjunction with MACADDR.

IPADDR=<address>
where <address> is the IP address.

LINKDELAY=<time>
where <time> is the number of seconds to wait for link negotiation before configuring the device.

MACADDR=<MAC-address>
where <MAC-address> is the hardware address of the Ethernet device in the form AA:BB:CC:DD:EE:FF. This directive is used to assign a MAC address to an interface, overriding the one assigned to the physical NIC. This directive should not be used in conjunction with HWADDR.

MASTER=<bond-interface>
where <bond-interface> is the channel bonding interface to which the Ethernet interface is linked.

This directive is used in conjunction with the SLAVE directive.

Refer to 15.2.3절. "채널 결합 인터페이스" for more information about channel bonding interfaces.

NETMASK=<mask>
where <mask> is the netmask value.

NETWORK=<address>
where <address> is the network address. This directive is deprecated, as the value is calculated automatically with ipcalc.

ONBOOT=<answer>
where <answer> is one of the following:

- yes — This device should be activated at boot-time.

- no — This device should not be activated at boot-time.

PEERDNS=<answer>
where <answer> is one of the following:

- yes — Modify /etc/resolv.conf if the DNS directive is set. If using DHCP, then yes is the default.

- no — Do not modify /etc/resolv.conf.

SLAVE=<answer>
where <answer> is one of the following:

- yes — This device is controlled by the channel bonding interface specified in the MASTER directive.

- no — This device is not controlled by the channel bonding interface specified in the MASTER directive.

This directive is used in conjunction with the MASTER directive.

Refer to 15.2.3절. "채널 결합 인터페이스" for more about channel bonding interfaces.

SRCADDR=<address>
where <address> is the specified source IP address for outgoing packets.

USERCTL=<answer>
where <answer> is one of the following:

- yes — Non-root users are allowed to control this device.

- no — Non-root users are not allowed to control this device.

15.2.2. IPsec 인터페이스

The following example shows the ifcfg file for a network-to-network IPsec connection for LAN A. The unique name to identify the connection in this example is ipsec1, so the resulting file is named /etc/sysconfig/network-scripts/ifcfg-ipsec1.

```
TYPE=IPsec
ONBOOT=yes
IKE_METHOD=PSK
SRCNET=192.168.1.0/24
DSTNET=192.168.2.0/24
DST=X.X.X.X
```

In the example above, X.X.X.X is the publicly routable IP address of the destination IPsec router.

Below is a listing of the configurable parameters for an IPsec interface:

DST=<address>
where <address> is the IP address of the IPsec destination host or router. This is used for both host-to-host and network-to-network IPsec configurations.

DSTNET=<network>
where <network> is the network address of the IPsec destination network. This is only used for network-to-network IPsec configurations.

SRC=<address>
where <address> is the IP address of the IPsec source host or router. This setting is optional and is only used for host-to-host IPsec configurations.

SRCNET=<network>

where <network> is the network address of the IPsec source network. This is only used for network-to-network IPsec configurations.

TYPE=<interface-type>

where <interface-type> is IPSEC. Both applications are part of the ipsec-tools package.

If manual key encryption with IPsec is being used, refer to /usr/share/doc/initscripts-<version-number>/sysconfig.txt (replace <version-number> with the version of the initscripts package installed) for configuration parameters.

The racoon IKEv1 key management daemon negotiates and configures a set of parameters for IPSec. It can use preshared keys, RSA signatures, or GSS-API. If racoon is used to automatically manage key encryption, the following options are required:

IKE_METHOD=<encryption-method>

where <encryption-method> is either PSK, X509, or GSSAPI. If PSK is specified, the IKE_PSK parameter must also be set. If X509 is specified, the IKE_CERTFILE parameter must also be set.

IKE_PSK=<shared-key>

where <shared-key> is the shared, secret value for the PSK (preshared keys) method.

IKE_CERTFILE=<cert-file>

where <cert-file> is a valid X.509 certificate file for the host.

IKE_PEER_CERTFILE=<cert-file>

where <cert-file> is a valid X.509 certificate file for the remote host.

IKE_DNSSEC=<answer>

where <answer> is yes. The racoon daemon retrieves the remote host's X.509 certificate via DNS. If a IKE_PEER_CERTFILE is specified, do not include this parameter.

For more information about the encryption algorithms available for IPsec, refer to the setkey man page. For more information about racoon, refer to the racoon and racoon.conf man pages.

15.2.3. 채널 결합 인터페이스

Red Hat Enterprise Linux allows administrators to bind multiple network interfaces together into a single channel using the bonding kernel module and a special network interface called a channel bonding interface. Channel bonding enables two or more network interfaces to act as one, simultaneously increasing the bandwidth and providing redundancy.

To create a channel bonding interface, create a file in the /etc/sysconfig/network-scripts/ directory called ifcfg-bond<N> , replacing <N> with the number for the interface, such as 0.

The contents of the file can be identical to whatever type of interface is getting bonded, such as an Ethernet interface. The only difference is that the DEVICE= directive must be bond<N> , replacing <N> with the number for the interface.

The following is a sample channel bonding configuration file, ifcfg-bond0:

```
DEVICE=bond0
IPADDR=192.168.1.1
NETMASK=255.255.255.0
ONBOOT=yes
BOOTPROTO=none
```

```
USERCTL=no
BONDING_OPTS="<bonding parameters separated by spaces>"
```

After the channel bonding interface is created, the network interfaces to be bound together must be configured by adding the MASTER= and SLAVE= directives to their configuration files. The configuration files for each of the channel-bonded interfaces can be nearly identical.

For example, if two Ethernet interfaces are being channel bonded, both eth0 and eth1 may look like the following example:

```
DEVICE=eth<N>
BOOTPROTO=none
ONBOOT=yes
MASTER=bond0
SLAVE=yes
USERCTL=no
```

In this example, replace <N> with the numerical value for the interface.

For a channel bonding interface to be valid, the kernel module must be loaded. To ensure that the module is loaded when the channel bonding interface is brought up, add the following line to /etc/modprobe.conf:

```
alias bond<N> bonding
```

Replace <N> with the number of the interface, such as 0.

중요

Do not place parameters for the bonding kernel module in the /etc/modprobe.conf file. Instead, specify them as a space-separated list in the BONDING_OPTS="<bonding parameters>" directive in the ifcfg-bond<N> interface file.

The only exception is the debug parameter, which cannot be used on a per-device basis, and which should therefore be specified in /etc/modprobe.conf as follows:

```
options bonding debug=1
```

For further instructions and advice on configuring the bonding module, as well as to view the list of bonding parameters, refer to 43.5.2절. "The Channel Bonding Module" .

15.2.4. 별칭 및 복제 파일

Two lesser-used types of interface configuration files are alias and clone files.

Alias interface configuration files, which are used to bind multiple addresses to a single interface, use the ifcfg-<if-name>:<alias-value> naming scheme.

For example, an ifcfg-eth0:0 file could be configured to specify DEVICE=eth0:0 and a static IP address of 10.0.0.2, serving as an alias of an Ethernet interface already configured to receive its IP information via DHCP in ifcfg-eth0. Under this configuration, eth0 is bound to a dynamic IP address, but the same physical network card can receive requests via the fixed, 10.0.0.2 IP address.

주의

별칭 인터페이스가 DHCP를 지원하지 않습니다.

A clone interface configuration file should use the following naming convention: ifcfg-<if-name>-<clone-name> . While an alias file allows multiple addresses for an existing interface, a clone file is used to specify additional options for an interface. For example, a standard DHCP Ethernet interface called eth0, may look similar to this:

```
DEVICE=eth0
ONBOOT=yes
BOOTPROTO=dhcp
```

Since the default value for the USERCTL directive is no if it is not specified, users cannot bring this interface up and down. To give users the ability to control the interface, create a clone by copying ifcfg-eth0 to ifcfg-eth0-user and add the following line to ifcfg-eth0-user:

```
USERCTL=yes
```

This way a user can bring up the eth0 interface using the /sbin/ifup eth0-user command because the configuration options from ifcfg-eth0 and ifcfg-eth0-user are combined. While this is a very basic example, this method can be used with a variety of options and interfaces.

The easiest way to create alias and clone interface configuration files is to use the graphical Network Administration Tool. For more information on using this tool, refer to 16장. 네트워크 설정.

15.2.5. 전화연결 인터페이스

If you are connecting to the Internet via a dialup connection, a configuration file is necessary for the interface.

PPP 인터페이스 파일은 다음의 포멧으로 사용됩니다:

ifcfg-ppp<X>

where <X> is a unique number corresponding to a specific interface.

The PPP interface configuration file is created automatically when wvdial, the Network Administration Tool or Kppp is used to create a dialup account. It is also possible to create and edit this file manually.

The following is a typical ifcfg-ppp0 file:

```
DEVICE=ppp0
NAME=test
WVDIALSECT=test
MODEMPORT=/dev/modem
LINESPEED=115200
PAPNAME=test
USERCTL=true
ONBOOT=no
PERSIST=no
DEFROUTE=yes
PEERDNS=yes
DEMAND=no
```

```
IDLETIMEOUT=600
```

Serial Line Internet Protocol (SLIP) is another dialup interface, although it is used less frequently. SLIP files have interface configuration file names such as ifcfg-sl0.

Other options that may be used in these files include:

DEFROUTE=<answer>
where <answer> is one of the following:

- yes — Set this interface as the default route.
- no — Do not set this interface as the default route.

DEMAND=<answer>
where <answer> is one of the following:

- yes — This interface allows pppd to initiate a connection when someone attempts to use it.
- no — A connection must be manually established for this interface.

IDLETIMEOUT=<value>
where <value> is the number of seconds of idle activity before the interface disconnects itself.

INITSTRING=<string>
where <string> is the initialization string passed to the modem device. This option is primarily used in conjunction with SLIP interfaces.

LINESPEED=<value>
where <value> is the baud rate of the device. Possible standard values include 57600, 38400, 19200, and 9600.

MODEMPORT=<device>
where <device> is the name of the serial device that is used to establish the connection for the interface.

MTU=<value>
where <value> is the Maximum Transfer Unit (MTU) setting for the interface. The MTU refers to the largest number of bytes of data a frame can carry, not counting its header information. In some dialup situations, setting this to a value of 576 results in fewer packets dropped and a slight improvement to the throughput for a connection.

NAME=<name>
where <name> is the reference to the title given to a collection of dialup connection configurations.

PAPNAME=<name>
where <name> is the username given during the Password Authentication Protocol (PAP) exchange that occurs to allow connections to a remote system.

PERSIST=<answer>
where <answer> is one of the following:

- yes — This interface should be kept active at all times, even if deactivated after a modem hang up.
- no — This interface should not be kept active at all times.

REMIP=<address>
where <address> is the IP address of the remote system. This is usually left unspecified.

WVDIALSECT=<name>
where <name> associates this interface with a dialer configuration in /etc/wvdial.conf. This file contains the phone number to be dialed and other important information for the interface.

15.2.6. 기타 인터페이스

기타 다른 일반적인 인터페이스 설정 파일은 다음과 같습니다:

ifcfg-lo
A local loopback interface is often used in testing, as well as being used in a variety of applications that require an IP address pointing back to the same system. Any data sent to the loopback device is immediately returned to the host's network layer.

> **주의**
>
> The loopback interface script, /etc/sysconfig/network-scripts/ifcfg-lo, should never be edited manually. Doing so can prevent the system from operating correctly.

ifcfg-irlan0
An infrared interface allows information between devices, such as a laptop and a printer, to flow over an infrared link. This works in a similar way to an Ethernet device except that it commonly occurs over a peer-to-peer connection.

ifcfg-plip0
A Parallel Line Interface Protocol (PLIP) connection works much the same way as an Ethernet device, except that it utilizes a parallel port.

ifcfg-tr0
Token Ring topologies are not as common on Local Area Networks (LANs) as they once were, having been eclipsed by Ethernet.

15.3. 인터페이스 제어 스크립트

The interface control scripts activate and deactivate system interfaces. There are two primary interface control scripts that call on control scripts located in the /etc/sysconfig/network-scripts/ directory: /sbin/ifdown and /sbin/ifup.

The ifup and ifdown interface scripts are symbolic links to scripts in the /sbin/ directory. When either of these scripts are called, they require the value of the interface to be specified, such as:

```
ifup eth0
```

주의

> The ifup and ifdown interface scripts are the only scripts that the user should use to bring up and take down network interfaces.
>
> The following scripts are described for reference purposes only.

Two files used to perform a variety of network initialization tasks during the process of bringing up a network interface are /etc/rc.d/init.d/functions and /etc/sysconfig/network-scripts/network-functions. Refer to 15.5절. “네트워크 기능 파일” for more information.

After verifying that an interface has been specified and that the user executing the request is allowed to control the interface, the correct script brings the interface up or down. The following are common interface control scripts found within the /etc/sysconfig/network-scripts/ directory:

ifup-aliases
: Configures IP aliases from interface configuration files when more than one IP address is associated with an interface.

ifup-ippp and ifdown-ippp
: Brings ISDN interfaces up and down.

ifup-ipsec and ifdown-ipsec
: Brings IPsec interfaces up and down.

ifup-ipv6 and ifdown-ipv6
: Brings IPv6 interfaces up and down.

ifup-ipx
: IPX 인터페이스 활성화

ifup-plip
: PLIP 인터페이스 활성화

ifup-plusb
: 네트워크 연결을 위한 USB 인터페이스 활성화

ifup-post and ifdown-post
: Contains commands to be executed after an interface is brought up or down.

ifup-ppp and ifdown-ppp
: Brings a PPP interface up or down.

ifup-routes
: Adds static routes for a device as its interface is brought up.

ifdown-sit and ifup-sit
: Contains function calls related to bringing up and down an IPv6 tunnel within an IPv4 connection.

ifup-sl and ifdown-sl
: Brings a SLIP interface up or down.

ifup-wireless
 무선 인터페이스 활성화

주의

Removing or modifying any scripts in the /etc/sysconfig/network-scripts/ directory can cause interface connections to act irregularly or fail. Only advanced users should modify scripts related to a network interface.

The easiest way to manipulate all network scripts simultaneously is to use the /sbin/service command on the network service (/etc/rc.d/init.d/network), as illustrated the following command:

```
service network <action>
```

Here, <action> can be either start, stop, or restart.

To view a list of configured devices and currently active network interfaces, use the following command:

```
service network status
```

15.4. Configuring Static Routes

Routing will be configured on routing devices, therefore it should not be necessary to configure static routes on Red Hat Enterprise Linux servers or clients. However, if static routes are required they can be configured for each interface. This can be useful if you have multiple interfaces in different subnets. Use the route command to display the IP routing table.

Static route configuration is stored in a /etc/sysconfig/network-scripts/route-interface file. For example, static routes for the eth0 interface would be stored in the /etc/sysconfig/network-scripts/route-eth0 file. The route-interface file has two formats: IP command arguments and network/netmask directives.

IP Command Arguments Format

Define a default gateway on the first line. This is only required if the default gateway is not set via DHCP:

```
default via X.X.X.X dev interface
```

X.X.X.X is the IP address of the default gateway. The interface is the interface that is connected to, or can reach, the default gateway.

Define a static route. Each line is parsed as an individual route:

```
X.X.X.X/X via X.X.X.X dev interface
```

X.X.X.X/X is the network number and netmask for the static route. X.X.X.X and interface are the IP address and interface for the default gateway respectively. The X.X.X.X address does not have to be the default gateway IP address. In most cases, X.X.X.X will be an IP address in a different subnet, and interface will be the interface that is connected to, or can reach, that subnet. Add as many static routes as required.

The following is a sample route-eth0 file using the IP command arguments format. The default gateway is 192.168.0.1, interface eth0. The two static routes are for the 10.10.10.0/24 and 172.16.1.0/24 networks:

```
default via 192.168.0.1 dev eth0
10.10.10.0/24 via 192.168.0.1 dev eth0
172.16.1.0/24 via 192.168.0.1 dev eth0
```

Static routes should only be configured for other subnets. The above example is not necessary, since packets going to the 10.10.10.0/24 and 172.16.1.0/24 networks will use the default gateway anyway. Below is an example of setting static routes to a different subnet, on a machine in a 192.168.0.0/24 subnet. The example machine has an eth0 interface in the 192.168.0.0/24 subnet, and an eth1 interface (10.10.10.1) in the 10.10.10.0/24 subnet:

```
10.10.10.0/24 via 10.10.10.1 dev eth1
```

Duplicate Default Gateways

If the default gateway is already assigned from DHCP, the IP command arguments format can cause one of two errors during start-up, or when bringing up an interface from the down state using the ifup command: "RTNETLINK answers: File exists" or 'Error: either "to" is a duplicate, or "X.X.X.X" is a garbage.', where X.X.X.X is the gateway, or a different IP address. These errors can also occur if you have another route to another network using the default gateway. Both of these errors are safe to ignore.

Network/Netmask Directives Format

You can also use the network/netmask directives format for route-interface files. The following is a template for the network/netmask format, with instructions following afterwards:

```
ADDRESS0=X.X.X.X
NETMASK0=X.X.X.X
GATEWAY0=X.X.X.X
```

- ADDRESS0=X.X.X.X is the network number for the static route.
- NETMASK0=X.X.X.X is the netmask for the network number defined with ADDRESS0=X.X.X.X .
- GATEWAY0=X.X.X.X is the default gateway, or an IP address that can be used to reach ADDRESS0=X.X.X.X

The following is a sample route-eth0 file using the network/netmask directives format. The default gateway is 192.168.0.1, interface eth0. The two static routes are for the 10.10.10.0/24 and 172.16.1.0/24 networks. However, as mentioned before, this example is not necessary as the 10.10.10.0/24 and 172.16.1.0/24 networks would use the default gateway anyway:

```
ADDRESS0=10.10.10.0
NETMASK0=255.255.255.0
```

```
GATEWAY0=192.168.0.1
ADDRESS1=172.16.1.0
NETMASK1=255.255.255.0
GATEWAY1=192.168.0.1
```

Subsequent static routes must be numbered sequentially, and must not skip any values. For example, ADDRESS0, ADDRESS1, ADDRESS2, and so on.

Below is an example of setting static routes to a different subnet, on a machine in the 192.168.0.0/24 subnet. The example machine has an eth0 interface in the 192.168.0.0/24 subnet, and an eth1 interface (10.10.10.1) in the 10.10.10.0/24 subnet:

```
ADDRESS0=10.10.10.0
NETMASK0=255.255.255.0
GATEWAY0=10.10.10.1
```

DHCP should assign these settings automatically, therefore it should not be necessary to configure static routes on Red Hat Enterprise Linux servers or clients.

15.5. 네트워크 기능 파일

Red Hat Enterprise Linux makes use of several files that contain important common functions used to bring interfaces up and down. Rather than forcing each interface control file to contain these functions, they are grouped together in a few files that are called upon when necessary.

The /etc/sysconfig/network-scripts/network-functions file contains the most commonly used IPv4 functions, which are useful to many interface control scripts. These functions include contacting running programs that have requested information about changes in the status of an interface, setting hostnames, finding a gateway device, verifying whether or not a particular device is down, and adding a default route.

As the functions required for IPv6 interfaces are different from IPv4 interfaces, a /etc/sysconfig/network-scripts/network-functions-ipv6 file exists specifically to hold this information. The functions in this file configure and delete static IPv6 routes, create and remove tunnels, add and remove IPv6 addresses to an interface, and test for the existence of an IPv6 address on an interface.

15.6. 추가 리소스

The following are resources which explain more about network interfaces.

15.6.1. 설치된 문서

/usr/share/doc/initscripts-<version>/sysconfig.txt

A guide to available options for network configuration files, including IPv6 options not covered in this chapter.

/usr/share/doc/iproute-<version>/ip-cref.ps

This file contains a wealth of information about the ip command, which can be used to manipulate routing tables, among other things. Use the ggv or kghostview application to view this file.

네트워크 설정

컴퓨터는 네트워크로 연결되어 다른 컴퓨터와 통신을 주고 받습니다. 이는 인터페이스 카드를 (예, 이더넷, ISDN 모뎀, 토큰 링)인식하는 운영 체제에서 네트워크에 연결하는 인터페이스를 설정하시면 됩니다.

The Network Administration Tool can be used to configure the following types of network interfaces:

- 이더넷 (Ethernet)
- ISDN
- 모뎀
- xDSL
- 토큰 링
- CIPE
- 무선 장치

It can also be used to configure IPsec connections, manage DNS settings, and manage the /etc/hosts file used to store additional hostnames and IP address combinations.

To use the Network Administration Tool, you must have root privileges. To start the application, go to the Applications (the main menu on the panel) > System Settings > Network, or type the command system-config-network at a shell prompt (for example, in an XTerm or a GNOME terminal). If you type the command, the graphical version is displayed if X is running; otherwise, the text-based version is displayed.

To use the command line version, execute the command system-config-network-cmd --help as root to view all of the options.

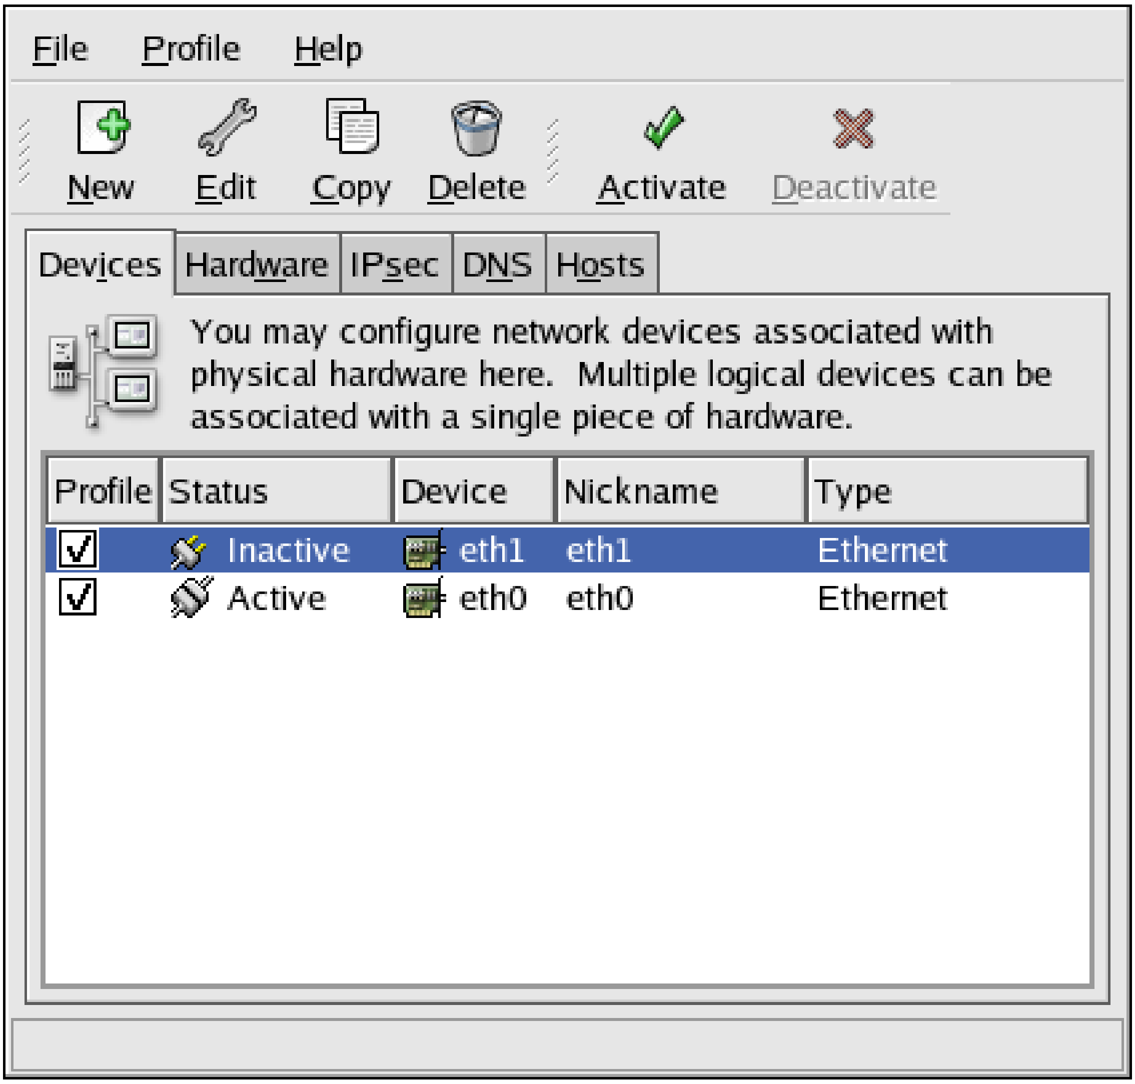

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	eth1	eth1	Ethernet
☑	Active	eth0	eth0	Ethernet

그림 16.1. Network Administration Tool

Tip

가지고 계신 하드웨어 장치가 Red Hat Enterprise Linux에서 지원되는지 알아보기 위하여 Red Hat 하드웨어 호환성 목록을 살펴보시기 바랍니다. (http://hardware.redhat.com/hcl/)

16.1. 개요

To configure a network connection with the Network Administration Tool, perform the following steps:

1. 물리적 하드웨어 장치와 연관된 네트워크 장치를 추가합니다.

2. 하드웨어 목록에 물리적 하드웨어 장치가 없을 경우 이를 추가합니다.

3. 호스트명과 DNS 셋팅을 설정합니다.

4. DNS를 통하여 검색할 수 없는 호스트를 설정하십시오.

이 장에서는 여러 유형의 네트워크 연결을 설정하는 방법에 대하여 설명해 보겠습니다.

16.2. 이더넷 연결 설정하기

이더넷 연결을 설정하기 위해서는, 네트워크 인터페이스 카드 (NIC)와 네트워크 케이블 (보통 CAT5 케이블), 그리고 연결할 네트워크가 필요합니다. 네트워크 종류에 따라 속도가 다르기 때문에, 가지고 계신 네트워크 인터페이스 카드가 연결하려는 네트워크과 호환 가능한지를 확인해 주십시오.

이더넷 연결을 추가하시려면, 다음의 단계를 따르십시오:

1. Click the Devices tab.

2. Click the New button on the toolbar.

3. Select Ethernet connection from the Device Type list, and click Forward.

4. If you have already added the network interface card to the hardware list, select it from the Ethernet card list. Otherwise, select Other Ethernet Card to add the hardware device.

   알림

   The installation program detects supported Ethernet devices and prompts you to configure them. If you configured any Ethernet devices during the installation, they are displayed in the hardware list on the Hardware tab.

5. If you selected Other Ethernet Card, the Select Ethernet Adapter window appears. Select the manufacturer and model of the Ethernet card. Select the device name. If this is the system's first Ethernet card, select eth0 as the device name; if this is the second Ethernet card, select eth1 (and so on). The Network Administration Tool also allows you to configure the resources for the NIC. Click Forward to continue.

6. In the Configure Network Settings window shown in 그림 16.2. “Ethernet Settings” , choose between DHCP and a static IP address. If the device receives a different IP address each time the network is started, do not specify a hostname.

7. Do not specify a value for the Set MTU to or Set MRU to fields. MTU stands for Maximum Transmission Unit and MRU for Maximum Receive Unit; the network configuration tool will choose appropriate values for both of these parameters. Click Forward to continue.

8. Click Apply on the Create Ethernet Device page.

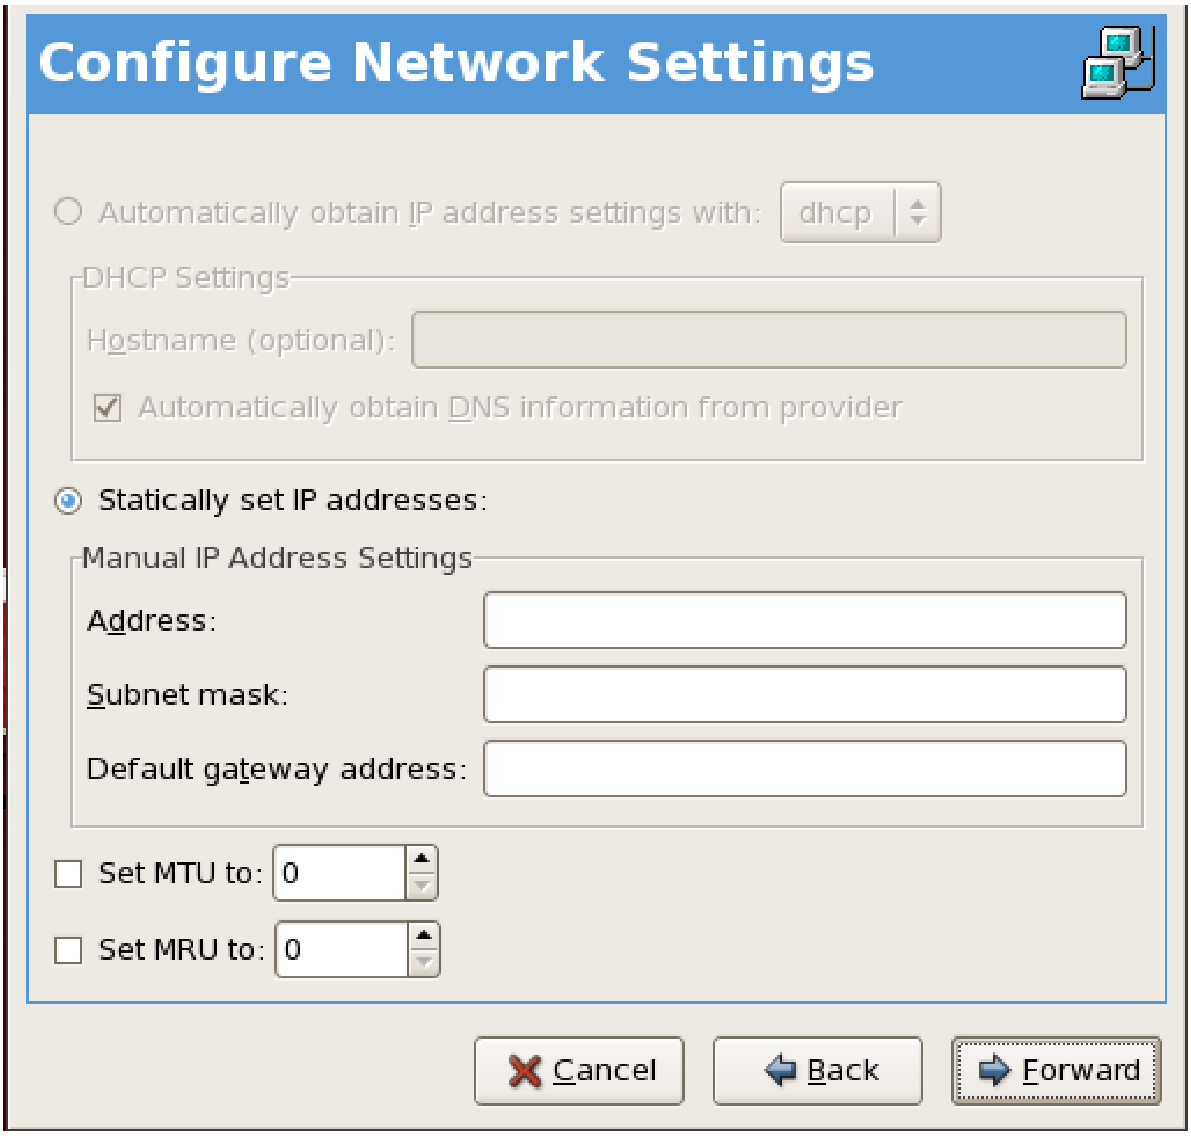

Configure Network Settings

○ Automatically obtain IP address settings with: dhcp

DHCP Settings

Hostname (optional):

☑ Automatically obtain DNS information from provider

● Statically set IP addresses:

Manual IP Address Settings

Address:

Subnet mask:

Default gateway address:

□ Set MTU to: 0

□ Set MRU to: 0

그림 16.2. Ethernet Settings

After configuring the Ethernet device, it appears in the device list as shown in 그림 16.3. "Ethernet Device" .

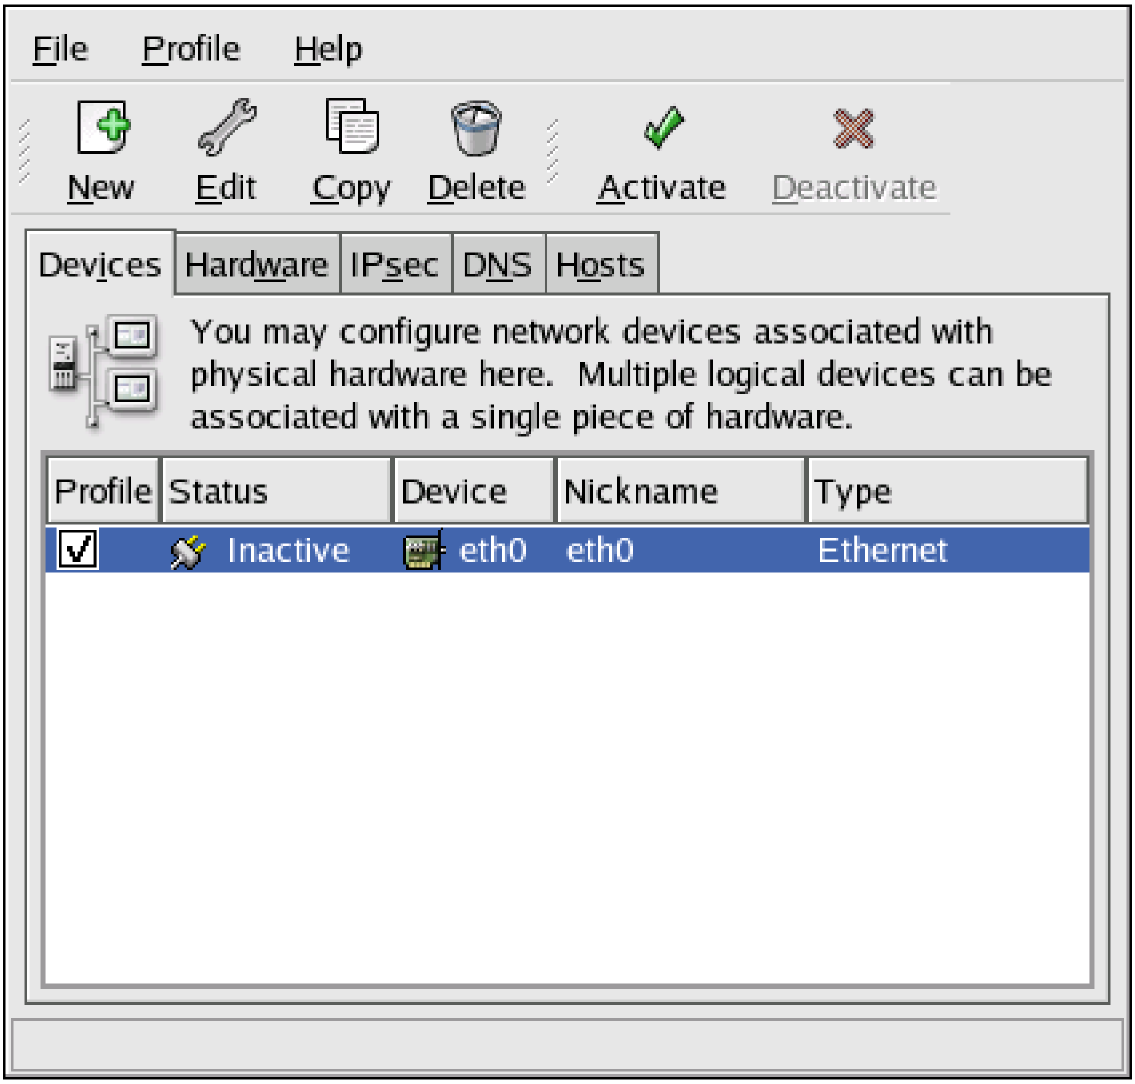

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	eth0	eth0	Ethernet

그림 16.3. Ethernet Device

Be sure to select File > Save to save the changes.

After adding the Ethernet device, you can edit its configuration by selecting the device from the device list and clicking Edit. For example, when the device is added, it is configured to start at boot time by default. To change this setting, select to edit the device, modify the Activate device when computer starts value, and save the changes.

When the device is added, it is not activated immediately, as seen by its Inactive status. To activate the device, select it from the device list, and click the Activate button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

If you associate more than one device with an Ethernet card, the subsequent devices are device aliases. A device alias allows you to setup multiple virtual devices for one physical device, thus giving the one physical device more than one IP address. For example, you can configure an eth1 device and an eth1:1 device. For details, refer to 16.11절. “장치 별칭” .

16.3. ISDN 연결 설정하기

ISDN 연결은 통신 회사에 의해 설치된 특별한 전화선을 통하여 ISDN 모뎀 카드를 이용하여 설정된 인터넷 연결을 의미합니다. ISDN 연결은 유럽에서 대중적입니다.

ISDN 연결을 추가하기 위해서는 다음의 단계를 따르십시오:

1. Click the Devices tab.

2. Click the New button on the toolbar.

3. Select ISDN connection from the Device Type list, and click Forward.

4. Select the ISDN adapter from the pulldown menu. Then configure the resources and D channel protocol for the adapter. Click Forward to continue.

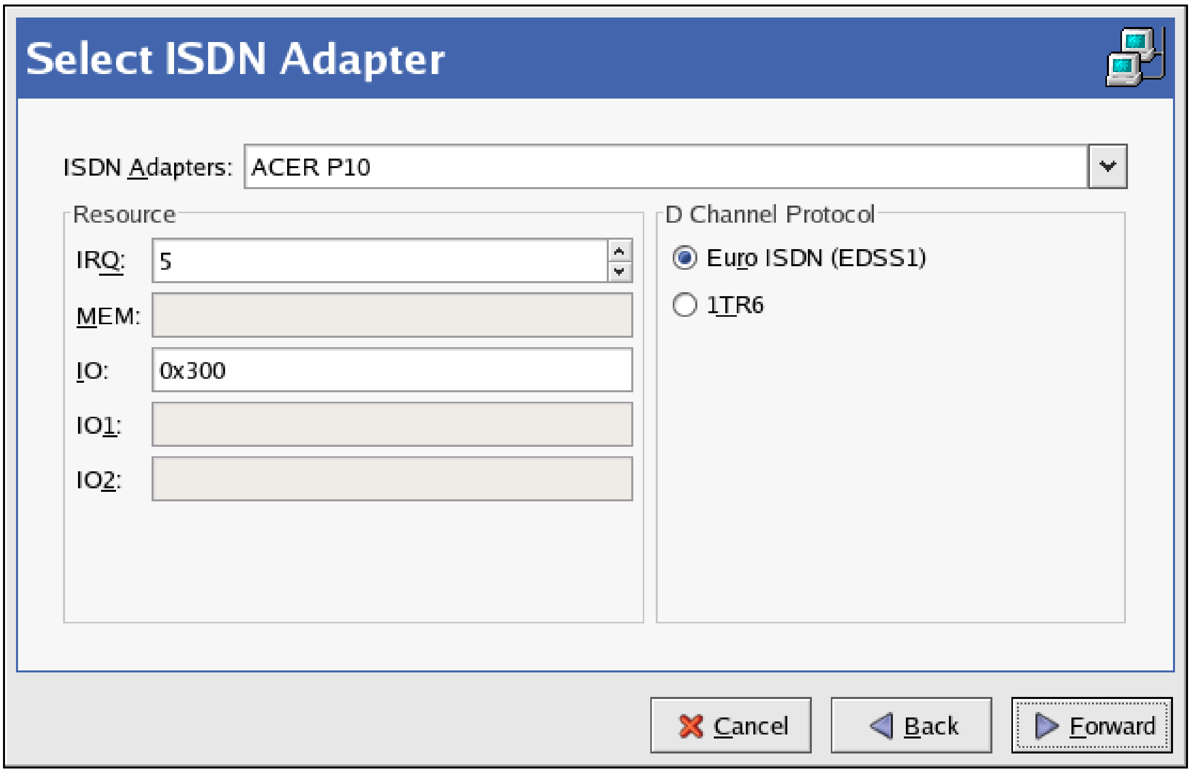

Select ISDN Adapter

ISDN Adapters: ACER P10

Resource

IRQ: 5

MEM:

IO: 0x300

IO1:

IO2:

D Channel Protocol

● Euro ISDN (EDSS1)

○ 1TR6

그림 16.4. ISDN Settings

5. If your Internet Service Provider (ISP) is in the pre-configured list, select it. Otherwise, enter the required information about your ISP account. If you do not know the values, contact your ISP. Click Forward.

6. In the IP Settings window, select the Encapsulation Mode and whether to obtain an IP address automatically or to set a static IP instead. Click Forward when finished.

7. On the Create Dialup Connection page, click Apply.

After configuring the ISDN device, it appears in the device list as a device with type ISDN as shown in 그림 16.5. "ISDN Device" .

Be sure to select File > Save to save the changes.

After adding the ISDN device, you can edit its configuration by selecting the device from the device list and clicking Edit. For example, when the device is added, it is configured not to start at boot time by default. Edit its configuration to modify this setting. Compression, PPP options, login name, password, and more can be changed.

When the device is added, it is not activated immediately, as seen by its Inactive status. To activate the device, select it from the device list, and click the Activate button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

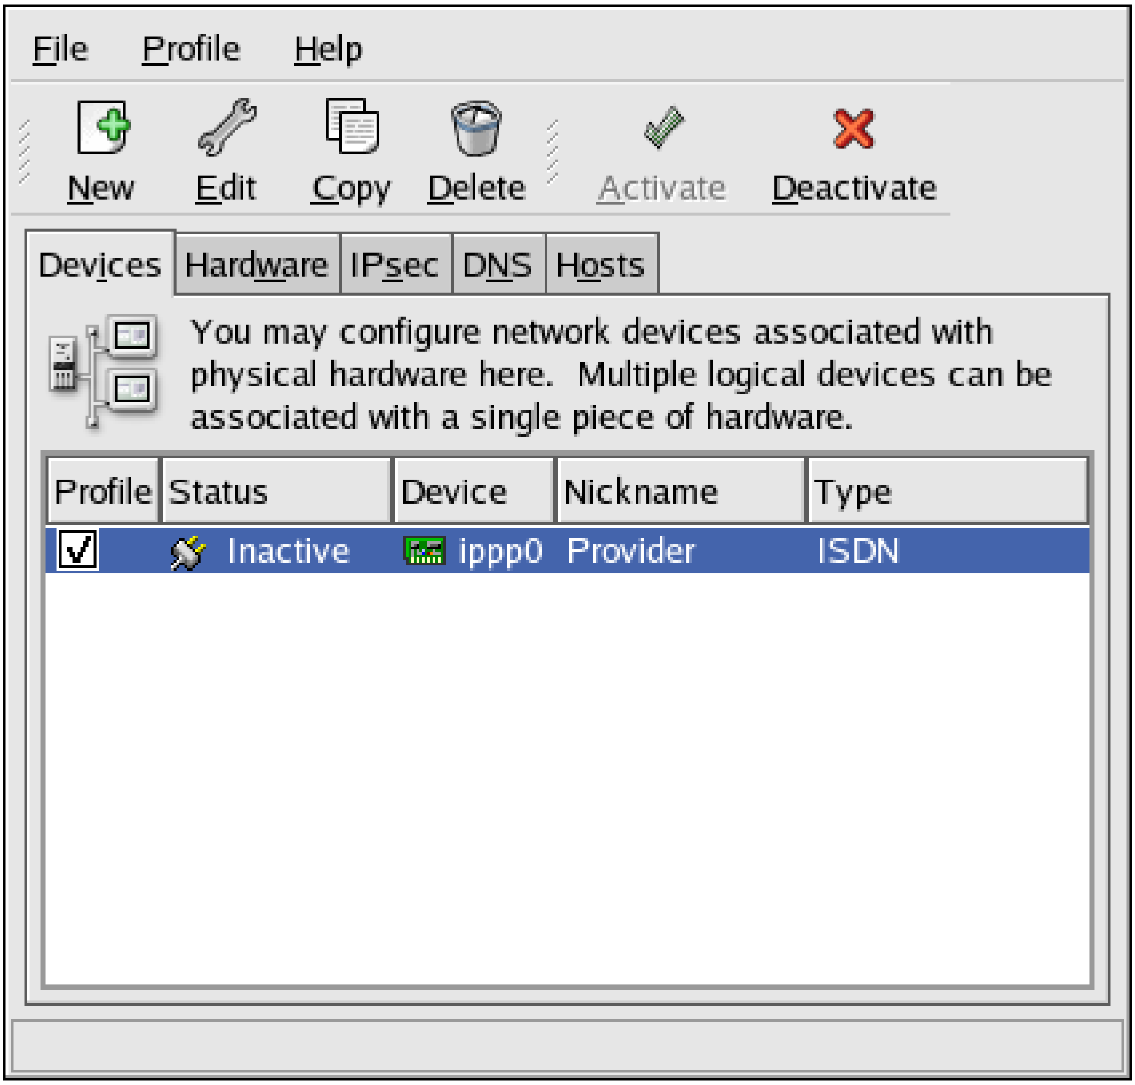

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	ippp0	Provider	ISDN

그림 16.5. ISDN Device

16.4. 모뎀 연결 설정하기

모뎀을 사용하여 전화선을 통한 인터넷 연결을 설정하실 수 있습니다. 전화걸기 계정이라고도 부르는 ISP (인터넷 제공 사업자) 계정이 필요합니다.

모뎀 연결을 추가하기 위해서는 다음의 단계를 따르십시오:

1. Click the Devices tab.

2. Click the New button on the toolbar.

3. Select Modem connection from the Device Type list, and click Forward.

4. If there is a modem already configured in the hardware list (on the Hardware tab), the Network Administration Tool assumes you want to use it to establish a modem connection. If there are no

modems already configured, it tries to detect any modems in the system. This probe might take a while. If a modem is not found, a message is displayed to warn you that the settings shown are not values found from the probe.

5. After probing, the window in 그림 16.6. "Modem Settings" appears.

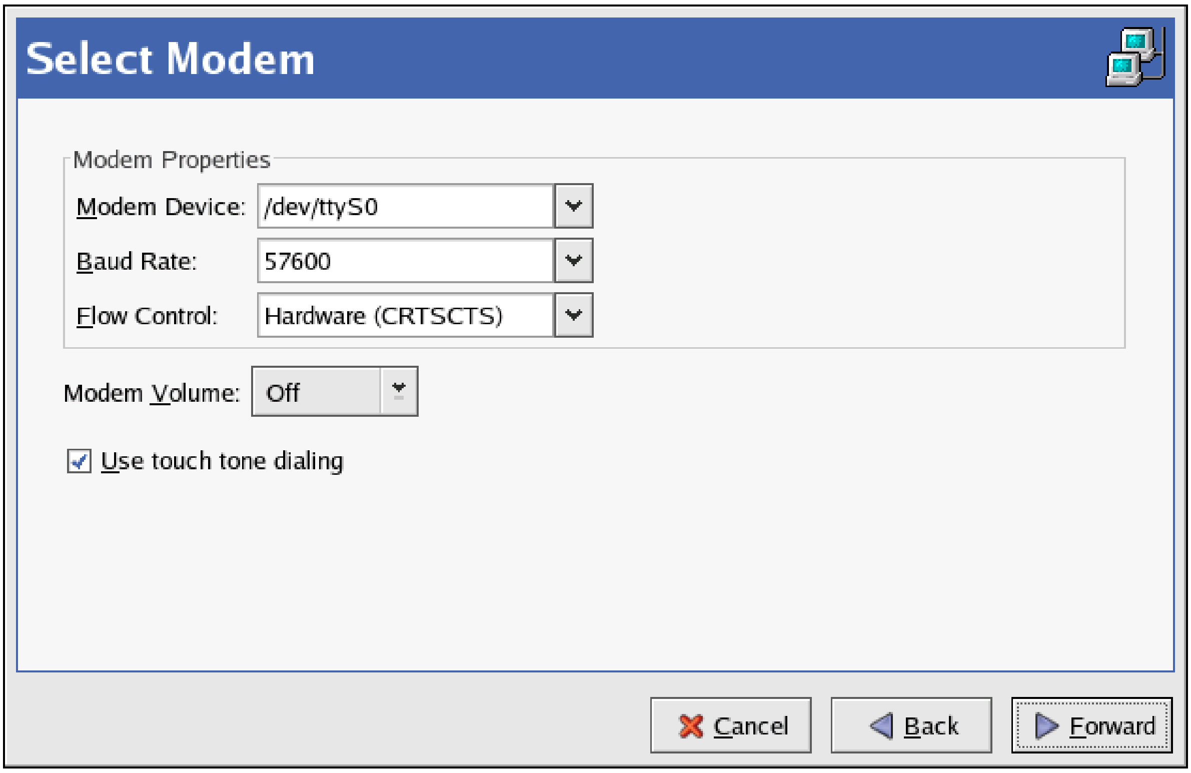

Select Modem

Modem Properties

Modem Device: /dev/ttyS0

Baud Rate: 57600

Flow Control: Hardware (CRTSCTS)

Modem Volume: Off

☑ Use touch tone dialing

그림 16.6. Modem Settings

6. Configure the modem device, baud rate, flow control, and modem volume. If you do not know these values, accept the defaults if the modem was probed successfully. If you do not have touch tone dialing, uncheck the corresponding checkbox. Click Forward.

7. If your ISP is in the pre-configured list, select it. Otherwise, enter the required information about your ISP account. If you do not know these values, contact your ISP. Click Forward.

8. On the IP Settings page, select whether to obtain an IP address automatically or whether to set one statically. Click Forward when finished.

9. On the Create Dialup Connection page, click Apply.

After configuring the modem device, it appears in the device list with the type Modem as shown in 그림 16.7. "Modem Device" .

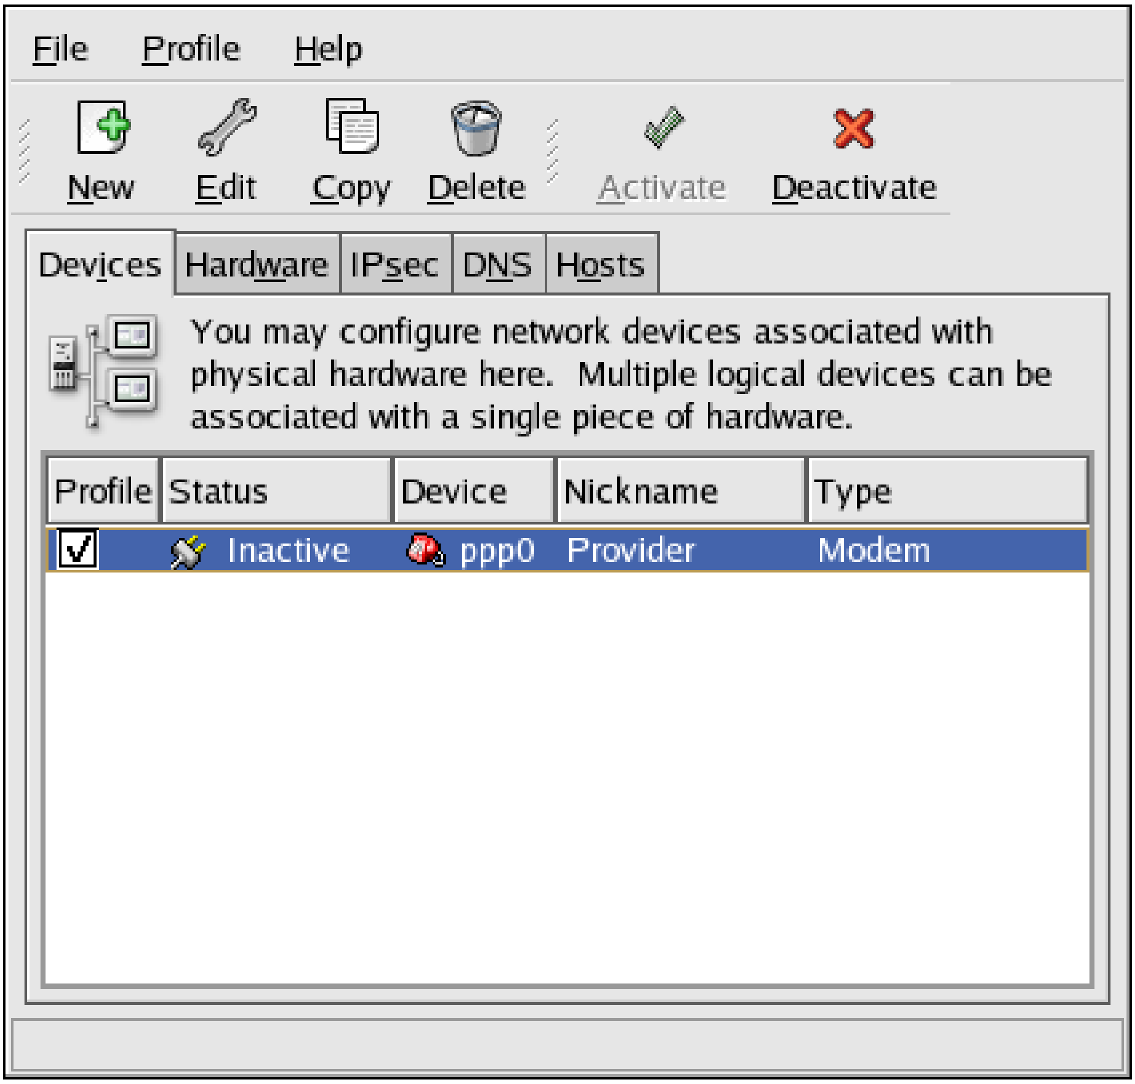

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	ppp0	Provider	Modem

그림 16.7. Modem Device

Be sure to select File > Save to save the changes.

After adding the modem device, you can edit its configuration by selecting the device from the device list and clicking Edit. For example, when the device is added, it is configured not to start at boot time by default. Edit its configuration to modify this setting. Compression, PPP options, login name, password, and more can also be changed.

When the device is added, it is not activated immediately, as seen by its Inactive status. To activate the device, select it from the device list, and click the Activate button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

16.5. xDSL 연결 설정하기

DSL stands for Digital Subscriber Lines. There are different types of DSL such as ADSL, IDSL, and SDSL. The Network Administration Tool uses the term xDSL to mean all types of DSL connections.

일부 DSL 제공 업체들은 이더넷 카드를 사용하여 DHCP를 통하여 IP 주소를 얻도록 시스템을 설정하도록 요구할 것입니다. 다른 DSL 제공 업체들은 이더넷 카드를 사용하여 PPPoE (Point-to-

Point Protocol over Ethernet) 연결을 설정하도록 요구할 수도 있습니다. 해당 DSL 제공 업체에 문의하셔서 어떤 방법을 사용할지 알아보십시오.

If you are required to use DHCP, refer to 16.2절. "이더넷 연결 설정하기" to configure your Ethernet card.

PPPoE를 사용하셔야 한다면 다음의 단계를 따르십시오:

1. Click the Devices tab.

2. Click the New button.

3. Select xDSL connection from the Device Type list, and click Forward as shown in 그림 16.8. "Select Device Type" .

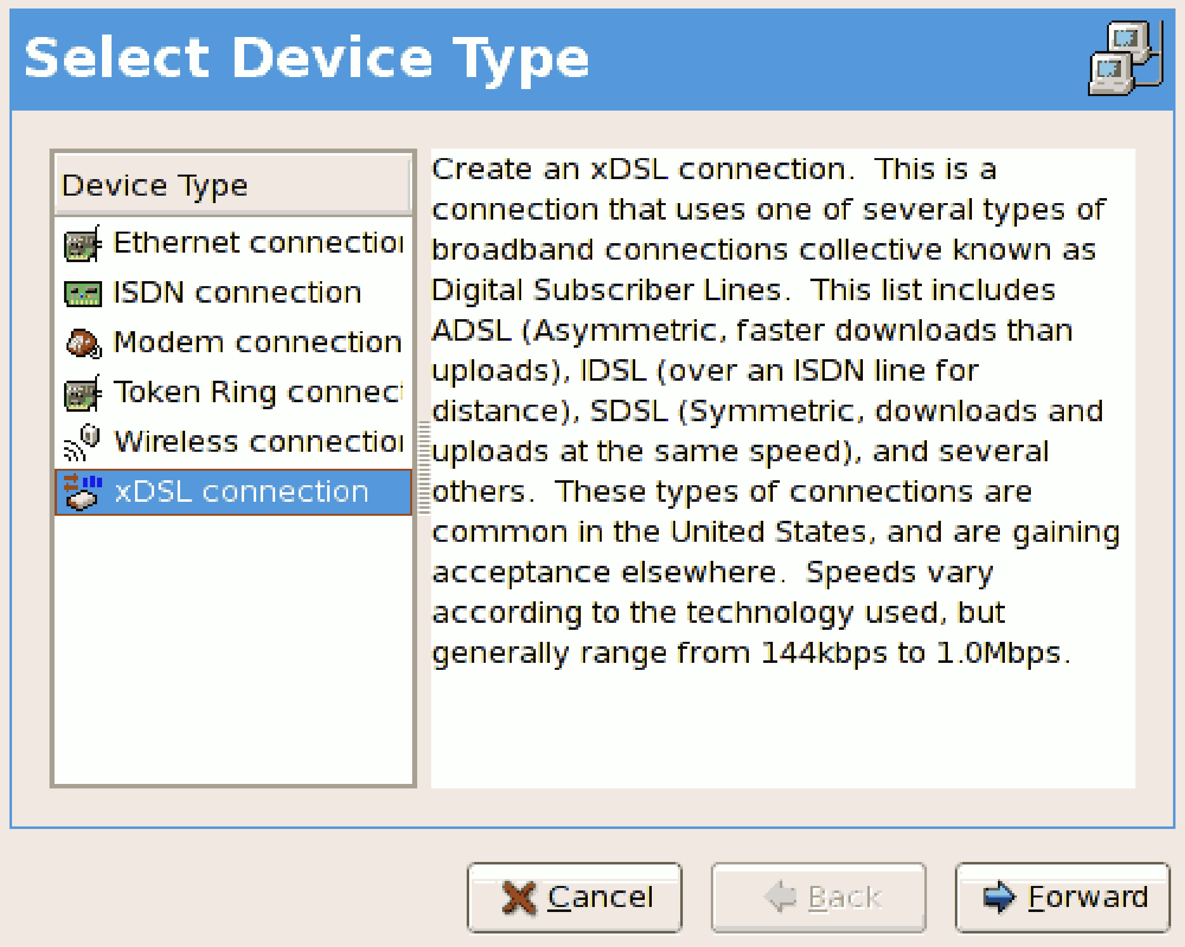

Create an xDSL connection. This is a connection that uses one of several types of broadband connections collective known as Digital Subscriber Lines. This list includes ADSL (Asymmetric, faster downloads than uploads), IDSL (over an ISDN line for distance), SDSL (Symmetric, downloads and uploads at the same speed), and several others. These types of connections are common in the United States, and are gaining acceptance elsewhere. Speeds vary according to the technology used, but generally range from 144kbps to 1.0Mbps.

그림 16.8. Select Device Type

4. If your Ethernet card is in the hardware list, select the Ethernet Device from the pulldown menu from the page shown in 그림 16.9. "xDSL Settings" . Otherwise, the Select Ethernet Adapter window appears.

알림

The installation program detects supported Ethernet devices and prompts you to configure them. If you configured any Ethernet devices during the installation, they are displayed in the hardware list on the Hardware tab.

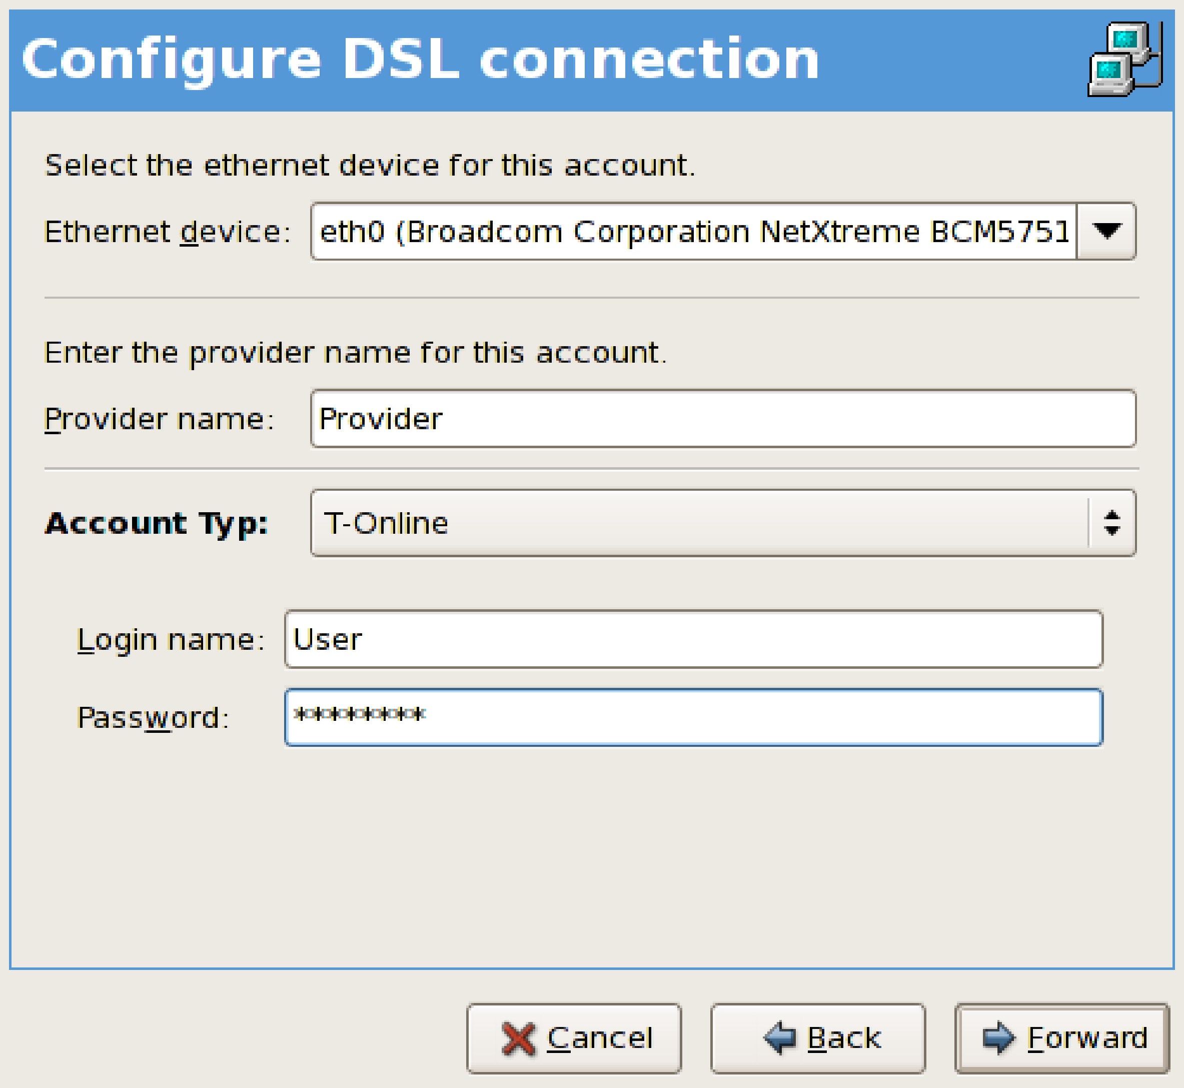

Configure DSL connection

Select the ethernet device for this account.

Ethernet device: eth0 (Broadcom Corporation NetXtreme BCM5751

Enter the provider name for this account.

Provider name: Provider

Account Typ: T-Online

Login name: User

Password: **********

그림 16.9. xDSL Settings

5. Enter the Provider Name, Login Name, and Password. If you are not setting up a T-Online account, select Normal from the Account Type pulldown menu.

   If you are setting up a T-Online account, select T-Online from the Account Type pulldown menu and enter any values in the Login name and Password field. You can further configure your T-Online account settings once the DSL connection has been fully configured (refer to T-Online 계정 설정).

6. 앞으로 버튼을 클릭하여 DSL 연결 생성 메뉴로 이동합니다. 셋팅을 확인한 후 적용 버튼을 클릭하여 설정을 완료합니다.

7. After configuring the DSL connection, it appears in the device list as shown in 그림 16.10. “xDSL Device” .

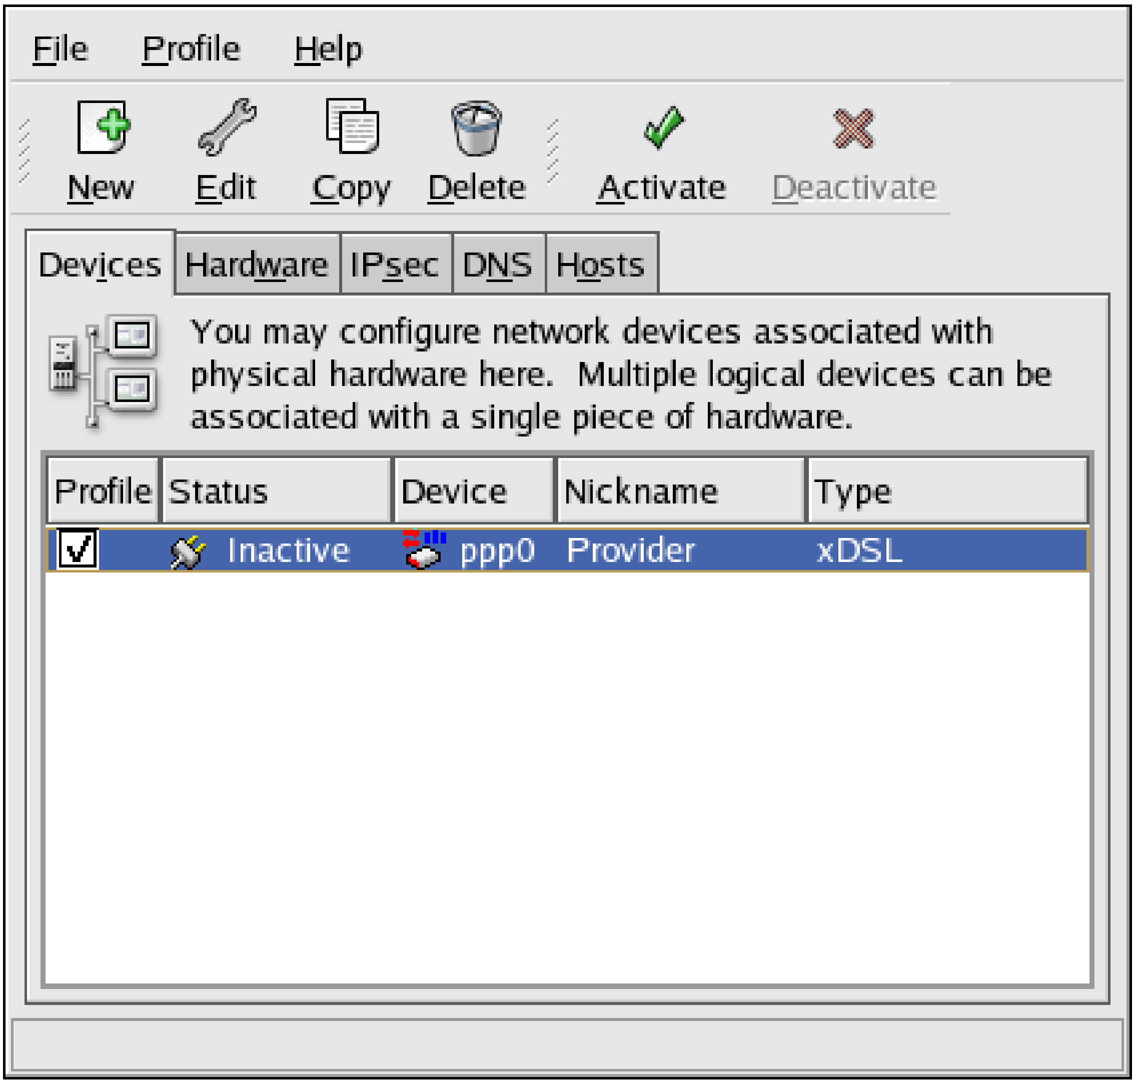

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	ppp0	Provider	xDSL

그림 16.10. xDSL Device

8. After adding the xDSL connection, you can edit its configuration by selecting the device from the device list and clicking Edit.

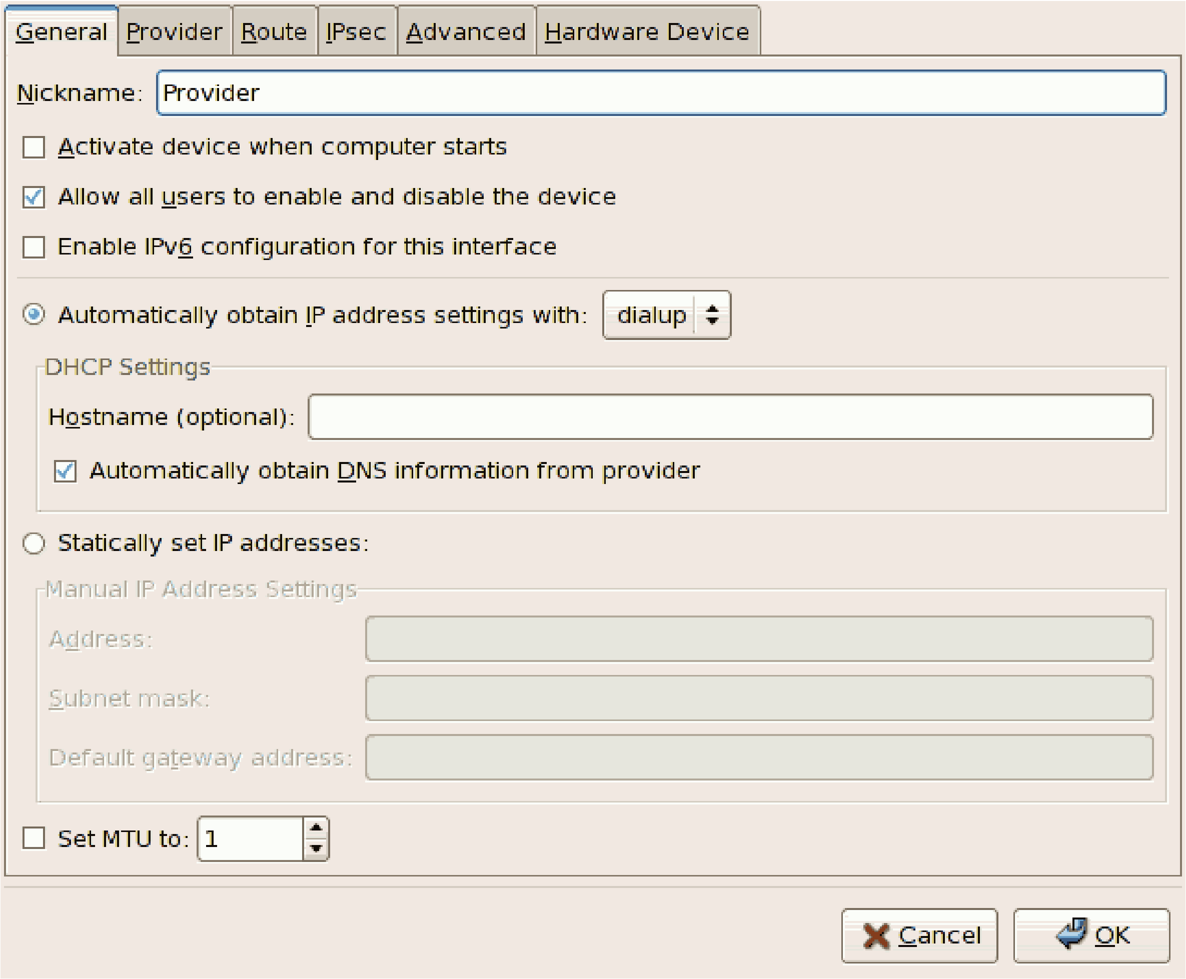

Nickname: Provider

□ Activate device when computer starts

☑ Allow all users to enable and disable the device

□ Enable IPv6 configuration for this interface

● Automatically obtain IP address settings with: dialup

DHCP Settings

Hostname (optional):

☑ Automatically obtain DNS information from provider

○ Statically set IP addresses:

Manual IP Address Settings

Address:

Subnet mask:

Default gateway address:

□ Set MTU to: 1

그림 16.11. xDSL Configuration

예를 들어 장치가 추가되었을 때 그 장치는 시스템 부팅시 시작하지 않도록 기본 설정되었다고 가정합니다. 설정을 편집하여 이러한 셋팅을 수정하실 수 있습니다. 설정을 완료하면 확인 버튼을 클릭합니다.

9. Once you are satisfied with your xDSL connection settings, select File > Save to save the changes.

T-Online 계정 설정

T-Online 계정을 설정하실 경우 다음의 단계를 따르십시오:

1. Select the device from the device list and click Edit.

2. Select the Provider tab from the xDSL Configuration menu as shown in 그림 16.12. “xDSL Configuration - Provider Tab” .

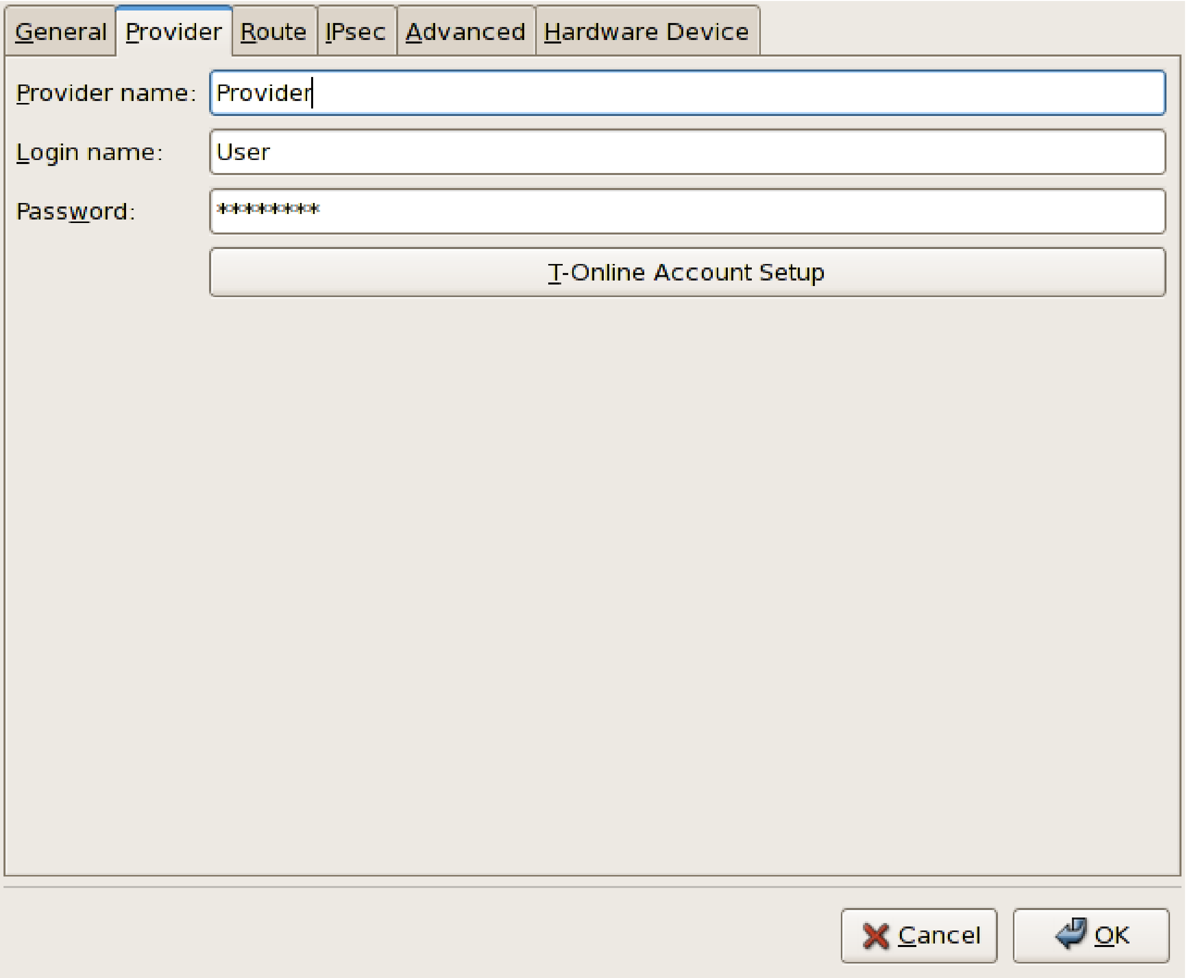

Provider name: Provider

Login name: User

Password: *********

그림 16.12. xDSL Configuration - Provider Tab

3. Click the T-Online Account Setup button. This will open the Account Setup window for your T-Online account as shown in 그림 16.13. “Account Setup” .

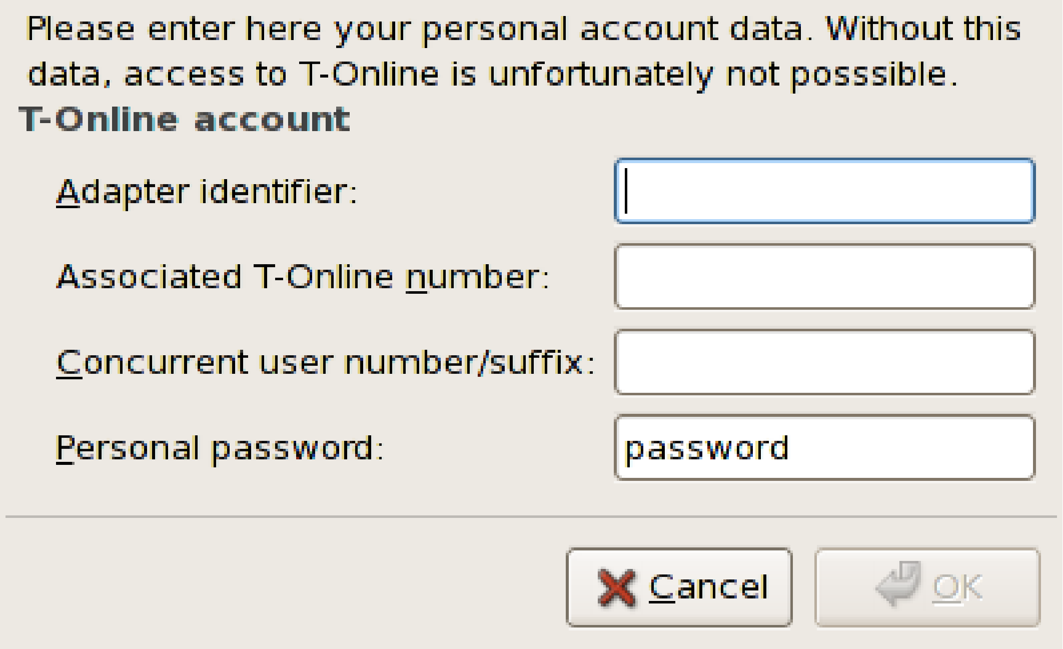

Please enter here your personal account data. Without this data, access to T-Online is unfortunately not posssible.

T-Online account

Adapter identifier:

Associated T-Online number:

Concurrent user number/suffix:

Personal password: password

그림 16.13. Account Setup

4. 아답터 식별자, 관련 T-Online 번호, 동시 사용자 번호/접미부, 개인 암호를 입력합니다. 설정을 완료하면 확인 버튼을 클릭하여 계정 설정 창을 닫습니다.

5. On the xDSL Configuration window, click OK. Be sure to select File > Save from the Network Administration Tool to save the changes.

When the device is added, it is not activated immediately, as seen by its Inactive status. To activate the device, select it from the device list, and click the Activate button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

16.6. 토큰 링 연결 설정하기

토큰 링 네트워크는 모든 컴퓨터가 원형으로 연결되어 있는 네트워크 입니다. 토큰 또는 특별 네트워크 패킷은 토큰 링 주위를 이동해 다니며 컴퓨터들 사이에서 서로 정보를 전송할 수 있게 합니다.

Tip

리눅스에서 토큰 링을 사용하는 방법에 대한 보다 자세한 내용은 http://www.linuxtr.net/에 있는 리눅스 토큰 링 프로젝트 웹사이트를 참조하시기 바랍니다.

토큰 링 연결을 추가하기 위해서는 다음의 단계를 따르십시오:

1. Click the Devices tab.

2. Click the New button on the toolbar.

3. Select Token Ring connection from the Device Type list and click Forward.

4. If you have already added the token ring card to the hardware list, select it from the Tokenring card list. Otherwise, select Other Tokenring Card to add the hardware device.

5. If you selected Other Tokenring Card, the Select Token Ring Adapter window as shown in 그림 16.14. "Token Ring Settings" appears. Select the manufacturer and model of the adapter. Select the device name. If this is the system's first token ring card, select tr0; if this is the second token ring card, select tr1 (and so on). The Network Administration Tool also allows the user to configure the resources for the adapter. Click Forward to continue.

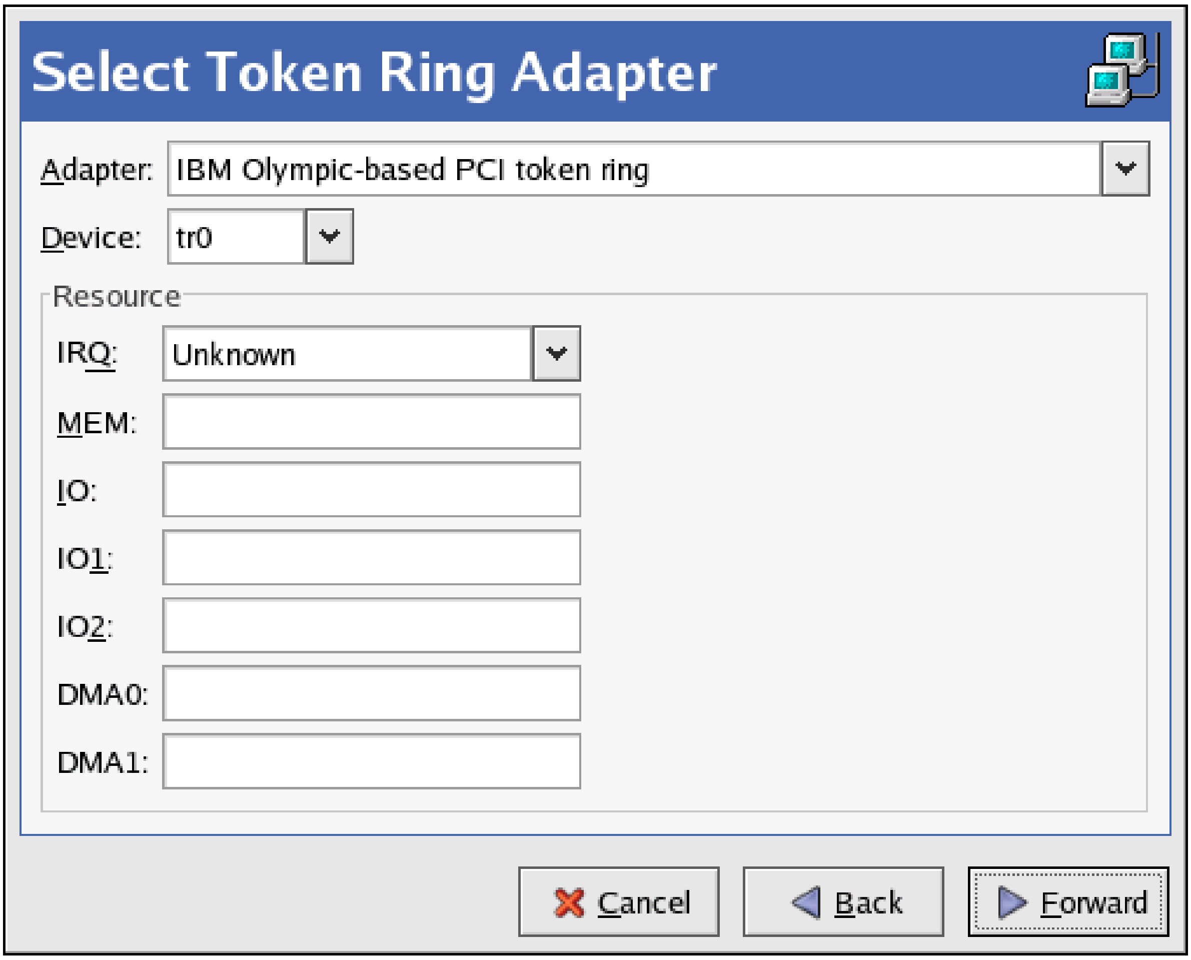

Adapter: IBM Olympic-based PCI token ring

Device: tr0

Resource

IRQ: Unknown

MEM:

IO:

IO1:

IO2:

DMA0:

DMA1:

그림 16.14. Token Ring Settings

6. On the Configure Network Settings page, choose between DHCP and static IP address. You may specify a hostname for the device. If the device receives a dynamic IP address each time the network is started, do not specify a hostname. Click Forward to continue.

7. Click Apply on the Create Tokenring Device page.

After configuring the token ring device, it appears in the device list as shown in 그림 16.15. “Token Ring Device” .

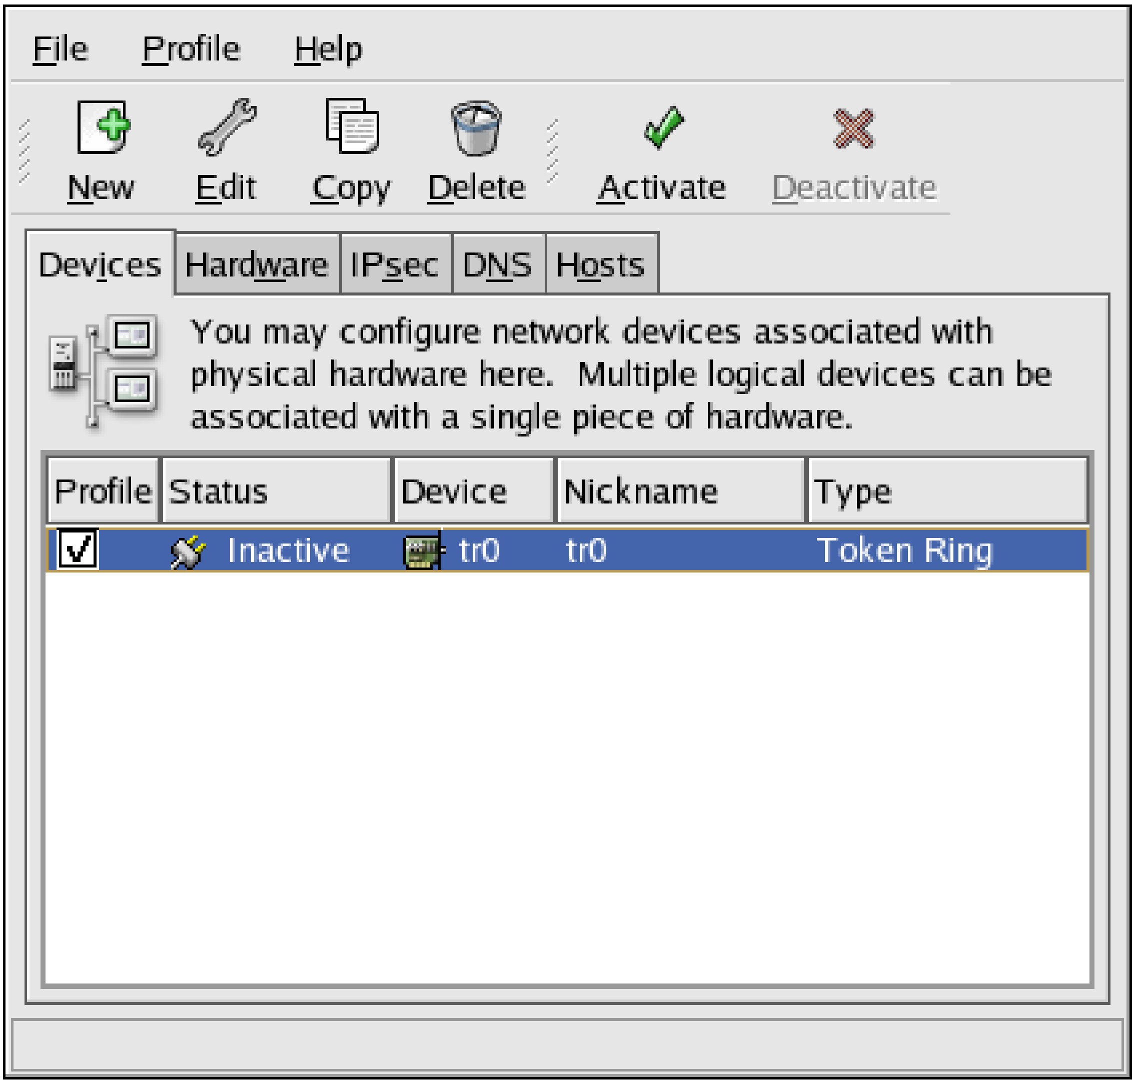

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	tr0	tr0	Token Ring

그림 16.15. Token Ring Device

Be sure to select File > Save to save the changes.

After adding the device, you can edit its configuration by selecting the device from the device list and clicking Edit. For example, you can configure whether the device is started at boot time.

When the device is added, it is not activated immediately, as seen by its Inactive status. To activate the device, select it from the device list, and click the Activate button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

16.7. 무선 연결 설정하기

무선 이더넷 장치들은 점점 더 대중적이 되고 있습니다. 무선 이더넷 장치 설정은 SSID와 무선 장치에 사용되는 키를 설정할 수 있다는 점을 제외하면 이더넷 설정과 유사합니다.

무선 이더넷 연결을 추가하기 위해서는, 다음의 단계를 따르십시오:

1. Click the Devices tab.

2. Click the New button on the toolbar.

3. Select Wireless connection from the Device Type list and click Forward.

4. If you have already added the wireless network interface card to the hardware list, select it from the Wireless card list. Otherwise, select Other Wireless Card to add the hardware device.

알림

The installation program usually detects supported wireless Ethernet devices and prompts you to configure them. If you configured them during the installation, they are displayed in the hardware list on the Hardware tab.

5. If you selected Other Wireless Card, the Select Ethernet Adapter window appears. Select the manufacturer and model of the Ethernet card and the device. If this is the first Ethernet card for the system, select eth0; if this is the second Ethernet card for the system, select eth1 (and so on). The Network Administration Tool also allows the user to configure the resources for the wireless network interface card. Click Forward to continue.

6. On the Configure Wireless Connection page as shown in 그림 16.16. "Wireless Settings", configure the settings for the wireless device.

Note: Open System and Shared Key Authentication

For the Authentication dropdown, note that wireless access points using WEP encryption have a choice between using open system and shared key authentication. Shared key authentication requires an exchange between the client and the access point during the association process that proves that the client has the correct WEP key. Open system authentication allows all wireless clients to connect. Counterintuitively, shared key authentication is less secure than open system, and thus is less widely deployed. It is therefore recommended to select Open System (open) as the authentication method when you do not know which method the access point requires. If connecting to the access point using open system fails, then try switching to shared key authentication.

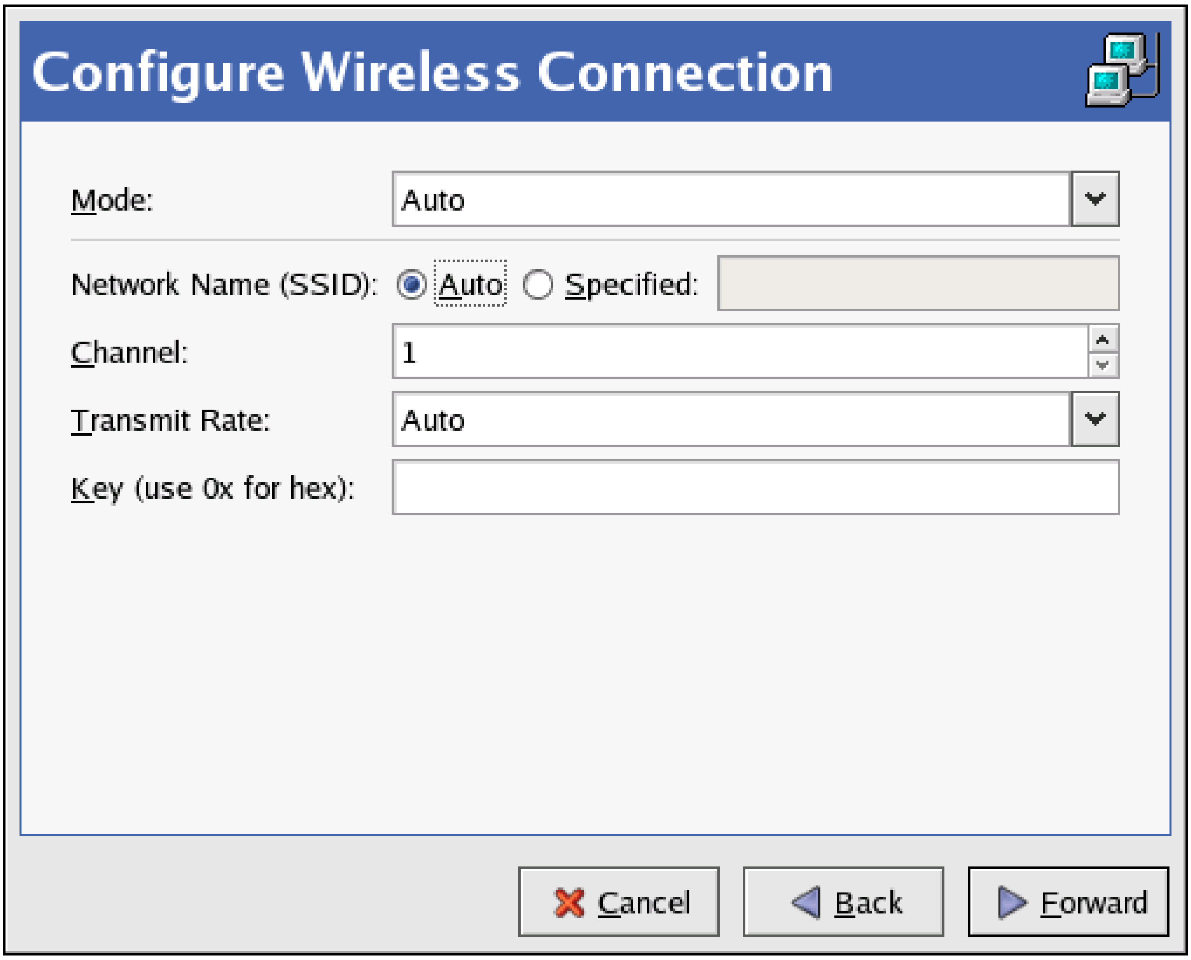

Configure Wireless Connection

Mode: Auto

Network Name (SSID): ● Auto ○ Specified:

Channel: 1

Transmit Rate: Auto

Key (use 0x for hex):

그림 16.16. Wireless Settings

7. On the Configure Network Settings page, choose between DHCP and static IP address. You may specify a hostname for the device. If the device receives a dynamic IP address each time the network is started, do not specify a hostname. Click Forward to continue.

8. Click Apply on the Create Wireless Device page.

After configuring the wireless device, it appears in the device list as shown in 그림 16.17. “Wireless Device” .

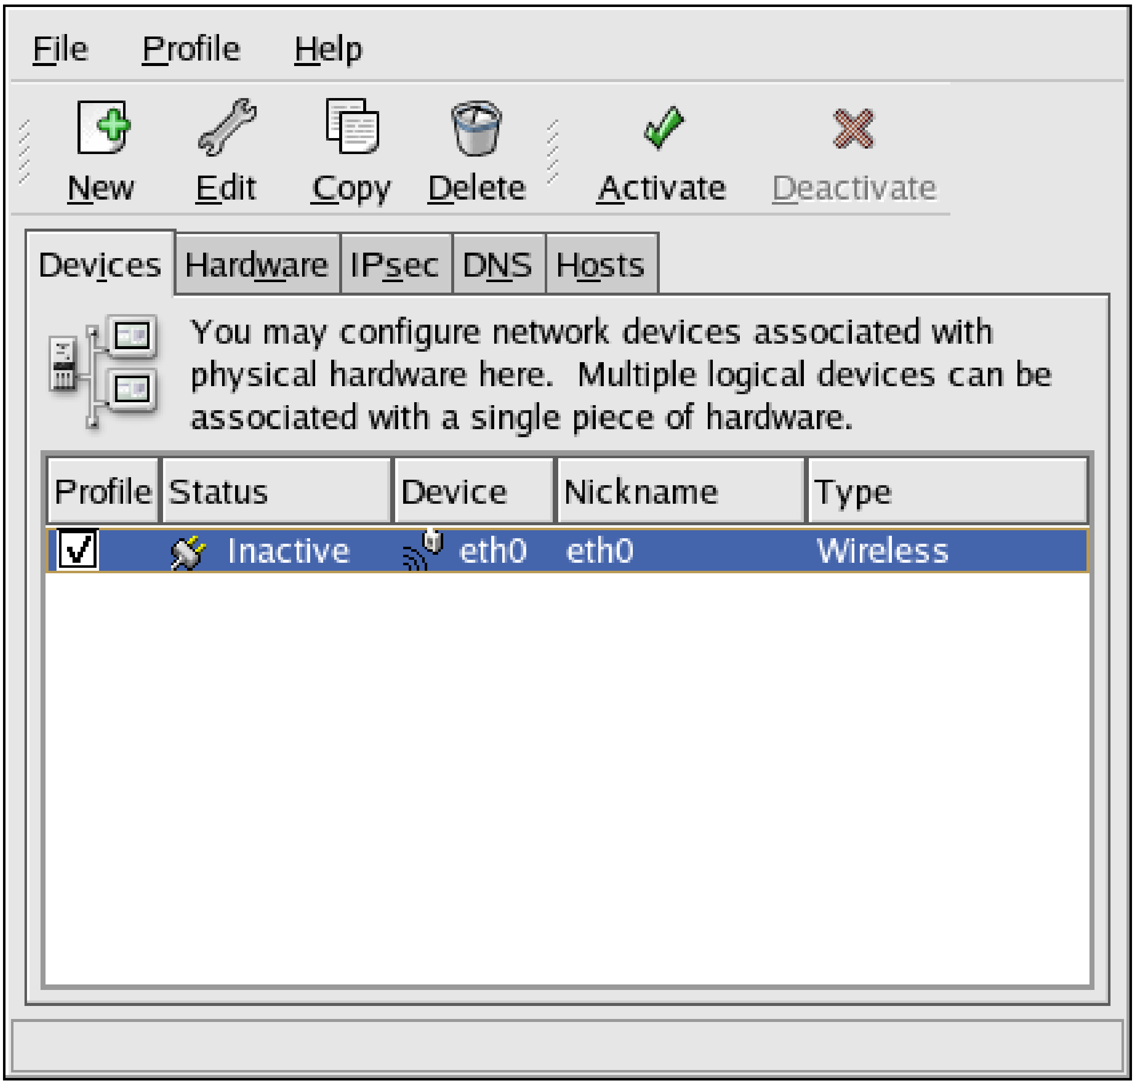

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Inactive	eth0	eth0	Wireless

그림 16.17. Wireless Device

Be sure to select File > Save to save the changes.

After adding the wireless device, you can edit its configuration by selecting the device from the device list and clicking Edit. For example, you can configure the device to activate at boot time.

When the device is added, it is not activated immediately, as seen by its Inactive status. To activate the device, select it from the device list, and click the Activate button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

16.8. DNS 셋팅 관리

The DNS tab allows you to configure the system's hostname, domain, name servers, and search domain. Name servers are used to look up other hosts on the network.

만일 DHCP 또는 PPPoE (또는 ISP)로부터 DNS 서버명이 검색되었다면, 1차, 2차, 3차 DNS 서버를 추가하지 마십시오.

만일 DHCP 또는 PPPoE (또는 ISP)로부터 동적으로 호스트명이 검색되었다면, 그 호스트명을 변경하지 마십시오.

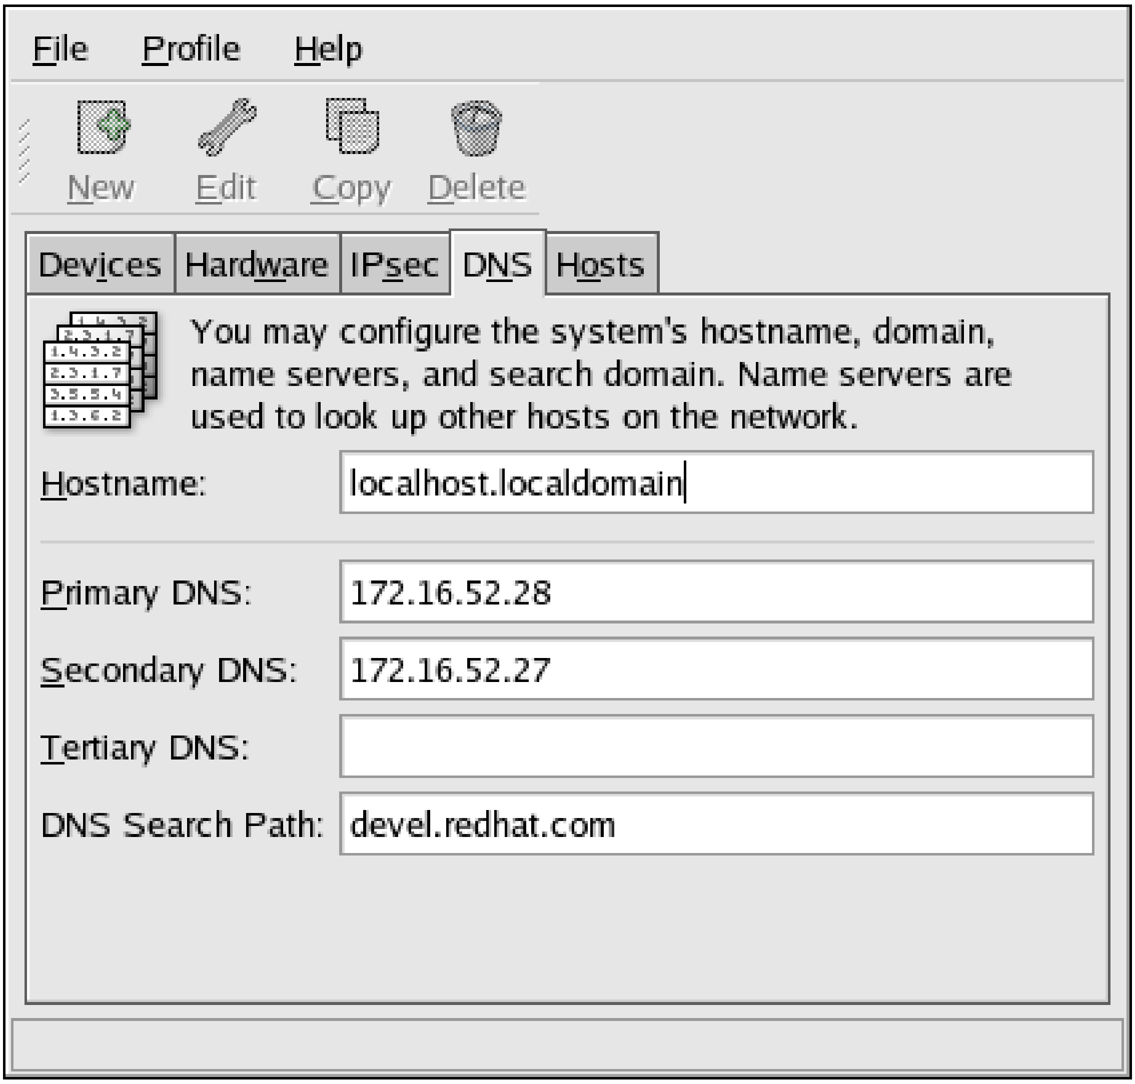

You may configure the system's hostname, domain, name servers, and search domain. Name servers are used to look up other hosts on the network.

Hostname: localhost.localdomain

Primary DNS: 172.16.52.28

Secondary DNS: 172.16.52.27

Tertiary DNS:

DNS Search Path: devel.redhat.com

그림 16.18. DNS Configuration

알림

네임 서버 섹션은 시스템이 네임 서버가 되도록 설정하지 않습니다. 대신, IP 주소를 호스트 명으로 변환하거나 호스트명을 IP 주소로 변환시 사용할 네임 서버를 설정합니다.

경고

호스트명을 변경하고 로컬 호스트에서 system-config-network 명령을 실행하면 다른 X11 응용 프로그램을 실행하실 수 없게 됩니다. 이러한 경우, 새로운 데스크탑 세션으로 다시 로그인하셔야 합니다.

16.9. 호스트 관리

The Hosts tab allows you to add, edit, or remove hosts from the /etc/hosts file. This file contains IP addresses and their corresponding hostnames.

When your system tries to resolve a hostname to an IP address or tries to determine the hostname for an IP address, it refers to the /etc/hosts file before using the name servers (if you are using the default Red Hat Enterprise Linux configuration). If the IP address is listed in the /etc/hosts file, the name servers are not used. If your network contains computers whose IP addresses are not listed in DNS, it is recommended that you add them to the /etc/hosts file.

To add an entry to the /etc/hosts file, go to the Hosts tab, click the New button on the toolbar, provide the requested information, and click OK. Select File > Save or press Ctrl+S to save the changes to the /etc/hosts file. The network or network services do not need to be restarted since the current version of the file is referred to each time an address is resolved.

경고

Do not remove the localhost entry. Even if the system does not have a network connection or have a network connection running constantly, some programs need to connect to the system via the localhost loopback interface.

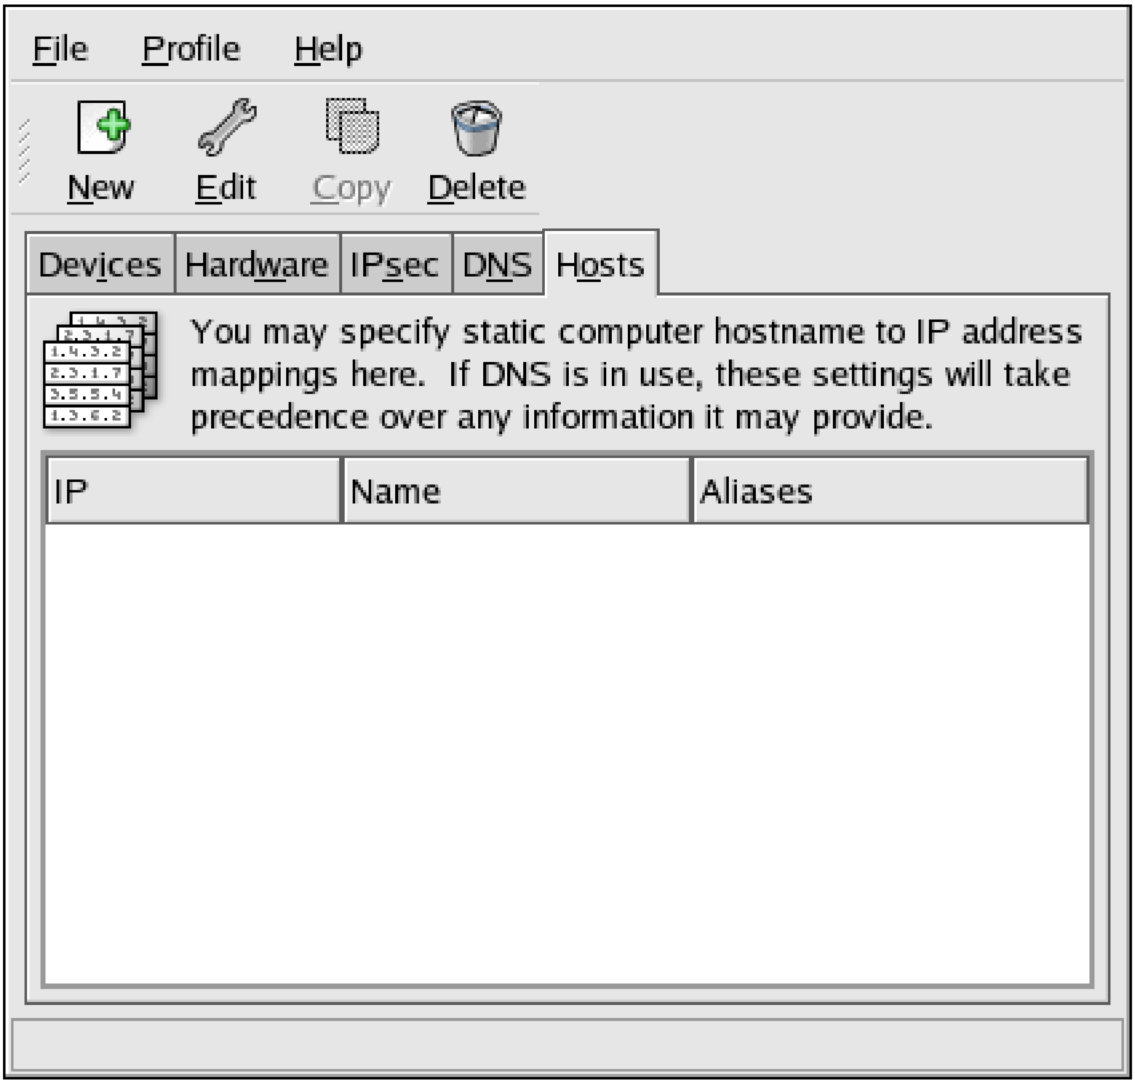

You may specify static computer hostname to IP address mappings here. If DNS is in use, these settings will take precedence over any information it may provide.

IP	Name	Aliases

그림 16.19. Hosts Configuration

Tip

To change lookup order, edit the /etc/host.conf file. The line order hosts, bind specifies that /etc/hosts takes precedence over the name servers. Changing the line to order bind, hosts configures the system to resolve hostnames and IP addresses using the name servers first. If the IP address cannot be resolved through the name servers, the system then looks for the IP address in the /etc/hosts file.

16.10. 프로파일 작업

한 개의 물리적 하드웨어 장치에 대하여 여러 개의 논리 네트워크 장치를 생성 가능합니다. 예를 들어 시스템에 한 개의 이더넷 카드 (eth0)를 가지고 계신 경우, 다른 별칭과 설정 옵션을 사용하여 eth0과 연관된 다양한 논리 네트워크 장치를 생성하실 수 있습니다.

Logical network devices are different from device aliases. Logical network devices associated with the same physical device must exist in different profiles and cannot be activated simultaneously. Device aliases are also associated with the same physical hardware device, but device aliases associated with the same physical hardware can be activated at the same time. Refer to 16.11절. “장치 별칭” for details about creating device aliases.

Profiles can be used to create multiple configuration sets for different networks. A configuration set can include logical devices as well as hosts and DNS settings. After configuring the profiles, you can use the Network Administration Tool to switch back and forth between them.

By default, there is one profile called Common. To create a new profile, select Profile > New from the pull-down menu, and enter a unique name for the profile.

이제 기본 창 아래쪽에 나타난 상태바에 표시된 새로운 프로파일을 수정하셔야 합니다.

Click on an existing device already in the list and click the Copy button to copy the existing device to a logical network device. If you use the New button, a network alias is created, which is incorrect. To change the properties of the logical device, select it from the list and click Edit. For example, the Nickname can be changed to a more descriptive name, such as eth0_office, so that it can be recognized more easily. Once you have finished editing your new profile, make sure to save it by clicking Save from the File menu. If you forget to save after creating a profile, that profile will be lost.

In the list of devices, there is a column of checkboxes labeled Profile. For each profile, you can check or uncheck devices. Only the checked devices are included for the currently selected profile. For example, if you create a logical device named eth0_office in a profile called Office and want to activate the logical device if the profile is selected, uncheck the eth0 device and check the eth0_office device.

For example, 그림 16.20. “사무용 (Office) 프로파일” shows a profile called Office with the logical device eth0_office. It is configured to activate the first Ethernet card using DHCP.

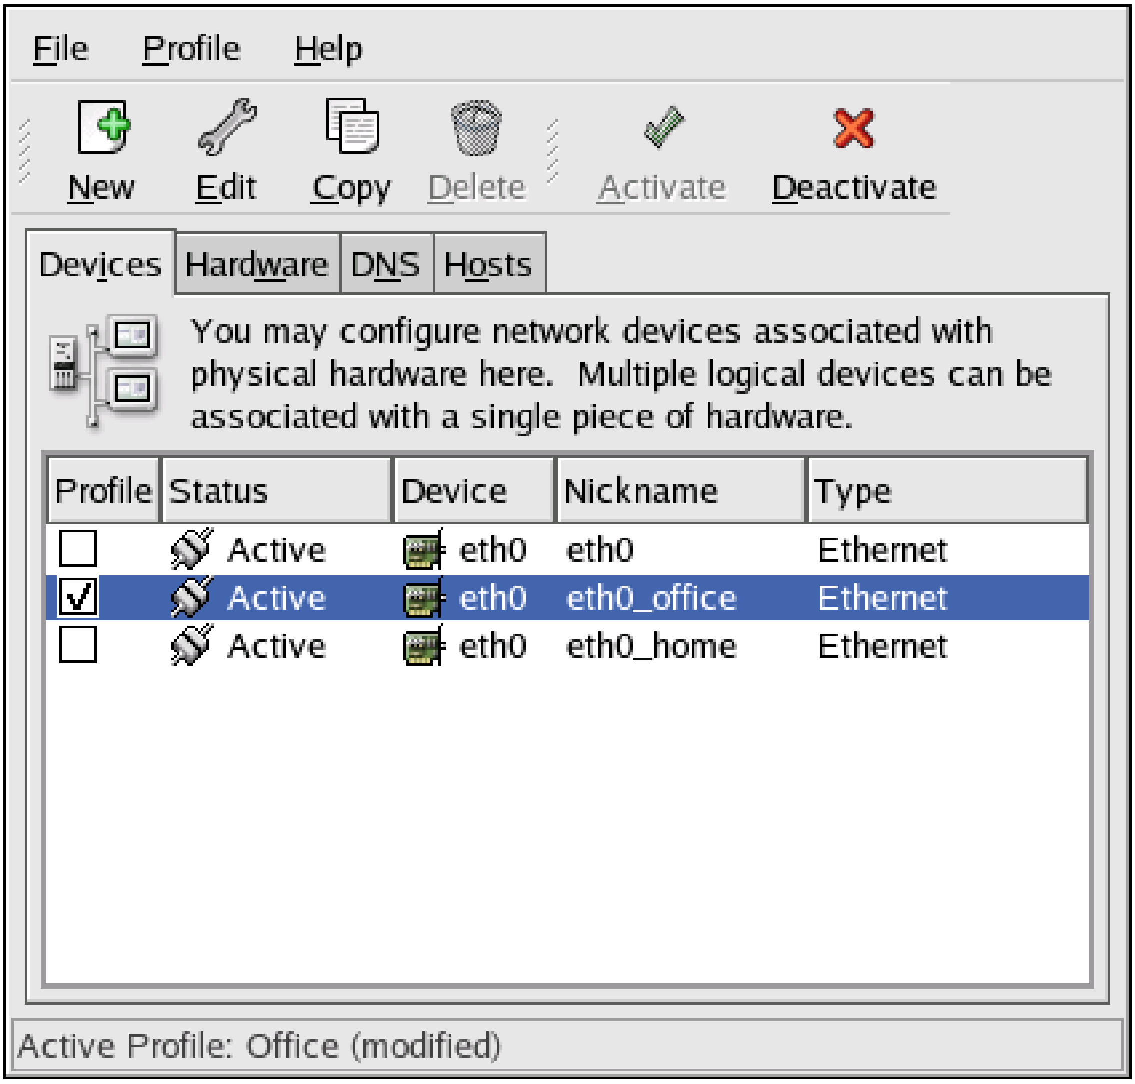

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
□	Active	eth0	eth0	Ethernet
☑	Active	eth0	eth0_office	Ethernet
□	Active	eth0	eth0_home	Ethernet

Active Profile: Office (modified)

그림 16.20. 사무용 (Office) 프로파일

Notice that the Home profile as shown in 그림 16.21. "홈 (Home) 프로파일" activates the eth0_home logical device, which is associated with eth0.

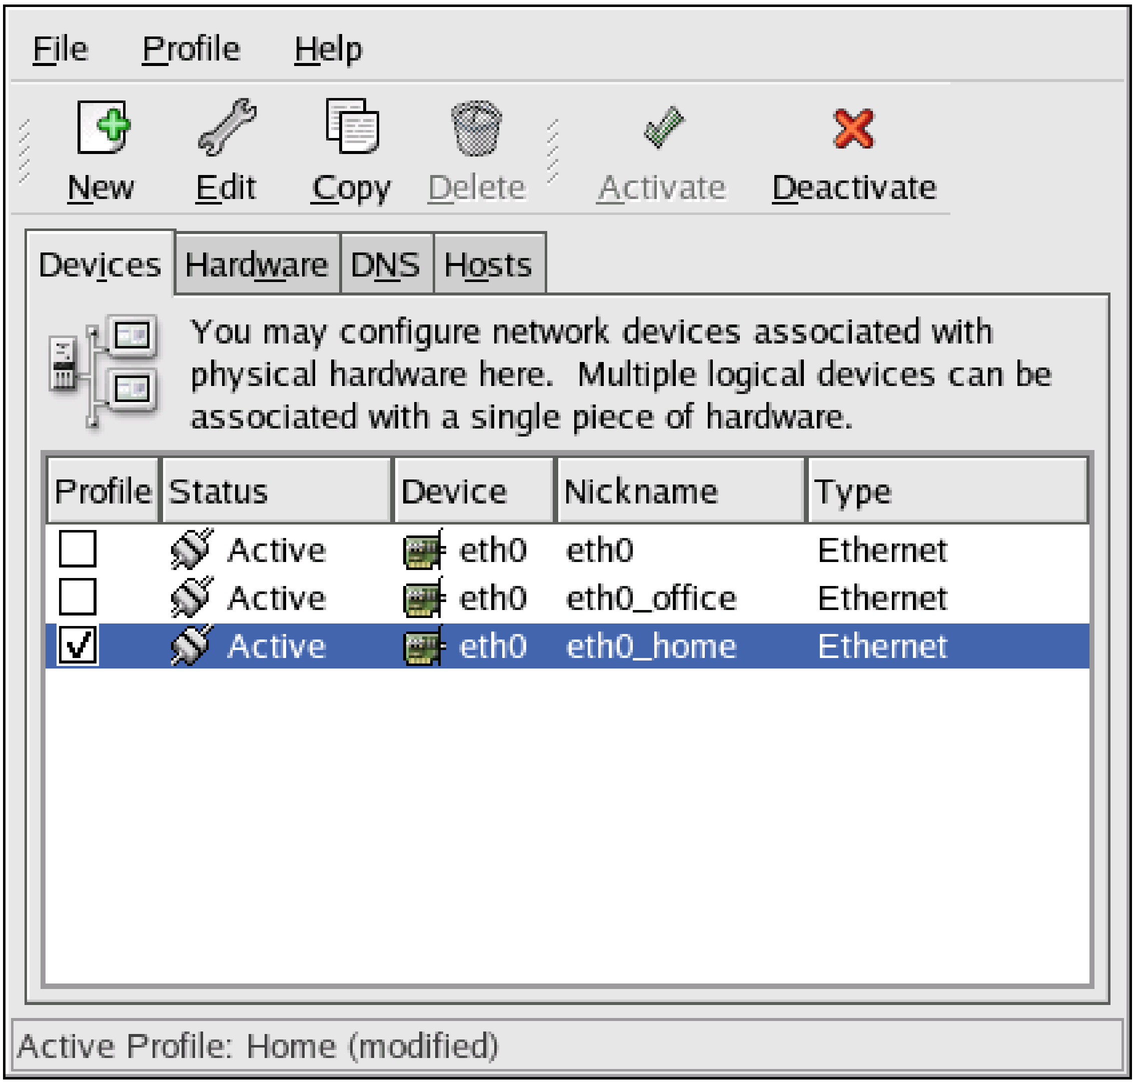

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
□	Active	eth0	eth0	Ethernet
□	Active	eth0	eth0_office	Ethernet
☑	Active	eth0	eth0_home	Ethernet

Active Profile: Home (modified)

그림 16.21. 홈 (Home) 프로파일

You can also configure eth0 to activate in the Office profile only and to activate a PPP (modem) device in the Home profile only. Another example is to have the Common profile activate eth0 and an Away profile activate a PPP device for use while traveling.

To activate a profile at boot time, modify the boot loader configuration file to include the netprofile=<profilename> option. For example, if the system uses GRUB as the boot loader and /boot/grub/grub.conf contains:

```
title Red Hat Enterprise Linux (2.6.9-5.EL)
          root (hd0,0)
  kernel /vmlinuz-2.6.9-5.EL ro root=/dev/VolGroup00/LogVol00 rhgb quiet
  initrd /initrd-2.6.9-5.EL.img
```

Modify it to the following (where <profilename> is the name of the profile to be activated at boot time):

```
title Red Hat Enterprise Linux (2.6.9-5.EL)
          root (hd0,0)
  kernel /vmlinuz-2.6.9-5.EL ro root=/dev/VolGroup00/LogVol00 \
   netprofile=<profilename>  \    rhgb quiet
```

```
    initrd /initrd-2.6.9-5.EL.img
```

To switch profiles after the system has booted, go to Applications (the main menu on the panel) > System Tools > Network Device Control (or type the command system-control-network) to select a profile and activate it. The activate profile section only appears in the Network Device Control interface if more than the default Common interface exists.

Alternatively, execute the following command to enable a profile (replace <profilename> with the name of the profile):

```
system-config-network-cmd --profile <profilename> --activate
```

16.11. 장치 별칭

장치 별칭이란 동일한 물리적 하드웨어에 연관되었지만, 동시에 다른 IP 주소를 갖도록 활성화 가능한 가상 장치를 말합니다. 일반적으로 장치 별칭은 장치명 다음에 콜론과 숫자가 오는 형식으로 나타납니다 (예, eth0:1). 한 개의 네트워크 카드를 갖는 시스템에 대해 여러 개의 IP 주소를 원하시는 경우, 장치 별칭이 유용합니다.

After configuring the Ethernet device —such as eth0 —to use a static IP address (DHCP does not work with aliases), go to the Devices tab and click New. Select the Ethernet card to configure with an alias, set the static IP address for the alias, and click Apply to create it. Since a device already exists for the Ethernet card, the one just created is the alias, such as eth0:1.

> **경고**
>
> 만일 이더넷 장치가 별칭을 갖도록 설정하시면, 이더넷 장치와 별칭은 DHCP를 사용할 수 없습니다. 따라서 직접 IP 주소를 설정하셔야 합니다.

그림 16.22. "네트워크 장치 별칭 예시" shows an example of one alias for the eth0 device. Notice the eth0:1 device — the first alias for eth0. The second alias for eth0 would have the device name eth0:2, and so on. To modify the settings for the device alias, such as whether to activate it at boot time and the alias number, select it from the list and click the Edit button.

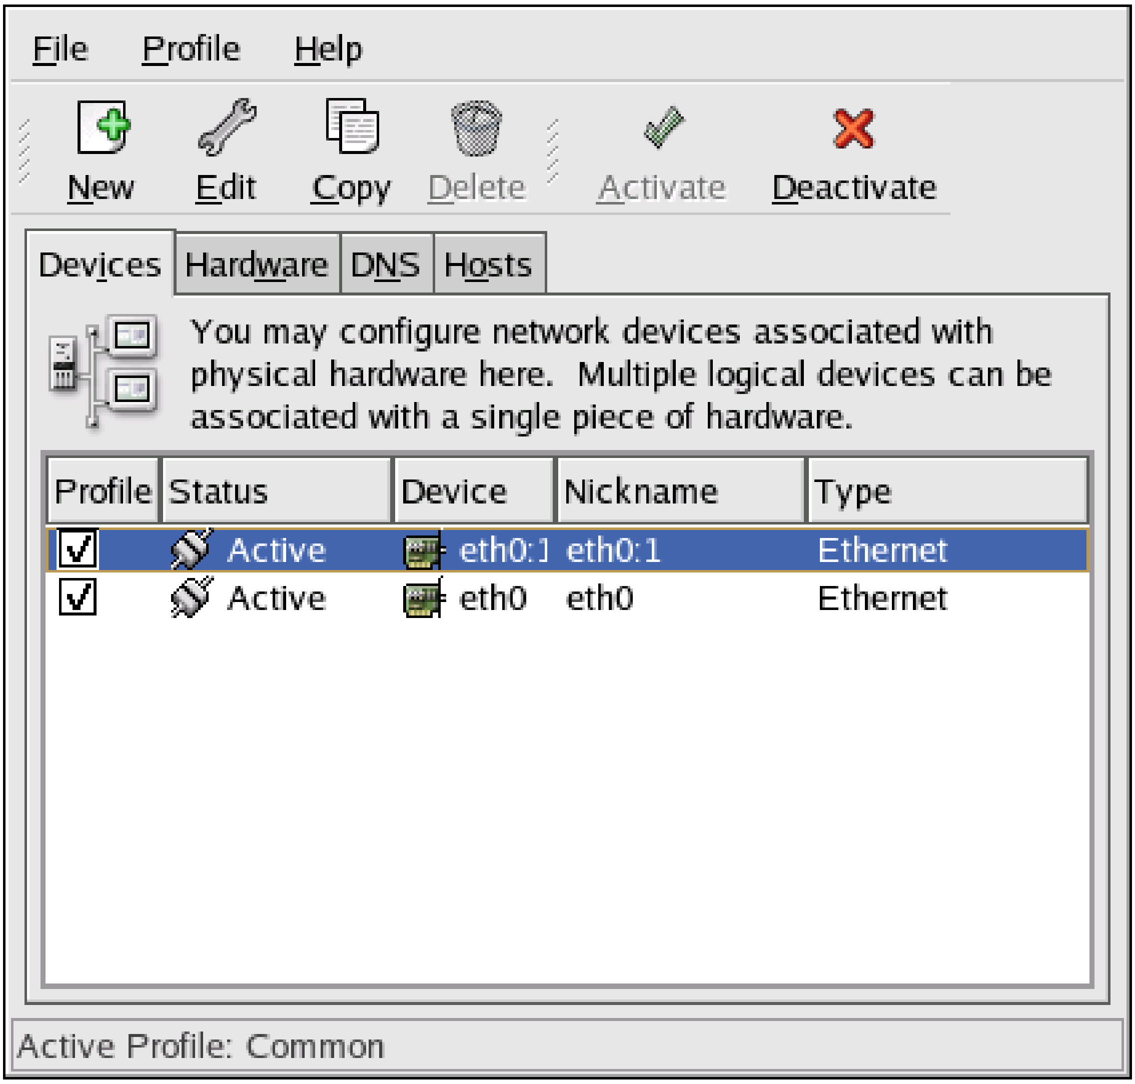

You may configure network devices associated with physical hardware here. Multiple logical devices can be associated with a single piece of hardware.

Profile	Status	Device	Nickname	Type
☑	Active	eth0:1	eth0:1	Ethernet
☑	Active	eth0	eth0	Ethernet

Active Profile: Common

그림 16.22. 네트워크 장치 별칭 예시

Select the alias and click the Activate button to activate the alias. If you have configured multiple profiles, select which profiles in which to include it.

To verify that the alias has been activated, use the command /sbin/ifconfig. The output should show the device and the device alias with different IP addresses:

```
eth0      Link encap:Ethernet
 HWaddr 00:A0:CC:60:B7:G4
 inet addr:192.168.100.5  Bcast:192.168.100.255  Mask:255.255.255.0
 UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
 RX packets:161930 errors:1 dropped:0 overruns:0 frame:0
 TX packets:244570 errors:0 dropped:0 overruns:0 carrier:0
 collisions:475 txqueuelen:100
 RX bytes:55075551 (52.5 Mb)  TX bytes:178108895 (169.8 Mb)
 Interrupt:10 Base address:0x9000  eth0:1    Link encap:Ethernet  HWaddr 00:A0:CC:60:B7:G4
 inet addr:192.168.100.42  Bcast:192.168.100.255  Mask:255.255.255.0
 UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
 Interrupt:10 Base address:0x9000  lo
 Link encap:Local Loopback
 inet addr:127.0.0.1  Mask:255.0.0.0
 UP LOOPBACK RUNNING  MTU:16436  Metric:1
 RX packets:5998 errors:0 dropped:0 overruns:0 frame:0
 TX packets:5998 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:1627579 (1.5 Mb)  TX bytes:1627579 (1.5 Mb)
```

16.12. 네트워크 설정 저장과 복구

The command line version of Network Administration Tool can be used to save the system's network configuration to a file. This file can then be used to restore the network settings to a Red Hat Enterprise Linux system.

이 기능은 업그레이드나 재설치 이전에 설정을 저장하거나 다른 Red Hat Enterprise Linux 시스템으로 설정을 복사하기 위해 자동 백업 스크립트의 일부로 사용될 수 있습니다.

To save, or export, the network configuration of a system to the file /tmp/network-config, execute the following command as root:

```
system-config-network-cmd -e > /tmp/network-config
```

앞의 명령을 사용하여 만들어진 파일에서 네트워크 설정을 복구하거나 가져오기하기 위해서는 루트로 로그인하여 다음 명령을 실행하시면 됩니다:

```
system-config-network-cmd -i -c -f /tmp/network-config
```

-i 옵션은 데이터를 가져오기(import) 한다는 것을 의미하며 -c 옵션은 가져오기 이전에 기존 설정을 지우고(clear), -f 옵션은 다음에 지정된 파일을 가져오기한다는 것을 의미합니다.

서비스로의 접근 통제

Maintaining security on your system is extremely important, and one approach for this task is to manage access to system services carefully. Your system may need to provide open access to particular services (for example, httpd if you are running a Web server). However, if you do not need to provide a service, you should turn it off to minimize your exposure to possible bug exploits.

There are several different methods for managing access to system services. Choose which method of management to use based on the service, your system's configuration, and your level of Linux expertise.

The easiest way to deny access to a service is to turn it off. Both the services managed by xinetd and the services in the /etc/rc.d/init.d hierarchy (also known as SysV services) can be configured to start or stop using three different applications:

Services Configuration Tool
: This is a graphical application that displays a description of each service, displays whether each service is started at boot time (for runlevels 3, 4, and 5), and allows services to be started, stopped, and restarted.

ntsysv
: This is a text-based application that allows you to configure which services are started at boot time for each runlevel. Non-xinetd services can not be started, stopped, or restarted using this program.

chkconfig
: This is a command line utility that allows you to turn services on and off for the different runlevels. Non-xinetd services can not be started, stopped, or restarted using this utility.

You may find that these tools are easier to use than the alternatives — editing the numerous symbolic links located in the directories below /etc/rc.d by hand or editing the xinetd configuration files in /etc/xinetd.d.

Another way to manage access to system services is by using iptables to configure an IP firewall. If you are a new Linux user, note that iptables may not be the best solution for you. Setting up iptables can be complicated, and is best tackled by experienced Linux system administrators.

On the other hand, the benefit of using iptables is flexibility. For example, if you need a customized solution which provides certain hosts access to certain services, iptables can provide it for you. Refer to 46.8.1절. "Netfilter and IPTables" and 46.8.3절. "Using IPTables" for more information about iptables.

Alternatively, if you are looking for a utility to set general access rules for your home machine, and/or if you are new to Linux, try the Security Level Configuration Tool (system-config-securitylevel), which allows you to select the security level for your system, similar to the Firewall Configuration screen in the installation program.

Refer to 46.8절. "Firewalls" for more information.

Important

When you allow access for new services, always remember that both the firewall and SELinux need to be configured as well. One of the most common mistakes committed when configuring a new service is neglecting to implement the necessary firewall configuration and SELinux policies to allow access for it. Refer to 46.8.2절. "Basic Firewall Configuration" for more information.

17.1. 런레벨 (runlevels)

Before you can configure access to services, you must understand Linux runlevels. A runlevel is a state, or mode, that is defined by the services listed in the directory /etc/rc.d/rc<x>.d, where <x> is the number of the runlevel.

사용 가능한 런레벨은 다음과 같습니다:

- 0 — 정지
- 1 — 단독-사용자 모드
- 2 — 사용안됨 (사용자-정의가능)
- 3 — 완전 다중-사용자 모드
- 4 — 사용안됨 (사용자-정의가능)
- 5 — (X-기반 로그인 화면을 사용한) 완전 다중-사용자 모드
- 6 — 재부팅

텍스트 로그인 화면을 사용하신 경우 런레벨 3으로 작동합니다. 만일 그래픽 로그인 화면을 선택하신 경우에는 런레벨 5에서 작동합니다.

The default runlevel can be changed by modifying the /etc/inittab file, which contains a line near the top of the file similar to the following:

```
id:5:initdefault:
```

Change the number in this line to the desired runlevel. The change does not take effect until you reboot the system.

17.2. TCP 래퍼 (Wrappers)

Many UNIX system administrators are accustomed to using TCP wrappers to manage access to certain network services. Any network services managed by xinetd (as well as any program with built-in support for libwrap) can use TCP wrappers to manage access. xinetd can use the /etc/hosts.allow and /etc/hosts.deny files to configure access to system services. As the names imply, hosts.allow contains a list of rules that allow clients to access the network services controlled by xinetd, and hosts.deny contains rules to deny access. The hosts.allow file takes precedence over the hosts.deny file. Permissions to grant or deny access can be based on individual IP address (or hostnames) or on a pattern of clients. Refer to hosts_access in section 5 of the man pages (man 5 hosts_access) for details.

17.2.1. xinetd

To control access to Internet services, use xinetd, which is a secure replacement for inetd. The xinetd daemon conserves system resources, provides access control and logging, and can be used to start special-purpose servers. xinetd can also be used to grant or deny access to particular hosts, provide service access at specific times, limit the rate of incoming connections, limit the load created by connections, and more.

xinetd runs constantly and listens on all ports for the services it manages. When a connection request arrives for one of its managed services, xinetd starts up the appropriate server for that service.

The configuration file for xinetd is /etc/xinetd.conf, but the file only contains a few defaults and an instruction to include the /etc/xinetd.d directory. To enable or disable an xinetd service, edit its configuration file in the /etc/xinetd.d directory. If the disable attribute is set to yes, the service is disabled. If the disable attribute is set to no, the service is enabled. You can edit any of the xinetd configuration files or change its enabled status using the Services Configuration Tool, ntsysv, or chkconfig. For a list of network services controlled by xinetd, review the contents of the /etc/xinetd.d directory with the command ls /etc/xinetd.d.

17.3. Services Configuration Tool

The Services Configuration Tool is a graphical application developed by Red Hat to configure which SysV services in the /etc/rc.d/init.d directory are started at boot time (for runlevels 3, 4, and 5) and which xinetd services are enabled. It also allows you to start, stop, and restart SysV services as well as reload xinetd.

To start the Services Configuration Tool from the desktop, go to the Applications (the main menu on the panel) > System Settings > Server Settings > Services or type the command system-config-services at a shell prompt (for example, in an XTerm or a GNOME terminal).

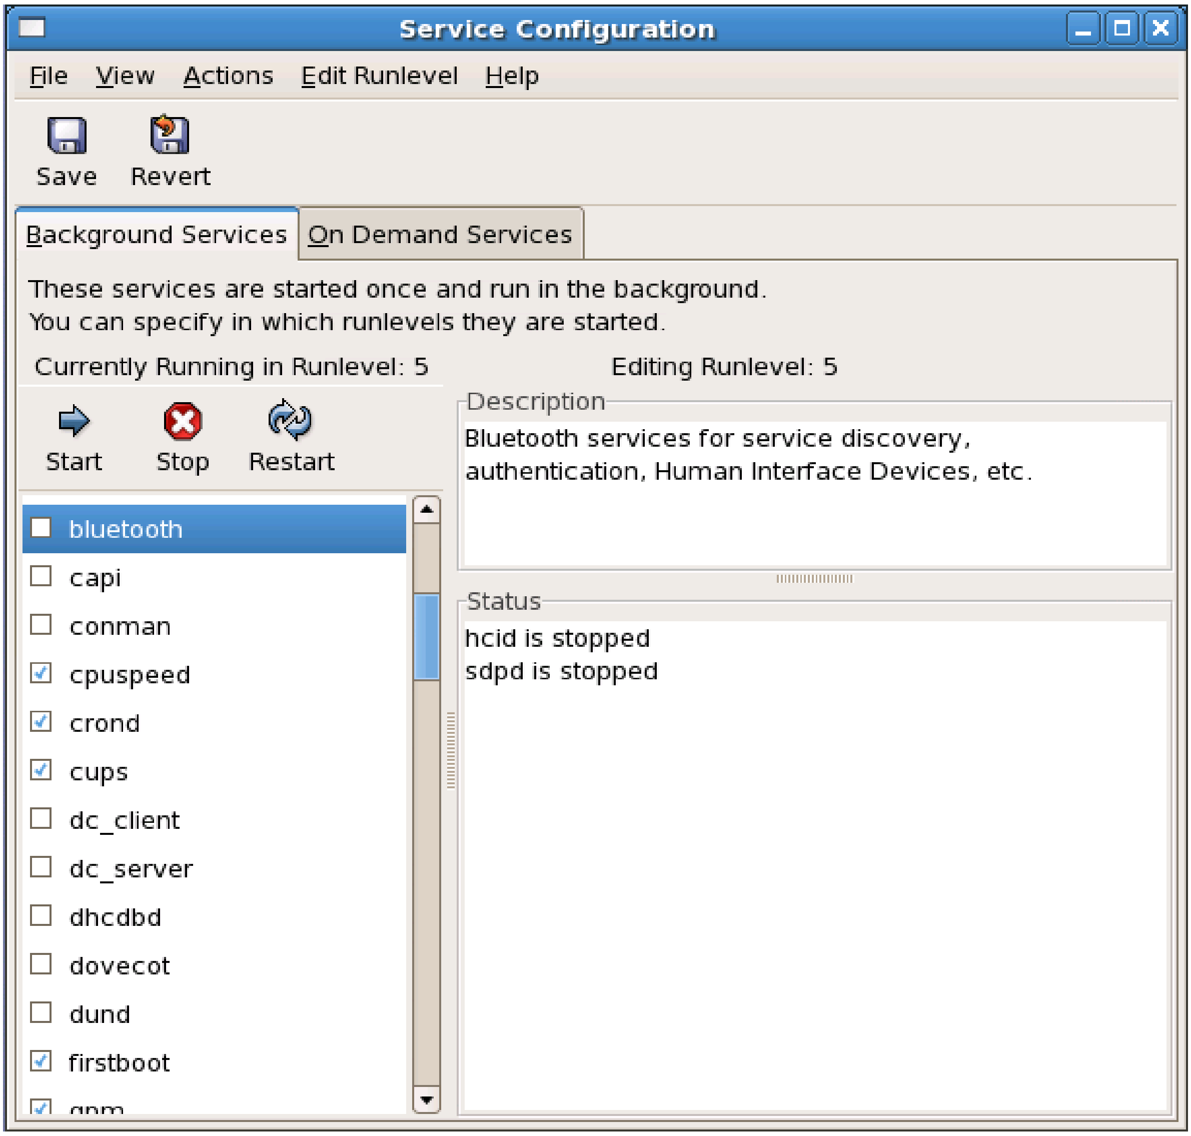

These services are started once and run in the background.
You can specify in which runlevels they are started.

Currently Running in Runlevel: 5

Editing Runlevel: 5

Bluetooth services for service discovery, authentication, Human Interface Devices, etc.

hcid is stopped
sdpd is stopped

- □ bluetooth
- □ capi
- □ conman
- ☑ cpuspeed
- ☑ crond
- ☑ cups
- □ dc_client
- □ dc_server
- □ dhcdbd
- □ dovecot
- □ dund
- ☑ firstboot
- ☑ gpm

그림 17.1. Services Configuration Tool

The Services Configuration Tool displays the current runlevel as well as the runlevel you are currently editing. To edit a different runlevel, select Edit Runlevel from the pulldown menu and select runlevel 3, 4, or 5. Refer to 17.1절. “런레벨 (runlevels)” for a description of runlevels.

The Services Configuration Tool lists the services from the /etc/rc.d/init.d directory as well as the services controlled by xinetd. Click on the name of the service from the list on the left-hand side of the application to display a brief description of that service as well as the status of the service. If the service is not an xinetd service, the status window shows whether the service is currently running. If the service is controlled by xinetd, the status window displays the phrase xinetd service.

To start, stop, or restart a service immediately, select the service from the list and click the appropriate button on the toolbar (or choose the action from the Actions pulldown menu). If the service is an xinetd service, the action buttons are disabled because they cannot be started or stopped individually.

If you enable/disable an xinetd service by checking or unchecking the checkbox next to the service name, you must select File > Save Changes from the pulldown menu (or the Save button above the tabs) to reload xinetd and immediately enable/disable the xinetd service that you changed. xinetd is also configured to remember the setting. You can enable/disable multiple xinetd services at a time and save the changes when you are finished.

For example, assume you check rsync to enable it in runlevel 3 and then save the changes. The rsync service is immediately enabled. The next time xinetd is started, rsync is still enabled.

> **Note**
>
> When you save changes to xinetd services, xinetd is reloaded, and the changes take place immediately. When you save changes to other services, the runlevel is reconfigured, but the changes do not take effect immediately.

To enable a non-xinetd service to start at boot time for the currently selected runlevel, check the box beside the name of the service in the list. After configuring the runlevel, apply the changes by selecting File > Save Changes from the pulldown menu. The runlevel configuration is changed, but the runlevel is not restarted; thus, the changes do not take place immediately.

For example, assume you are configuring runlevel 3. If you change the value for the httpd service from checked to unchecked and then select Save Changes, the runlevel 3 configuration changes so that httpd is not started at boot time. However, runlevel 3 is not reinitialized, so httpd is still running. Select one of following options at this point:

1. Stop the httpd service — Stop the service by selecting it from the list and clicking the Stop button. A message appears stating that the service was stopped successfully.

2. Reinitialize the runlevel — Reinitialize the runlevel by going to a shell prompt and typing the command telinit x (where x is the runlevel number; in this example, 3.). This option is recommended if you change the Start at Boot value of multiple services and want to activate the changes immediately.

3. Do nothing else — You do not have to stop the httpd service. You can wait until the system is rebooted for the service to stop. The next time the system is booted, the runlevel is initialized without the httpd service running.

To add a service to a runlevel, select the runlevel from the Edit Runlevel pulldown menu, and then select Actions > Add Service. To delete a service from a runlevel, select the runlevel from the Edit Runlevel pulldown menu, select the service to be deleted from the list on the left, and select Actions > Delete Service.

17.4. ntsysv

The ntsysv utility provides a simple interface for activating or deactivating services. You can use ntsysv to turn an xinetd-managed service on or off. You can also use ntsysv to configure runlevels. By default, only the current runlevel is configured. To configure a different runlevel, specify one or more runlevels with the --level option. For example, the command ntsysv --level 345 configures runlevels 3, 4, and 5.

The ntsysv interface works like the text mode installation program. Use the up and down arrows to navigate up and down the list. The space bar selects/unselects services and is also used to "press" the Ok and Cancel buttons. To move between the list of services and the Ok and Cancel buttons, use the Tab key. An asterisk (*) signifies that a service is set to on. Pressing the F1 key displays a short description of the selected service.

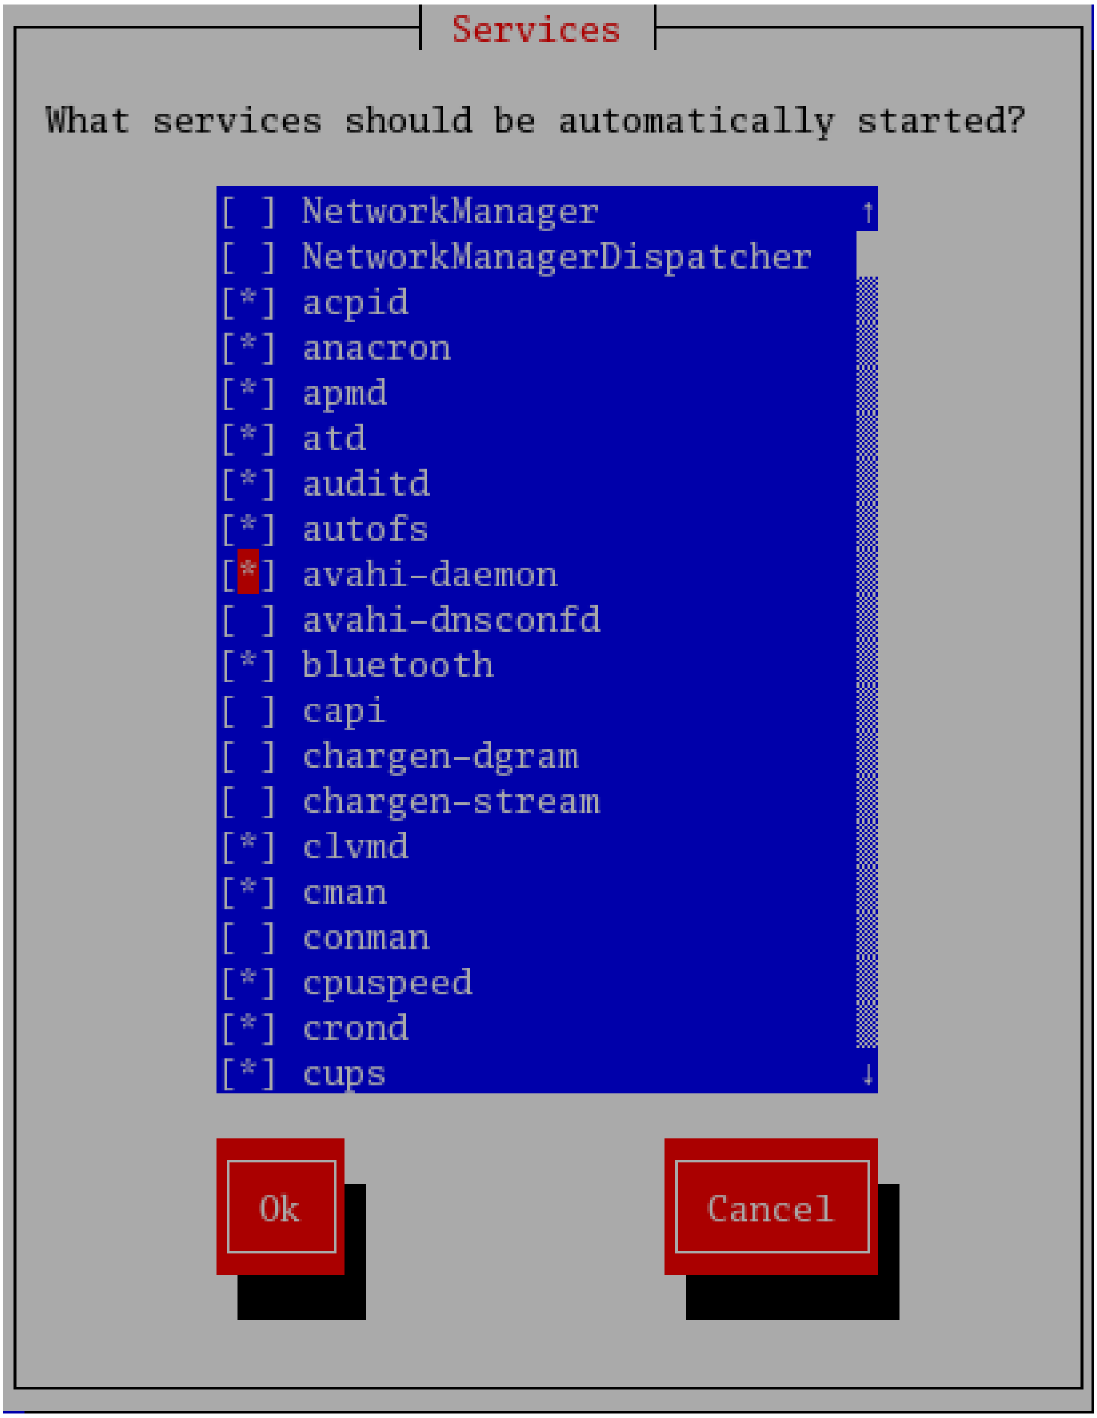

```
Services

What services should be automatically started?

[ ] NetworkManager
[ ] NetworkManagerDispatcher
[*] acpid
[*] anacron
[*] apmd
[*] atd
[*] auditd
[*] autofs
[*] avahi-daemon
[ ] avahi-dnsconfd
[*] bluetooth
[ ] capi
[ ] chargen-dgram
[ ] chargen-stream
[*] clvmd
[*] cman
[ ] conman
[*] cpuspeed
[*] crond
[*] cups
```

그림 17.2. The ntsysv utility

경고

Services managed by xinetd are immediately affected by ntsysv. For all other services, changes do not take effect immediately. You must stop or start the individual service with the command service <daemon> stop (where <daemon> is the name of the service you want to stop; for example, httpd). Replace stop with start or restart to start or restart the service.

17.5. chkconfig

The chkconfig command can also be used to activate and deactivate services. The chkconfig --list command displays a list of system services and whether they are started (on) or stopped (off) in runlevels 0-6. At the end of the list is a section for the services managed by xinetd.

If the chkconfig --list command is used to query a service managed by xinetd, it displays whether the xinetd service is enabled (on) or disabled (off). For example, the command chkconfig --list rsync returns the following output:

```
rsync           on
```

As shown, rsync is enabled as an xinetd service. If xinetd is running, rsync is enabled.

If you use chkconfig --list to query a service in /etc/rc.d, that service's settings for each runlevel are displayed. For example, the command chkconfig --list httpd returns the following output:

```
httpd           0:off   1:off   2:on    3:on    4:on    5:on    6:off
```

chkconfig can also be used to configure a service to be started (or not) in a specific runlevel. For example, to turn nscd off in runlevels 3, 4, and 5, use the following command:

```
chkconfig --level 345 nscd off
```

경고

Services managed by xinetd are immediately affected by chkconfig. For example, if xinetd is running while rsync is disabled, and the command chkconfig rsync on is executed, then rsync is immediately enabled without having to restart xinetd manually. Changes for other services do not take effect immediately after using chkconfig. You must stop or start the individual service with the command service <daemon> stop (where <daemon> is the name of the service you want to stop; for example, httpd). Replace stop with start or restart to start or restart the service.

17.6. 추가 자료

보다 많은 정보를 원하신다면, 다음에 나온 자료를 참조하시기 바랍니다.

17.6.1. 설치된 문서 자료

- The man pages for ntsysv, chkconfig, xinetd, and xinetd.conf.
- man 5 hosts_access — The man page for the format of host access control files (in section 5 of the man pages).

17.6.2. 유용한 웹사이트

- http://www.xinetd.org — The xinetd webpage. It contains sample configuration files and a more detailed list of features.

Berkeley Internet Name Domain (BIND)

On most modern networks, including the Internet, users locate other computers by name. This frees users from the daunting task of remembering the numerical network address of network resources. The most effective way to configure a network to allow such name-based connections is to set up a Domain Name Service (DNS) or a nameserver, which resolves hostnames on the network to numerical addresses and vice versa.

This chapter reviews the nameserver included in Red Hat Enterprise Linux and the Berkeley Internet Name Domain (BIND) DNS server, with an emphasis on the structure of its configuration files and how it may be administered both locally and remotely.

Note

BIND is also known as the service named in Red Hat Enterprise Linux. You can manage it via the Services Configuration Tool (system-config-service).

18.1. Introduction to DNS

DNS associates hostnames with their respective IP addresses, so that when users want to connect to other machines on the network, they can refer to them by name, without having to remember IP addresses.

Use of DNS and FQDNs also has advantages for system administrators, allowing the flexibility to change the IP address for a host without affecting name-based queries to the machine. Conversely, administrators can shuffle which machines handle a name-based query.

DNS is normally implemented using centralized servers that are authoritative for some domains and refer to other DNS servers for other domains.

When a client host requests information from a nameserver, it usually connects to port 53. The nameserver then attempts to resolve the FQDN based on its resolver library, which may contain authoritative information about the host requested or cached data from an earlier query. If the nameserver does not already have the answer in its resolver library, it queries other nameservers, called root nameservers, to determine which nameservers are authoritative for the FQDN in question. Then, with that information, it queries the authoritative nameservers to determine the IP address of the requested host. If a reverse lookup is performed, the same procedure is used, except that the query is made with an unknown IP address rather than a name.

18.1.1. Nameserver Zones

On the Internet, the FQDN of a host can be broken down into different sections. These sections are organized into a hierarchy (much like a tree), with a main trunk, primary branches, secondary branches, and so forth. Consider the following FQDN:

```
bob.sales.example.com
```

When looking at how an FQDN is resolved to find the IP address that relates to a particular system, read the name from right to left, with each level of the hierarchy divided by periods (.). In this example, com defines the top level domain for this FQDN. The name example is a sub-domain under com, while sales is a sub-domain under example. The name furthest to the left, bob, identifies a specific machine hostname.

Except for the hostname, each section is called a zone, which defines a specific namespace. A namespace controls the naming of the sub-domains to its left. While this example only contains two sub-domains, an FQDN must contain at least one sub-domain but may include many more, depending upon how the namespace is organized.

Zones are defined on authoritative nameservers through the use of zone files (which describe the namespace of that zone), the mail servers to be used for a particular domain or sub-domain, and more. Zone files are stored on primary nameservers (also called master nameservers), which are truly authoritative and where changes are made to the files, and secondary nameservers (also called slave nameservers), which receive their zone files from the primary nameservers. Any nameserver can be a primary and secondary nameserver for different zones at the same time, and they may also be considered authoritative for multiple zones. It all depends on how the nameserver is configured.

18.1.2. Nameserver Types

There are four primary nameserver configuration types:

master
: Stores original and authoritative zone records for a namespace, and answers queries about the namespace from other nameservers.

slave
: Answers queries from other nameservers concerning namespaces for which it is considered an authority. However, slave nameservers get their namespace information from master nameservers.

caching-only
: Offers name-to-IP resolution services, but is not authoritative for any zones. Answers for all resolutions are cached in memory for a fixed period of time, which is specified by the retrieved zone record.

forwarding
: Forwards requests to a specific list of nameservers for name resolution. If none of the specified nameservers can perform the resolution, the resolution fails.

A nameserver may be one or more of these types. For example, a nameserver can be a master for some zones, a slave for others, and only offer forwarding resolutions for others.

18.1.3. BIND as a Nameserver

BIND performs name resolution services through the /usr/sbin/named daemon. BIND also includes an administration utility called /usr/sbin/rndc. More information about rndc can be found in 18.4절. “Using rndc ” .

BIND stores its configuration files in the following locations:

/etc/named.conf
: The configuration file for the named daemon

/var/named/ directory
: The named working directory which stores zone, statistic, and cache files

Note

If you have installed the bind-chroot package, the BIND service will run in the /var/named/chroot environment. All configuration files will be moved there. As such, named.conf will be located in /var/named/chroot/etc/named.conf, and so on.

Tip

If you have installed the caching-nameserver package, the default configuration file is /etc/named.caching-nameserver.conf. To override this default configuration, you can create your own custom configuration file in /etc/named.conf. BIND will use the /etc/named.conf custom file instead of the default configuration file after you restart.

The next few sections review the BIND configuration files in more detail.

18.2. /etc/named.conf

The named.conf file is a collection of statements using nested options surrounded by opening and closing ellipse characters, { }. Administrators must be careful when editing named.conf to avoid syntax errors as many seemingly minor errors prevent the named service from starting.

A typical named.conf file is organized similar to the following example:

```
<statement-1> ["<statement-1-name>"] [<statement-1-class>] {
 <option-1>;
 <option-2>;
 <option-N>;
};
<statement-2> ["<statement-2-name>"] [<statement-2-class>] {
 <option-1>;
 <option-2>;
 <option-N>;
};
<statement-N> ["<statement-N-name>"] [<statement-N-class>] {
 <option-1>;
 <option-2>;
 <option-N>;
};
```

18.2.1. Common Statement Types

The following types of statements are commonly used in /etc/named.conf:

18.2.1.1. acl Statement

The acl statement (or access control statement) defines groups of hosts which can then be permitted or denied access to the nameserver.

An acl statement takes the following form:

```
acl <acl-name> {
 <match-element>;
 [<match-element>; ...]
};
```

In this statement, replace <acl-name> with the name of the access control list and replace <match-element> with a semi-colon separated list of IP addresses. Most of the time, an individual IP address or IP network notation (such as 10.0.1.0/24) is used to identify the IP addresses within the acl statement.

The following access control lists are already defined as keywords to simplify configuration:

- any — Matches every IP address
- localhost — Matches any IP address in use by the local system
- localnets — Matches any IP address on any network to which the local system is connected
- none — Matches no IP addresses

When used in conjunction with other statements (such as the options statement), acl statements can be very useful in preventing the misuse of a BIND nameserver.

The following example defines two access control lists and uses an options statement to define how they are treated by the nameserver:

```
acl black-hats {
 10.0.2.0/24;
 192.168.0.0/24;
};
acl red-hats {
 10.0.1.0/24;
};
options {
 blackhole { black-hats; };
 allow-query { red-hats; };
 allow-recursion { red-hats; };
};
```

This example contains two access control lists, black-hats and red-hats. Hosts in the black-hats list are denied access to the nameserver, while hosts in the red-hats list are given normal access.

18.2.1.2. include Statement

The include statement allows files to be included in a named.conf file. In this way, sensitive configuration data (such as keys) can be placed in a separate file with restrictive permissions.

An include statement takes the following form:

```
include "<file-name>"
```

In this statement, <file-name> is replaced with an absolute path to a file.

18.2.1.3. options Statement

The options statement defines global server configuration options and sets defaults for other statements. It can be used to specify the location of the named working directory, the types of queries allowed, and much more.

The options statement takes the following form:

```
options {
 <option>;
 [<option>; ...]
};
```

In this statement, the <option> directives are replaced with a valid option.

The following are commonly used options:

allow-query

Specifies which hosts are allowed to query this nameserver. By default, all hosts are allowed to query. An access control list, or collection of IP addresses or networks, may be used here to allow only particular hosts to query the nameserver.

allow-recursion

Similar to allow-query, this option applies to recursive queries. By default, all hosts are allowed to perform recursive queries on the nameserver.

blackhole

Specifies which hosts are not allowed to query the server.

directory

Specifies the named working directory if different from the default value, /var/named/.

forwarders

Specifies a list of valid IP addresses for nameservers where requests should be forwarded for resolution.

forward

Specifies the forwarding behavior of a forwarders directive.

The following options are accepted:

- first — Specifies that the nameservers listed in the forwarders directive be queried before named attempts to resolve the name itself.
- only — Specifies that named does not attempt name resolution itself in the event that queries to nameservers specified in the forwarders directive fail.

listen-on

Specifies the network interface on which named listens for queries. By default, all interfaces are used.

Using this directive on a DNS server which also acts a gateway, BIND can be configured to only answer queries that originate from one of the networks.

The following is an example of a listen-on directive:

```
options {
 listen-on { 10.0.1.1; };
};
```

In this example, only requests that arrive from the network interface serving the private network (10.0.1.1) are accepted.

notify

Controls whether named notifies the slave servers when a zone is updated. It accepts the following options:

- yes — Notifies slave servers.
- no — Does not notify slave servers.
- explicit — Only notifies slave servers specified in an also-notify list within a zone statement.

pid-file

Specifies the location of the process ID file created by named.

root-delegation-only

Turns on the enforcement of delegation properties in top-level domains (TLDs) and root zones with an optional exclude list. Delegation is the process of dividing a single zone into multiple subzones. In order to create a delegated zone, items known as NS records are used. NameServer records (delegation records) announce the authoritative nameservers for a particular zone.

The following root-delegation-only example specifies an exclude list of TLDs from whom undelegated responses are expected and trusted:

```
options {
 root-delegation-only exclude { "ad"; "ar"; "biz"; "cr"; "cu"; "de"; "dm"; "id";
  "lu"; "lv"; "md"; "ms"; "museum"; "name"; "no"; "pa";
  "pf"; "se"; "sr"; "to"; "tw"; "us"; "uy"; };
};
```

statistics-file

Specifies an alternate location for statistics files. By default, named statistics are saved to the /var/named/named.stats file.

There are several other options also available, many of which rely upon one another to work properly. Refer to the BIND 9 Administrator Reference Manual referenced in 18.7.1절. "Installed Documentation" and the bind.conf man page for more details.

18.2.1.4. zone Statement

A zone statement defines the characteristics of a zone, such as the location of its configuration file and zone-specific options. This statement can be used to override the global options statements.

A zone statement takes the following form:

```
zone <zone-name> <zone-class> {
 <zone-options>;
 [<zone-options>; ...]
};
```

In this statement, <zone-name> is the name of the zone, <zone-class> is the optional class of the zone, and <zone-options> is a list of options characterizing the zone.

The <zone-name> attribute for the zone statement is particularly important. It is the default value assigned for the $ORIGIN directive used within the corresponding zone file located in the /var/named/ directory. The named daemon appends the name of the zone to any non-fully qualified domain name listed in the zone file.

Note

If you have installed the caching-nameserver package, the default configuration file will be in /etc/named.rfc1912.zones.

For example, if a zone statement defines the namespace for example.com, use example.com as the <zone-name> so it is placed at the end of hostnames within the example.com zone file.

For more information about zone files, refer to 18.3절. “Zone Files” .

The most common zone statement options include the following:

allow-query
: Specifies the clients that are allowed to request information about this zone. The default is to allow all query requests.

allow-transfer
: Specifies the slave servers that are allowed to request a transfer of the zone's information. The default is to allow all transfer requests.

allow-update
: Specifies the hosts that are allowed to dynamically update information in their zone. The default is to deny all dynamic update requests.

  Be careful when allowing hosts to update information about their zone. Do not enable this option unless the host specified is completely trusted. In general, it is better to have an administrator manually update the records for a zone and reload the named service.

file
: Specifies the name of the file in the named working directory that contains the zone's configuration data.

masters
: Specifies the IP addresses from which to request authoritative zone information and is used only if the zone is defined as type slave.

notify
: Specifies whether or not named notifies the slave servers when a zone is updated. This directive accepts the following options:

  - yes — Notifies slave servers.
  - no — Does not notify slave servers.
  - explicit — Only notifies slave servers specified in an also-notify list within a zone statement.

type
: Defines the type of zone.

  Below is a list of valid options:

  - delegation-only — Enforces the delegation status of infrastructure zones such as COM, NET, or ORG. Any answer that is received without an explicit or implicit delegation is treated as

NXDOMAIN. This option is only applicable in TLDs or root zone files used in recursive or caching implementations.

- forward — Forwards all requests for information about this zone to other nameservers.
- hint — A special type of zone used to point to the root nameservers which resolve queries when a zone is not otherwise known. No configuration beyond the default is necessary with a hint zone.
- master — Designates the nameserver as authoritative for this zone. A zone should be set as the master if the zone's configuration files reside on the system.
- slave — Designates the nameserver as a slave server for this zone. Also specifies the IP address of the master nameserver for the zone.

zone-statistics

Configures named to keep statistics concerning this zone, writing them to either the default location (/var/named/named.stats) or the file listed in the statistics-file option in the server statement. Refer to 18.2.2절. "Other Statement Types" for more information about the server statement.

18.2.1.5. Sample zone Statements

Most changes to the /etc/named.conf file of a master or slave nameserver involves adding, modifying, or deleting zone statements. While these zone statements can contain many options, most nameservers require only a small subset to function efficiently. The following zone statements are very basic examples illustrating a master-slave nameserver relationship.

The following is an example of a zone statement for the primary nameserver hosting example.com (192.168.0.1):

```
zone "example.com" IN {
  type master;
  file "example.com.zone";
  allow-update { none; };
};
```

In the statement, the zone is identified as example.com, the type is set to master, and the named service is instructed to read the /var/named/example.com.zone file. It also tells named not to allow any other hosts to update.

A slave server's zone statement for example.com is slightly different from the previous example. For a slave server, the type is set to slave and in place of the allow-update line is a directive telling named the IP address of the master server.

The following is an example slave server zone statement for example.com zone:

```
zone "example.com" {
  type slave;
  file "example.com.zone";
  masters { 192.168.0.1; };
};
```

This zone statement configures named on the slave server to query the master server at the 192.168.0.1 IP address for information about the example.com zone. The information that the slave server receives from the master server is saved to the /var/named/example.com.zone file.

18.2.2. Other Statement Types

The following is a list of lesser used statement types available within named.conf:

controls

Configures various security requirements necessary to use the rndc command to administer the named service.

Refer to 18.4.1절. "Configuring /etc/named.conf " to learn more about how the controls statement is structured and what options are available.

key "<key-name>"

Defines a particular key by name. Keys are used to authenticate various actions, such as secure updates or the use of the rndc command. Two options are used with key:

- algorithm <algorithm-name> — The type of algorithm used, such as dsa or hmac-md5.
- secret "<key-value>" — The encrypted key.

Refer to 18.4.2절. "Configuring /etc/rndc.conf " for instructions on how to write a key statement.

logging

Allows for the use of multiple types of logs, called channels. By using the channel option within the logging statement, a customized type of log can be constructed — with its own file name (file), size limit (size), versioning (version), and level of importance (severity). Once a customized channel is defined, a category option is used to categorize the channel and begin logging when named is restarted.

By default, named logs standard messages to the syslog daemon, which places them in /var/log/messages. This occurs because several standard channels are built into BIND with various severity levels, such as default_syslog (which handles informational logging messages) and default_debug (which specifically handles debugging messages). A default category, called default, uses the built-in channels to do normal logging without any special configuration.

Customizing the logging process can be a very detailed process and is beyond the scope of this chapter. For information on creating custom BIND logs, refer to the BIND 9 Administrator Reference Manual referenced in 18.7.1절. "Installed Documentation" .

server

Specifies options that affect how named should respond to remote nameservers, especially with regard to notifications and zone transfers.

The transfer-format option controls whether one resource record is sent with each message (one-answer) or multiple resource records are sent with each message (many-answers). While many-answers is more efficient, only newer BIND nameservers understand it.

trusted-keys

Contains assorted public keys used for secure DNS (DNSSEC). Refer to 18.5.3절. "Security" for more information concerning BIND security.

view "<view-name>"

Creates special views depending upon which network the host querying the nameserver is on. This allows some hosts to receive one answer regarding a zone while other hosts receive totally different information. Alternatively, certain zones may only be made available to particular trusted hosts while non-trusted hosts can only make queries for other zones.

Multiple views may be used, but their names must be unique. The match-clients option specifies the IP addresses that apply to a particular view. Any options statement may also be used within a view, overriding the global options already configured for named. Most view statements contain multiple zone statements that apply to the match-clients list. The order in which view statements are listed is important, as the first view statement that matches a particular client's IP address is used.

Refer to 18.5.2절. "Multiple Views" for more information about the view statement.

18.2.3. Comment Tags

The following is a list of valid comment tags used within named.conf:

- // — When placed at the beginning of a line, that line is ignored by named.
- # — When placed at the beginning of a line, that line is ignored by named.
- /* and */ — When text is enclosed in these tags, the block of text is ignored by named.

18.3. Zone Files

Zone files contain information about a namespace and are stored in the named working directory (/var/named/) by default. Each zone file is named according to the file option data in the zone statement, usually in a way that relates to the domain in question and identifies the file as containing zone data, such as example.com.zone.

Note

If you have installed the bind-chroot package, the BIND service will run in the /var/named/ chroot environment. All configuration files will be moved there. As such, you can find the zone files in /var/named/chroot/var/named.

Each zone file may contain directives and resource records. Directives tell the nameserver to perform tasks or apply special settings to the zone. Resource records define the parameters of the zone and assign identities to individual hosts. Directives are optional, but resource records are required to provide name service to a zone.

All directives and resource records should be entered on individual lines.

Comments can be placed after semicolon characters (;) in zone files.

18.3.1. Zone File Directives

Directives begin with the dollar sign character ($) followed by the name of the directive. They usually appear at the top of the zone file.

The following are commonly used directives:

$INCLUDE

Configures named to include another zone file in this zone file at the place where the directive appears. This allows additional zone settings to be stored apart from the main zone file.

$ORIGIN

Appends the domain name to unqualified records, such as those with the hostname and nothing more.

For example, a zone file may contain the following line:

```
$ORIGIN example.com.
```

Any names used in resource records that do not end in a trailing period (.) are appended with example.com.

Note

The use of the $ORIGIN directive is unnecessary if the zone is specified in /etc/named.conf because the zone name is used as the value for the $ORIGIN directive by default.

$TTL

Sets the default Time to Live (TTL) value for the zone. This is the length of time, in seconds, that a zone resource record is valid. Each resource record can contain its own TTL value, which overrides this directive.

Increasing this value allows remote nameservers to cache the zone information for a longer period of time, reducing the number of queries for the zone and lengthening the amount of time required to proliferate resource record changes.

18.3.2. Zone File Resource Records

The primary component of a zone file is its resource records.

There are many types of zone file resource records. The following are used most frequently:

A

This refers to the Address record, which specifies an IP address to assign to a name, as in this example:

```
<host> IN A <IP-address>
```

If the <host> value is omitted, then an A record points to a default IP address for the top of the namespace. This system is the target for all non-FQDN requests.

Consider the following A record examples for the example.com zone file:

```
server1   IN A 10.0.1.3
   IN A 10.0.1.5
```

Requests for example.com are pointed to 10.0.1.3 or 10.0.1.5.

CNAME

This refers to the Canonical Name record, which maps one name to another. This type of record can also be referred to as an alias record.

The next example tells named that any requests sent to the <alias-name> should point to the host, <real-name>. CNAME records are most commonly used to point to services that use a common naming scheme, such as www for Web servers.

```
<alias-name> IN CNAME <real-name>
```

In the following example, an A record binds a hostname to an IP address, while a CNAME record points the commonly used www hostname to it.

```
server1   IN  A  10.0.1.5
www   IN  CNAME  server1
```

MX

This refers to the Mail eXchange record, which tells where mail sent to a particular namespace controlled by this zone should go.

```
  IN  MX  <preference-value>  <email-server-name>
```

Here, the <preference-value> allows numerical ranking of the email servers for a namespace, giving preference to some email systems over others. The MX resource record with the lowest <preference-value> is preferred over the others. However, multiple email servers can possess the same value to distribute email traffic evenly among them.

The <email-server-name> may be a hostname or FQDN.

```
IN      MX      10      mail.example.com.
IN      MX      20      mail2.example.com.
```

In this example, the first mail.example.com email server is preferred to the mail2.example.com email server when receiving email destined for the example.com domain.

NS

This refers to the NameServer record, which announces the authoritative nameservers for a particular zone.

The following illustrates the layout of an NS record:

```
  IN  NS  <nameserver-name>
```

Here, <nameserver-name> should be an FQDN.

Next, two nameservers are listed as authoritative for the domain. It is not important whether these nameservers are slaves or if one is a master; they are both still considered authoritative.

```
IN      NS      dns1.example.com.
IN      NS      dns2.example.com.
```

PTR

This refers to the PoinTeR record, which is designed to point to another part of the namespace.

PTR records are primarily used for reverse name resolution, as they point IP addresses back to a particular name. Refer to 18.3.4절. "Reverse Name Resolution Zone Files" for more examples of PTR records in use.

SOA

This refers to the Start Of Authority resource record, which proclaims important authoritative information about a namespace to the nameserver.

Located after the directives, an SOA resource record is the first resource record in a zone file.

The following shows the basic structure of an SOA resource record:

```
@  IN SOA  <primary-name-server>  <hostmaster-email> (
 <serial-number>
 <time-to-refresh>
 <time-to-retry>
 <time-to-expire>
 <minimum-TTL> )
```

The @ symbol places the $ORIGIN directive (or the zone's name, if the $ORIGIN directive is not set) as the namespace being defined by this SOA resource record. The hostname of the primary nameserver that is authoritative for this domain is the <primary-name-server> directive, and the email of the person to contact about this namespace is the <hostmaster-email> directive.

The <serial-number> directive is a numerical value incremented every time the zone file is altered to indicate it is time for named to reload the zone. The <time-to-refresh> directive is the numerical value slave servers use to determine how long to wait before asking the master nameserver if any changes have been made to the zone. The <serial-number> directive is a numerical value used by the slave servers to determine if it is using outdated zone data and should therefore refresh it.

The <time-to-retry> directive is a numerical value used by slave servers to determine the length of time to wait before issuing a refresh request in the event that the master nameserver is not answering. If the master has not replied to a refresh request before the amount of time specified in the <time-to-expire> directive elapses, the slave servers stop responding as an authority for requests concerning that namespace.

In BIND 4 and 8, the <minimum-TTL> directive is the amount of time other nameservers cache the zone's information. However, in BIND 9, the <minimum-TTL> directive defines how long negative answers are cached for. Caching of negative answers can be set to a maximum of 3 hours (3H).

When configuring BIND, all times are specified in seconds. However, it is possible to use abbreviations when specifying units of time other than seconds, such as minutes (M), hours (H), days (D), and weeks (W). The table in 표 18.1. "Seconds compared to other time units" shows an amount of time in seconds and the equivalent time in another format.

표 18.1. Seconds compared to other time units

Seconds	Other Time Units
60	1M
1800	30M
3600	1H
10800	3H
21600	6H

Seconds	Other Time Units
43200	12H
86400	1D
259200	3D
604800	1W
31536000	365D

The following example illustrates the form an SOA resource record might take when it is populated with real values.

```
@ IN SOA dns1.example.com. hostmaster.example.com. (
  2001062501 ; serial
  21600      ; refresh after 6 hours
  3600       ; retry after 1 hour
  604800     ; expire after 1 week
  86400 )    ; minimum TTL of 1 day
```

18.3.3. Example Zone File

Seen individually, directives and resource records can be difficult to grasp. However, when placed together in a single file, they become easier to understand.

The following example shows a very basic zone file.

```
$ORIGIN example.com.
$TTL 86400
@   IN SOA dns1.example.com. hostmaster.example.com. (
    2001062501 ; serial
    21600      ; refresh after 6 hours
    3600       ; retry after 1 hour
    604800     ; expire after 1 week
    86400 )    ; minimum TTL of 1 day
;
;
  IN NS dns1.example.com.
  IN NS dns2.example.com.
dns1   IN A 10.0.1.1
  IN AAAA aaaa:bbbb::1
dns2   IN A 10.0.1.2
  IN AAAA aaaa:bbbb::2
;
;
@   IN MX 10 mail.example.com.
  IN MX 20 mail2.example.com.
mail   IN A 10.0.1.5
  IN AAAA aaaa:bbbb::5
mail2   IN A 10.0.1.6
  IN AAAA aaaa:bbbb::6
;
;
; This sample zone file illustrates sharing the same IP addresses
; for multiple services:
;
services IN A 10.0.1.10
  IN AAAA aaaa:bbbb::10
  IN A 10.0.1.11
  IN AAAA aaaa:bbbb::11
```

```
ftp   IN CNAME services.example.com.
www   IN CNAME services.example.com.
;
;
```

In this example, standard directives and SOA values are used. The authoritative nameservers are set as dns1.example.com and dns2.example.com, which have A records that tie them to 10.0.1.1 and 10.0.1.2, respectively.

The email servers configured with the MX records point to mail and mail2 via A records. Since the mail and mail2 names do not end in a trailing period (.), the $ORIGIN domain is placed after them, expanding them to mail.example.com and mail2.example.com. Through the related A resource records, their IP addresses can be determined.

Services available at the standard names, such as www.example.com (WWW), are pointed at the appropriate servers using a CNAME record.

This zone file would be called into service with a zone statement in the named.conf similar to the following:

```
zone "example.com" IN {
 type master;
 file "example.com.zone";
 allow-update { none; };
};
```

18.3.4. Reverse Name Resolution Zone Files

A reverse name resolution zone file is used to translate an IP address in a particular namespace into an FQDN. It looks very similar to a standard zone file, except that PTR resource records are used to link the IP addresses to a fully qualified domain name.

The following illustrates the layout of a PTR record:

```
<last-IP-digit> IN PTR <FQDN-of-system>
```

The <last-IP-digit> is the last number in an IP address which points to a particular system's FQDN.

In the following example, IP addresses 10.0.1.1 through 10.0.1.6 are pointed to corresponding FQDNs. It can be located in /var/named/example.com.rr.zone.

```
$ORIGIN 1.0.10.in-addr.arpa.
$TTL 86400
@ IN SOA dns1.example.com. hostmaster.example.com. (
    2001062501 ; serial
    21600      ; refresh after 6 hours
    3600       ; retry after 1 hour
    604800     ; expire after 1 week
    86400 )    ; minimum TTL of 1 day
;
@ IN NS dns1.example.com.
;
1 IN PTR dns1.example.com.
2 IN PTR dns2.example.com.
```

```
;
5 IN PTR server1.example.com.
6 IN PTR server2.example.com.
;
3 IN PTR ftp.example.com.
4 IN PTR ftp.example.com.
```

This zone file would be called into service with a zone statement in the named.conf file similar to the following:

```

zone "1.0.10.in-addr.arpa" IN {
 type master;
 file "example.com.rr.zone";
 allow-update { none; };
};
```

There is very little difference between this example and a standard zone statement, except for the zone name. Note that a reverse name resolution zone requires the first three blocks of the IP address reversed followed by .in-addr.arpa. This allows the single block of IP numbers used in the reverse name resolution zone file to be associated with the zone.

18.4. Using rndc

BIND includes a utility called rndc which allows command line administration of the named daemon from the localhost or a remote host.

In order to prevent unauthorized access to the named daemon, BIND uses a shared secret key authentication method to grant privileges to hosts. This means an identical key must be present in both /etc/named.conf and the rndc configuration file, /etc/rndc.conf.

Note

If you have installed the bind-chroot package, the BIND service will run in the /var/named/chroot environment. All configuration files will be moved there. As such, the rndc.conf file is located in /var/named/chroot/etc/rndc.conf.

Note that since the rndc utility does not run in a chroot environment, /etc/rndc.conf is a symlink to /var/named/chroot/etc/rndc.conf.

18.4.1. Configuring /etc/named.conf

In order for rndc to connect to a named service, there must be a controls statement in the BIND server's /etc/named.conf file.

The controls statement, shown in the following example, allows rndc to connect from the localhost.

```

controls {
 inet 127.0.0.1
  allow { localhost; } keys { <key-name>; };
};
```

This statement tells named to listen on the default TCP port 953 of the loopback address and allow rndc commands coming from the localhost, if the proper key is given. The <key-name> specifies a name in the key statement within the /etc/named.conf file. The next example illustrates a sample key statement.

```
key "<key-name>" {
 algorithm hmac-md5;
 secret "<key-value>";
};
```

In this case, the <key-value> uses the HMAC-MD5 algorithm. Use the following command to generate keys using the HMAC-MD5 algorithm:

```
dnssec-keygen -a hmac-md5 -b <bit-length> -n HOST <key-file-name>
```

A key with at least a 256-bit length is a good idea. The actual key that should be placed in the <key-value> area can be found in the <key-file-name> file generated by this command.

Warning

Because /etc/named.conf is world-readable, it is advisable to place the key statement in a separate file, readable only by root, and then use an include statement to reference it. For example:

```
include "/etc/rndc.key";
```

18.4.1.1. Firewall Blocking Communication

If a firewall is blocking connections from the named daemon to other nameservers, the recommended best practice is to change the firewall settings whenever possible.

Warning: Avoid Using Fixed UDP Source Ports

DNS resolvers, that are not configured to perform DNSSEC validation or that need to query DNS zones that are not protected by DNSSEC only, use a 16-bit transaction identifier (TXID) and the destination UDP port number to check whether the DNS reply was sent by the server they queried for DNS data.

Previously, BIND always used a fixed UDP source port when sending DNS queries. BIND used either a port configured using the query-source (and query-source-v6) directive, or one randomly chosen at startup. When a static query source port is used, TXID offers insufficient protection against spoofed replies and allows an attacker to efficiently perform cache-poisoning attacks. To address this issue, BIND was updated to allow the use of a randomly-selected source port for each DNS query, making it more difficult for an attacker to spoof replies, when the query packets cannot be detected. A security update [1] was released for all the affected Red Hat Enterprise Linux versions. Additionally, the default configuration provided by the caching-nameserver package was updated to no longer specify a fixed query source port.

When deploying BIND as a DNS resolver, ensure that BIND is not forced, by the aforementioned configuration directives, to use a fixed query source port. Your firewall configuration must also permit the use of random query source ports. Previously, it was common practice to configure BIND to use port 53 as a query source port, and only allow DNS queries from that port on the firewall.

18.4.2. Configuring /etc/rndc.conf

The key is the most important statement in /etc/rndc.conf.

```
key "<key-name>" {
 algorithm hmac-md5;
 secret "<key-value>";
};
```

The <key-name> and <key-value> should be exactly the same as their settings in /etc/named.conf.

To match the keys specified in the target server's /etc/named.conf, add the following lines to /etc/rndc.conf.

```
options {
 default-server   localhost;
 default-key        "<key-name>";
};
```

This directive sets a global default key. However, the rndc configuration file can also specify different keys for different servers, as in the following example:

[1] The security update was RHSA-2008:0533 [https://rhn.redhat.com/errata/RHSA-2008-0533.html].

```
server localhost {
 key  "<key-name>";
};
```

Important

Make sure that only the root user can read or write to the /etc/rndc.conf file.

For more information about the /etc/rndc.conf file, refer to the rndc.conf man page.

18.4.3. Command Line Options

An rndc command takes the following form:

```

rndc <options> <command> <command-options>
```

When executing rndc on a properly configured localhost, the following commands are available:

- halt — Stops the named service immediately.
- querylog — Logs all queries made to this nameserver.
- refresh — Refreshes the nameserver's database.
- reload — Reloads the zone files but keeps all other previously cached responses. This command also allows changes to zone files without losing all stored name resolutions.

  If changes made only affect a specific zone, reload only that specific zone by adding the name of the zone after the reload command.
- stats — Dumps the current named statistics to the /var/named/named.stats file.
- stop — Stops the server gracefully, saving any dynamic update and Incremental Zone Transfers (IXFR) data before exiting.

Occasionally, it may be necessary to override the default settings in the /etc/rndc.conf file. The following options are available:

- -c <configuration-file> — Specifies the alternate location of a configuration file.
- -p <port-number> — Specifies a port number to use for the rndc connection other than the default port 953.
- -s <server> — Specifies a server other than the default-server listed in /etc/rndc.conf.
- -y <key-name> — Specifies a key other than the default-key option in /etc/rndc.conf.

Additional information about these options can be found in the rndc man page.

18.5. Advanced Features of BIND

Most BIND implementations only use named to provide name resolution services or to act as an authority for a particular domain or sub-domain. However, BIND version 9 has a number of advanced features that allow for a more secure and efficient DNS service.

Caution

Some of these advanced features, such as DNSSEC, TSIG, and IXFR (which are defined in the following section), should only be used in network environments with nameservers that support the features. If the network environment includes non-BIND or older BIND nameservers, verify that each advanced feature is supported before attempting to use it.

All of the features mentioned are discussed in greater detail in the BIND 9 Administrator Reference Manual referenced in 18.7.1절. "Installed Documentation" .

18.5.1. DNS Protocol Enhancements

BIND supports Incremental Zone Transfers (IXFR), where a slave nameserver only downloads the updated portions of a zone modified on a master nameserver. The standard transfer process requires that the entire zone be transferred to each slave nameserver for even the smallest change. For very popular domains with very lengthy zone files and many slave nameservers, IXFR makes the notification and update process much less resource-intensive.

Note that IXFR is only available when using dynamic updating to make changes to master zone records. If manually editing zone files to make changes, Automatic Zone Transfer (AXFR) is used. More information on dynamic updating is available in the BIND 9 Administrator Reference Manual referenced in 18.7.1절. "Installed Documentation" .

18.5.2. Multiple Views

Through the use of the view statement in named.conf, BIND can present different information depending on which network a request originates from.

This is primarily used to deny sensitive DNS entries from clients outside of the local network, while allowing queries from clients inside the local network.

The view statement uses the match-clients option to match IP addresses or entire networks and give them special options and zone data.

18.5.3. Security

BIND supports a number of different methods to protect the updating and transfer of zones, on both master and slave nameservers:

DNSSEC
: Short for DNS SECurity, this feature allows for zones to be cryptographically signed with a zone key.

In this way, the information about a specific zone can be verified as coming from a nameserver that has signed it with a particular private key, as long as the recipient has that nameserver's public key.

BIND version 9 also supports the SIG(0) public/private key method of message authentication.

TSIG

Short for Transaction SIGnatures, this feature allows a transfer from master to slave only after verifying that a shared secret key exists on both nameservers.

This feature strengthens the standard IP address-based method of transfer authorization. An attacker would not only need to have access to the IP address to transfer the zone, but they would also need to know the secret key.

BIND version 9 also supports TKEY, which is another shared secret key method of authorizing zone transfers.

18.5.4. IP version 6

BIND version 9 supports name service in IP version 6 (IPv6) environments through the use of A6 zone records.

If the network environment includes both IPv4 and IPv6 hosts, use the lwresd lightweight resolver daemon on all network clients. This daemon is a very efficient, caching-only nameserver which understands the new A6 and DNAME records used under IPv6. Refer to the lwresd man page for more information.

18.6. Common Mistakes to Avoid

It is very common for beginners to make mistakes when editing BIND configuration files. Be sure to avoid the following issues:

- Take care to increment the serial number when editing a zone file.

  If the serial number is not incremented, the master nameserver has the correct, new information, but the slave nameservers are never notified of the change and do not attempt to refresh their data of that zone.

- Be careful to use ellipses and semi-colons correctly in the /etc/named.conf file.

  An omitted semi-colon or unclosed ellipse section can cause named to refuse to start.

- Remember to place periods (.) in zone files after all FQDNs and omit them on hostnames.

  A period at the end of a domain name denotes a fully qualified domain name. If the period is omitted, then named appends the name of the zone or the $ORIGIN value to complete it.

- If a firewall is blocking connections from the named daemon to other nameservers, the recommended best practice is to change the firewall settings whenever possible. For important security information regarding fixed UDP source ports, refer to 18.4.1.1절. "Firewall Blocking Communication"

18.7. Additional Resources

The following sources of information provide additional resources regarding BIND.

18.7.1. Installed Documentation

BIND features a full range of installed documentation covering many different topics, each placed in its own subject directory. For each item below, replace <version-number> with the version of bind installed on the system:

/usr/share/doc/bind-<version-number>/
: This directory lists the most recent features.

/usr/share/doc/bind-<version-number>/arm/
: This directory contains the BIND 9 Administrator Reference Manual in HTML and SGML formats, which details BIND resource requirements, how to configure different types of nameservers, how to perform load balancing, and other advanced topics. For most new users of BIND, this is the best place to start.

/usr/share/doc/bind-<version-number>/draft/
: This directory contains assorted technical documents that review issues related to DNS service and propose some methods to address them.

/usr/share/doc/bind-<version-number>/misc/
: This directory contains documents designed to address specific advanced issues. Users of BIND version 8 should consult the migration document for specific changes they must make when moving to BIND 9. The options file lists all of the options implemented in BIND 9 that are used in /etc/named.conf.

/usr/share/doc/bind-<version-number>/rfc/
: This directory provides every RFC document related to BIND.

There are also a number of man pages for the various applications and configuration files involved with BIND. The following lists some of the more important man pages.

Administrative Applications

- man rndc — Explains the different options available when using the rndc command to control a BIND nameserver.

Server Applications

- man named — Explores assorted arguments that can be used to control the BIND nameserver daemon.

- man lwresd — Describes the purpose of and options available for the lightweight resolver daemon.

Configuration Files

- man named.conf — A comprehensive list of options available within the named configuration file.

- man rndc.conf — A comprehensive list of options available within the rndc configuration file.

18.7.2. Useful Websites

- http://www.isc.org/index.pl?/sw/bind/ — The home page of the BIND project containing information about current releases as well as a PDF version of the BIND 9 Administrator Reference Manual.

- http://www.redhat.com/mirrors/LDP/HOWTO/DNS-HOWTO.html — Covers the use of BIND as a resolving, caching nameserver and the configuration of various zone files necessary to serve as the primary nameserver for a domain.

18.7.3. Related Books

- DNS and BIND by Paul Albitz and Cricket Liu; O'Reilly & Associates — A popular reference that explains both common and esoteric BIND configuration options, as well as providing strategies for securing a DNS server.
- The Concise Guide to DNS and BIND by Nicolai Langfeldt; Que — Looks at the connection between multiple network services and BIND, with an emphasis on task-oriented, technical topics.

OpenSSH

SSH™ (or Secure SHell) is a protocol which facilitates secure communications between two systems using a client/server architecture and allows users to log into server host systems remotely. Unlike other remote communication protocols, such as FTP or Telnet, SSH encrypts the login session, rendering the connection difficult for intruders to collect unencrypted passwords.

SSH is designed to replace older, less secure terminal applications used to log into remote hosts, such as telnet or rsh. A related program called scp replaces older programs designed to copy files between hosts, such as rcp. Because these older applications do not encrypt passwords transmitted between the client and the server, avoid them whenever possible. Using secure methods to log into remote systems decreases the risks for both the client system and the remote host.

19.1. Features of SSH

The SSH protocol provides the following safeguards:

- After an initial connection, the client can verify that it is connecting to the same server it had connected to previously.
- The client transmits its authentication information to the server using strong, 128-bit encryption.
- All data sent and received during a session is transferred using 128-bit encryption, making intercepted transmissions extremely difficult to decrypt and read.
- The client can forward X11[1] applications from the server. This technique, called X11 forwarding, provides a secure means to use graphical applications over a network.

Because the SSH protocol encrypts everything it sends and receives, it can be used to secure otherwise insecure protocols. Using a technique called port forwarding, an SSH server can become a conduit to securing otherwise insecure protocols, like POP, and increasing overall system and data security.

The OpenSSH server and client can also be configured to create a tunnel similar to a virtual private network for traffic between server and client machines.

Finally, OpenSSH servers and clients can be configured to authenticate using the GSSAPI implementation of the Kerberos network authentication protocol. For more information on configuring Kerberos authentication services, refer to 46.6절. “Kerberos” .

Red Hat Enterprise Linux includes the general OpenSSH package (openssh) as well as the OpenSSH server (openssh-server) and client (openssh-clients) packages. Note, the OpenSSH packages require the OpenSSL package (openssl) which installs several important cryptographic libraries, enabling OpenSSH to provide encrypted communications.

19.1.1. Why Use SSH?

Nefarious computer users have a variety of tools at their disposal enabling them to disrupt, intercept, and re-route network traffic in an effort to gain access to a system. In general terms, these threats can be categorized as follows:

[1] X11 refers to the X11R7 windowing display system, traditionally referred to as the X Window System or X. Red Hat Enterprise Linux includes X11R7, an open source X Window System.

- Interception of communication between two systems — In this scenario, the attacker can be somewhere on the network between the communicating parties, copying any information passed between them. The attacker may intercept and keep the information, or alter the information and send it on to the intended recipient.

  This attack can be mounted through the use of a packet sniffer — a common network utility.

- Impersonation of a particular host — Using this strategy, an attacker's system is configured to pose as the intended recipient of a transmission. If this strategy works, the user's system remains unaware that it is communicating with the wrong host.

  This attack can be mounted through techniques known as DNS poisoning[2] or IP spoofing[3].

Both techniques intercept potentially sensitive information and, if the interception is made for hostile reasons, the results can be disastrous.

If SSH is used for remote shell login and file copying, these security threats can be greatly diminished. This is because the SSH client and server use digital signatures to verify their identity. Additionally, all communication between the client and server systems is encrypted. Attempts to spoof the identity of either side of a communication does not work, since each packet is encrypted using a key known only by the local and remote systems.

19.2. SSH Protocol Versions

The SSH protocol allows any client and server programs built to the protocol's specifications to communicate securely and to be used interchangeably.

Two varieties of SSH (version 1 and version 2) currently exist. The OpenSSH suite under Red Hat Enterprise Linux uses SSH version 2 which has an enhanced key exchange algorithm not vulnerable to the exploit in version 1. However, the OpenSSH suite does support version 1 connections.

Important

It is recommended that only SSH version 2-compatible servers and clients are used whenever possible.

19.3. Event Sequence of an SSH Connection

The following series of events help protect the integrity of SSH communication between two hosts.

1. A cryptographic handshake is made so that the client can verify that it is communicating with the correct server.

2. The transport layer of the connection between the client and remote host is encrypted using a symmetric cipher.

3. The client authenticates itself to the server.

[2] DNS poisoning occurs when an intruder cracks a DNS server, pointing client systems to a maliciously duplicated host.

[3] IP spoofing occurs when an intruder sends network packets which falsely appear to be from a trusted host on the network.

4. The remote client interacts with the remote host over the encrypted connection.

19.3.1. Transport Layer

The primary role of the transport layer is to facilitate safe and secure communication between the two hosts at the time of authentication and during subsequent communication. The transport layer accomplishes this by handling the encryption and decryption of data, and by providing integrity protection of data packets as they are sent and received. The transport layer also provides compression, speeding the transfer of information.

Once an SSH client contacts a server, key information is exchanged so that the two systems can correctly construct the transport layer. The following steps occur during this exchange:

- Keys are exchanged
- The public key encryption algorithm is determined
- The symmetric encryption algorithm is determined
- The message authentication algorithm is determined
- The hash algorithm is determined

During the key exchange, the server identifies itself to the client with a unique host key. If the client has never communicated with this particular server before, the server's host key is unknown to the client and it does not connect. OpenSSH gets around this problem by accepting the server's host key. This is done after the user is notified and has both accepted and verified the new host key. In subsequent connections, the server's host key is checked against the saved version on the client, providing confidence that the client is indeed communicating with the intended server. If, in the future, the host key no longer matches, the user must remove the client's saved version before a connection can occur.

Caution

It is possible for an attacker to masquerade as an SSH server during the initial contact since the local system does not know the difference between the intended server and a false one set up by an attacker. To help prevent this, verify the integrity of a new SSH server by contacting the server administrator before connecting for the first time or in the event of a host key mismatch.

SSH is designed to work with almost any kind of public key algorithm or encoding format. After an initial key exchange creates a hash value used for exchanges and a shared secret value, the two systems immediately begin calculating new keys and algorithms to protect authentication and future data sent over the connection.

After a certain amount of data has been transmitted using a given key and algorithm (the exact amount depends on the SSH implementation), another key exchange occurs, generating another set of hash values and a new shared secret value. Even if an attacker is able to determine the hash and shared secret value, this information is only useful for a limited period of time.

19.3.2. Authentication

Once the transport layer has constructed a secure tunnel to pass information between the two systems, the server tells the client the different authentication methods supported, such as using a private key-encoded signature or typing a password. The client then tries to authenticate itself to the server using one of these supported methods.

SSH servers and clients can be configured to allow different types of authentication, which gives each side the optimal amount of control. The server can decide which encryption methods it supports based on its security model, and the client can choose the order of authentication methods to attempt from the available options.

19.3.3. Channels

After a successful authentication over the SSH transport layer, multiple channels are opened via a technique called multiplexing[4]. Each of these channels handles communication for different terminal sessions and for forwarded X11 sessions.

Both clients and servers can create a new channel. Each channel is then assigned a different number on each end of the connection. When the client attempts to open a new channel, the clients sends the channel number along with the request. This information is stored by the server and is used to direct communication to that channel. This is done so that different types of sessions do not affect one another and so that when a given session ends, its channel can be closed without disrupting the primary SSH connection.

Channels also support flow-control, which allows them to send and receive data in an orderly fashion. In this way, data is not sent over the channel until the client receives a message that the channel is open.

The client and server negotiate the characteristics of each channel automatically, depending on the type of service the client requests and the way the user is connected to the network. This allows great flexibility in handling different types of remote connections without having to change the basic infrastructure of the protocol.

19.4. OpenSSH 서버 설정

To run an OpenSSH server, you must first make sure that you have the proper RPM packages installed. The openssh-server package is required and is dependent on the openssh package.

The OpenSSH daemon uses the configuration file /etc/ssh/sshd_config. The default configuration file should be sufficient for most purposes. If you want to configure the daemon in ways not provided by the default sshd_config, read the sshd man page for a list of the keywords that can be defined in the configuration file.

To start the OpenSSH service, use the command /sbin/service sshd start. To stop the OpenSSH server, use the command /sbin/service sshd stop. If you want the daemon to start automatically at boot time, refer to 17장. 서비스로의 접근 통제 for information on how to manage services.

시스템을 재설치하시면 새로운 인증키가 생성됩니다. 재설치하기 전에 OpenSSH 도구를 사용하여 시스템에 연결했었던 클라이언트는 다시 재접속을 시도시 다음과 같은 메시지를 보게될 것입니다:

[4] A multiplexed connection consists of several signals being sent over a shared, common medium. With SSH, different channels are sent over a common secure connection.

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@     WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!      @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA host key has just been changed.
```

If you want to keep the host keys generated for the system, backup the /etc/ssh/ssh_host*key* files and restore them after the reinstall. This process retains the system's identity, and when clients try to connect to the system after the reinstall, they will not receive the warning message.

19.4.1. Requiring SSH for Remote Connections

For SSH to be truly effective, using insecure connection protocols, such as Telnet and FTP, should be prohibited. Otherwise, a user's password may be protected using SSH for one session, only to be captured later while logging in using Telnet.

Some services to disable include:

- telnet
- rsh
- rlogin
- vsftpd

To disable insecure connection methods to the system, use the command line program chkconfig, the ncurses-based program /usr/sbin/ntsysv, or the Services Configuration Tool (system-config-services) graphical application. All of these tools require root level access.

For more information on runlevels and configuring services with chkconfig, /usr/sbin/ntsysv, and the Services Configuration Tool, refer to 17장. 서비스로의 접근 통제.

19.5. OpenSSH Configuration Files

OpenSSH has two different sets of configuration files: one for client programs (ssh, scp, and sftp) and one for the server daemon (sshd).

System-wide SSH configuration information is stored in the /etc/ssh/ directory:

- moduli — Contains Diffie-Hellman groups used for the Diffie-Hellman key exchange which is critical for constructing a secure transport layer. When keys are exchanged at the beginning of an SSH session, a shared, secret value is created which cannot be determined by either party alone. This value is then used to provide host authentication.
- ssh_config — The system-wide default SSH client configuration file. It is overridden if one is also present in the user's home directory (~/.ssh/config).
- sshd_config — The configuration file for the sshd daemon.
- ssh_host_dsa_key — The DSA private key used by the sshd daemon.
- ssh_host_dsa_key.pub — The DSA public key used by the sshd daemon.
- ssh_host_key — The RSA private key used by the sshd daemon for version 1 of the SSH protocol.

- ssh_host_key.pub — The RSA public key used by the sshd daemon for version 1 of the SSH protocol.
- ssh_host_rsa_key — The RSA private key used by the sshd daemon for version 2 of the SSH protocol.
- ssh_host_rsa_key.pub — The RSA public key used by the sshd for version 2 of the SSH protocol.

User-specific SSH configuration information is stored in the user's home directory within the ~/.ssh/ directory:

- authorized_keys — This file holds a list of authorized public keys for servers. When the client connects to a server, the server authenticates the client by checking its signed public key stored within this file.
- id_dsa — Contains the DSA private key of the user.
- id_dsa.pub — The DSA public key of the user.
- id_rsa — The RSA private key used by ssh for version 2 of the SSH protocol.
- id_rsa.pub — The RSA public key used by ssh for version 2 of the SSH protocol
- identity — The RSA private key used by ssh for version 1 of the SSH protocol.
- identity.pub — The RSA public key used by ssh for version 1 of the SSH protocol.
- known_hosts — This file contains DSA host keys of SSH servers accessed by the user. This file is very important for ensuring that the SSH client is connecting the correct SSH server.

> **Important**
>
> If an SSH server's host key has changed, the client notifies the user that the connection cannot proceed until the server's host key is deleted from the known_hosts file using a text editor. Before doing this, however, contact the system administrator of the SSH server to verify the server is not compromised.

Refer to the ssh_config and sshd_config man pages for information concerning the various directives available in the SSH configuration files.

19.6. OpenSSH 클라이언트 설정

To connect to an OpenSSH server from a client machine, you must have the openssh-clients and openssh packages installed on the client machine.

19.6.1. Using the ssh Command

The ssh command is a secure replacement for the rlogin, rsh, and telnet commands. It allows you to log in to a remote machine as well as execute commands on a remote machine.

Logging in to a remote machine with ssh is similar to using telnet. To log in to a remote machine named penguin.example.net, type the following command at a shell prompt:

```
ssh penguin.example.net
```

The first time you ssh to a remote machine, you will see a message similar to the following:

```
The authenticity of host 'penguin.example.net' can't be established.
DSA key fingerprint is 94:68:3a:3a:bc:f3:9a:9b:01:5d:b3:07:38:e2:11:0c.
Are you sure you want to continue connecting (yes/no)?
```

Type yes to continue. This will add the server to your list of known hosts (~/.ssh/known_hosts) as seen in the following message:

```
Warning: Permanently added 'penguin.example.net' (RSA) to the list of known hosts.
```

다음으로 원격 컴퓨터에 접속하기 위한 암호를 입력하셔야 합니다. 암호를 입력하시면, 원격 컴퓨터의 쉘 프롬프트가 나타납니다. 특별히 사용자명을 지정하지 않으시면, 로컬 클라이언트 컴퓨터에서 로그인하셨던 사용자명이 원격 컴퓨터로 전달됩니다. 다른 사용자명을 지정하기 위해서는 다음 명령을 사용하십시오:

```
ssh username@penguin.example.net
```

You can also use the syntax ssh -l username penguin.example.net.

The ssh command can be used to execute a command on the remote machine without logging in to a shell prompt. The syntax is ssh hostname command. For example, if you want to execute the command ls /usr/share/doc on the remote machine penguin.example.net, type the following command at a shell prompt:

```
ssh penguin.example.net ls /usr/share/doc
```

After you enter the correct password, the contents of the remote directory /usr/share/doc will be displayed, and you will return to your local shell prompt.

19.6.2. Using the scp Command

The scp command can be used to transfer files between machines over a secure, encrypted connection. It is similar to rcp.

원격 시스템으로 로컬 파일을 전송할 때 사용되는 일반적인 명령 구문은 다음과 같습니다:

```
scp <localfile> username@tohostname:<remotefile>
```

The <localfile> specifies the source including path to the file, such as /var/log/maillog. The <remotefile> specifies the destination, which can be a new filename such as /tmp/hostname-maillog. For the remote system, if you do not have a preceding /, the path will be relative to the home directory of username, typically /home/username/.

To transfer the local file shadowman to the home directory of your account on penguin.example.net, type the following at a shell prompt (replace username with your username):

```
scp shadowman username@penguin.example.net:shadowman
```

This will transfer the local file shadowman to /home/username/shadowman on penguin.example.net. Alternately, you can leave off the final shadowman in the scp command.

원격 컴퓨터에서 로컬 시스템으로 파일을 전송하기 위해서는 다음과 같은 명령 구문이 사용됩니다:

```
scp username@tohostname:<remotefile> <newlocalfile>
```

The <remotefile> specifies the source including path, and <newlocalfile> specifies the destination including path.

Multiple files can be specified as the source files. For example, to transfer the contents of the directory downloads/ to an existing directory called uploads/ on the remote machine penguin.example.net, type the following at a shell prompt:

```
scp downloads/* username@penguin.example.net:uploads/
```

19.6.3. Using the sftp Command

The sftp utility can be used to open a secure, interactive FTP session. It is similar to ftp except that it uses a secure, encrypted connection. The general syntax is sftp username@hostname.com. Once authenticated, you can use a set of commands similar to those used by FTP. Refer to the sftp man page for a list of these commands. To read the man page, execute the command man sftp at a shell prompt. The sftp utility is only available in OpenSSH version 2.5.0p1 and higher.

19.7. More Than a Secure Shell

A secure command line interface is just the beginning of the many ways SSH can be used. Given the proper amount of bandwidth, X11 sessions can be directed over an SSH channel. Or, by using TCP/IP forwarding, previously insecure port connections between systems can be mapped to specific SSH channels.

19.7.1. X11 Forwarding

Opening an X11 session over an SSH connection is as easy as connecting to the SSH server using the -Y option and running an X program on a local machine.

```
ssh -Y <user>@example.com
```

When an X program is run from the secure shell prompt, the SSH client and server create a new secure channel, and the X program data is sent over that channel to the client machine transparently.

X11 forwarding can be very useful. For example, X11 forwarding can be used to create a secure, interactive session of the Printer Configuration Tool. To do this, connect to the server using ssh and type:

```
system-config-printer &
```

After supplying the root password for the server, the Printer Configuration Tool appears and allows the remote user to safely configure printing on the remote system.

19.7.2. Port Forwarding

SSH can secure otherwise insecure TCP/IP protocols via port forwarding. When using this technique, the SSH server becomes an encrypted conduit to the SSH client.

Port forwarding works by mapping a local port on the client to a remote port on the server. SSH can map any port from the server to any port on the client; port numbers do not need to match for this technique to work.

To create a TCP/IP port forwarding channel which listens for connections on the localhost, use the following command:

```
ssh -L local-port:remote-hostname:remote-port username@hostname
```

Note

Setting up port forwarding to listen on ports below 1024 requires root level access.

To check email on a server called mail.example.com using POP3 through an encrypted connection, use the following command:

```
ssh -L 1100:mail.example.com:110 mail.example.com
```

Once the port forwarding channel is in place between the client machine and the mail server, direct a POP3 mail client to use port 1100 on the localhost to check for new mail. Any requests sent to port 1100 on the client system are directed securely to the mail.example.com server.

If mail.example.com is not running an SSH server, but another machine on the same network is, SSH can still be used to secure part of the connection. However, a slightly different command is necessary:

```
ssh -L 1100:mail.example.com:110 other.example.com
```

In this example, POP3 requests from port 1100 on the client machine are forwarded through the SSH connection on port 22 to the SSH server, other.example.com. Then, other.example.com connects to port 110 on mail.example.com to check for new mail. Note, when using this technique only the connection between the client system and other.example.com SSH server is secure.

Port forwarding can also be used to get information securely through network firewalls. If the firewall is configured to allow SSH traffic via its standard port (22) but blocks access to other ports, a connection between two hosts using the blocked ports is still possible by redirecting their communication over an established SSH connection.

Note

Using port forwarding to forward connections in this manner allows any user on the client system to connect to that service. If the client system becomes compromised, the attacker also has access to forwarded services.

System administrators concerned about port forwarding can disable this functionality on the server by specifying a No parameter for the AllowTcpForwarding line in /etc/ssh/sshd_config and restarting the sshd service.

19.7.3. 키 쌍 생성하기

If you do not want to enter your password every time you use ssh, scp, or sftp to connect to a remote machine, you can generate an authorization key pair.

각각의 사용자를 위한 키를 생성하셔야 합니다. 원격 컴퓨터에 접속하려는 사용자를 위한 키를 생성하시려면, 다음과 같은 단계를 따르십시오. 만일 루트로서 키를 생성하시면, 루트 사용자만이 그 키를 사용할 수 있습니다.

Starting with OpenSSH version 3.0, ~/.ssh/authorized_keys2, ~/.ssh/known_hosts2, and /etc/ssh_known_hosts2 are obsolete. SSH Protocol 1 and 2 share the ~/.ssh/authorized_keys, ~/.ssh/known_hosts, and /etc/ssh/ssh_known_hosts files.

Red Hat Enterprise Linux 5.8 uses SSH Protocol 2 and RSA keys by default.

Tip

If you reinstall and want to save your generated key pair, backup the .ssh directory in your home directory. After reinstalling, copy this directory back to your home directory. This process can be done for all users on your system, including root.

19.7.3.1. 2.0 버전에 사용되는 RSA 키 쌍 생성하기

SSH 프로토콜 2.0 버전에 사용되는 RSA 키 쌍을 생성하기 위해서는 다음과 같은 단계를 따르십시오. OpenSSH 2.9 이후 버전에서 2.0 버전 SSH 프로토콜이 기본으로 사용됩니다.

1. 2.0 버전 프로토콜과 함께 사용될 RSA 키 쌍을 생성하기 위해서는 쉘 프롬프트에서 다음과 같은 명령을 입력합니다:

   ```
   ssh-keygen -t rsa
   ```

   Accept the default file location of ~/.ssh/id_rsa. Enter a passphrase different from your account password and confirm it by entering it again.

   The public key is written to ~/.ssh/id_rsa.pub. The private key is written to ~/.ssh/id_rsa. Never distribute your private key to anyone.

2. Change the permissions of the .ssh directory using the following command:

```
chmod 755 ~/.ssh
```

3. Copy the contents of ~/.ssh/id_rsa.pub into the file ~/.ssh/authorized_keys on the machine to which you want to connect. If the file ~/.ssh/authorized_keys exist, append the contents of the file ~/.ssh/id_rsa.pub to the file ~/.ssh/authorized_keys on the other machine.

4. Change the permissions of the authorized_keys file using the following command:

```
chmod 644 ~/.ssh/authorized_keys
```

5. If you are running GNOME or are running in a graphical desktop with GTK2+ libraries installed, skip to 19.7.3.4절. "Configuring ssh-agent with a GUI" . If you are not running the X Window System, skip to 19.7.3.5절. "Configuring ssh-agent" .

19.7.3.2. 2.0 버전에 사용되는 DSA 키 쌍 생성하기

SSH 프로토콜 2.0 버전에 사용될 DSA 키 쌍을 생성하기 위해서는 다음과 같은 단계를 따르십시오.

1. 2.0 버전 프로토콜과 함께 사용될 DSA 키 쌍을 생성하기 위해서는 쉘 프롬프트에서 다음과 같은 명령을 입력합니다:

```
ssh-keygen -t dsa
```

Accept the default file location of ~/.ssh/id_dsa. Enter a passphrase different from your account password and confirm it by entering it again.

Tip

암호 문구는 사용자 인증을 위해 사용되는 문자열로서 단어와 문자로 이루어 졌습니다. 문자열은 스페이스와 탭을 사용할 수 있다는 점에서 암호와 차이가 있습니다. 암호 문구는 한 단어가 아닌 하나의 문구로 이루어졌기 때문에 보통 암호보다 더 깁니다.

The public key is written to ~/.ssh/id_dsa.pub. The private key is written to ~/.ssh/id_dsa. It is important never to give anyone the private key.

2. Change the permissions of the .ssh directory with the following command:

```
chmod 755 ~/.ssh
```

3. Copy the contents of ~/.ssh/id_dsa.pub into the file ~/.ssh/authorized_keys on the machine to which you want to connect. If the file ~/.ssh/authorized_keys exist, append the contents of the file ~/.ssh/id_dsa.pub to the file ~/.ssh/authorized_keys on the other machine.

4. Change the permissions of the authorized_keys file using the following command:

```
chmod 644 ~/.ssh/authorized_keys
```

5. If you are running GNOME or a graphical desktop environment with the GTK2+ libraries installed, skip to 19.7.3.4절. “Configuring ssh-agent with a GUI” . If you are not running the X Window System, skip to 19.7.3.5절. “Configuring ssh-agent” .

19.7.3.3. 1.3 과 1.5 버전에 사용되는 RSA 키 쌍 생성하기

SSH 프로토콜 1.0 버전에서 사용되는 RSA 키 쌍을 생성하기 위해서는 다음과 같은 단계를 따르십시오. 단순히 DSA를 사용하는 시스템들을 연결하는 경우라면 RSA 1.3 버전과 RSA 1.5 버전 키 쌍이 필요하지 않습니다.

1. RSA (버전 1.3과 1.5 프로토콜) 키 쌍을 생성하기 위해서는 쉘 프롬프트 상에서 다음과 같은 명령을 입력하십시오:

   ```
   ssh-keygen -t rsa1
   ```

   Accept the default file location (~/.ssh/identity). Enter a passphrase different from your account password. Confirm the passphrase by entering it again.

   The public key is written to ~/.ssh/identity.pub. The private key is written to ~/.ssh/identity. Do not give anyone the private key.

2. Change the permissions of your .ssh directory and your key with the commands chmod 755 ~/.ssh and chmod 644 ~/.ssh/identity.pub.

3. Copy the contents of ~/.ssh/identity.pub into the file ~/.ssh/authorized_keys on the machine to which you wish to connect. If the file ~/.ssh/authorized_keys does not exist, you can copy the file ~/.ssh/identity.pub to the file ~/.ssh/authorized_keys on the remote machine.

4. If you are running GNOME, skip to 19.7.3.4절. “Configuring ssh-agent with a GUI” . If you are not running GNOME, skip to 19.7.3.5절. “Configuring ssh-agent” .

19.7.3.4. Configuring ssh-agent with a GUI

The ssh-agent utility can be used to save your passphrase so that you do not have to enter it each time you initiate an ssh or scp connection. If you are using GNOME, the gnome-ssh-askpass package contains the application used to prompt you for your passphrase when you log in to GNOME and save it until you log out of GNOME. You will not have to enter your password or passphrase for any ssh or scp connection made during that GNOME session. If you are not using GNOME, refer to 19.7.3.5절. “Configuring ssh-agent” .

GNOME 세션을 사용하는 동안 암호 문구를 저장하기 위해서는 다음의 단계를 따르십시오

1. You will need to have the package gnome-ssh-askpass installed; you can use the command rpm -q openssh-askpass to determine if it is installed or not. If it is not installed, install it from your Red Hat Enterprise Linux CD-ROM set, from a Red Hat FTP mirror site, or using Red Hat Network.

2. Select Main Menu Button (on the Panel) > Preferences > More Preferences > Sessions, and click on the Startup Programs tab. Click Add and enter /usr/bin/ssh-add in the Startup Command text area. Set it a priority to a number higher than any existing commands to ensure that it

is executed last. A good priority number for ssh-add is 70 or higher. The higher the priority number, the lower the priority. If you have other programs listed, this one should have the lowest priority. Click Close to exit the program.

3. Log out and then log back into GNOME; in other words, restart X. After GNOME is started, a dialog box will appear prompting you for your passphrase(s). Enter the passphrase requested. If you have both DSA and RSA key pairs configured, you will be prompted for both. From this point on, you should not be prompted for a password by ssh, scp, or sftp.

19.7.3.5. Configuring ssh-agent

The ssh-agent can be used to store your passphrase so that you do not have to enter it each time you make a ssh or scp connection. If you are not running the X Window System, follow these steps from a shell prompt. If you are running GNOME but you do not want to configure it to prompt you for your passphrase when you log in (refer to 19.7.3.4절. "Configuring ssh-agent with a GUI"), this procedure will work in a terminal window, such as an XTerm. If you are running X but not GNOME, this procedure will work in a terminal window. However, your passphrase will only be remembered for that terminal window; it is not a global setting.

1. 쉘 프롬프트에서 다음의 명령을 입력하십시오:

   ```
   exec /usr/bin/ssh-agent $SHELL
   ```

2. 그 후 다음의 명령을 입력합니다:

   ```
   ssh-add
   ```

   그리고 암호 문구를 입력하십시오. 한개 이상의 키 쌍이 설정되었다면, 각각의 키 쌍에 대한 암호 문구가 요구될 것입니다.

3. 로그 아웃 후에는 암호 문구는 더 이상 기억되지 않습니다. 가상 콘솔에 로그인할 때마다 또는 터미널 창을 열 때마다 앞에서 언급된 두 명령을 실행하셔야 합니다.

19.8. 추가 자료

OpenSSH와 OpenSSL 프로젝트는 계속적으로 개발 중이며 따라서 가장 최근의 정보는 웹사이트에서 찾으실 수 있습니다. OpenSSH와 OpenSSL 도구에 대한 메뉴얼 페이지도 자세한 정보를 찾기 위한 좋은 자료가 될 수 있습니다.

19.8.1. 설치된 문서 자료

- The ssh, scp, sftp, sshd, and ssh-keygen man pages — These man pages include information on how to use these commands as well as all the parameters that can be used with them.

19.8.2. 유용한 웹사이트

- http://www.openssh.com/ — OpenSSH FAQ 페이지, 버그 리포트, 메일링 리스트, 프로젝트 목표와 보안 기능에 대한 더욱 기술적인 설명을 찾으실 수 있습니다.

- http://www.openssl.org/ — OpenSSL FAQ 페이지, 메일링 리스트와 프로젝트 목표에 대한 설명을 찾으실 수 있습니다.

- http://www.freesshd.com/ — SSH client software for other platforms.

네트워크 파일 시스템 (NFS)

A Network File System (NFS) allows remote hosts to mount file systems over a network and interact with those file systems as though they are mounted locally. This enables system administrators to consolidate resources onto centralized servers on the network.

This chapter focuses on fundamental NFS concepts and supplemental information.

20.1. How It Works

Currently, there are three versions of NFS. NFS version 2 (NFSv2) is older and is widely supported. NFS version 3 (NFSv3) has more features, including 64bit file handles, Safe Async writes and more robust error handling. NFS version 4 (NFSv4) works through firewalls and on the Internet, no longer requires portmapper, supports ACLs, and utilizes stateful operations. Red Hat Enterprise Linux supports NFSv2, NFSv3, and NFSv4 clients, and when mounting a file system via NFS, Red Hat Enterprise Linux uses NFSv3 by default, if the server supports it.

All versions of NFS can use Transmission Control Protocol (TCP) running over an IP network, with NFSv4 requiring it. NFSv2 and NFSv3 can use the User Datagram Protocol (UDP) running over an IP network to provide a stateless network connection between the client and server.

When using NFSv2 or NFSv3 with UDP, the stateless UDP connection under normal conditions has less Protocol overhead than TCP which can translate into better performance on very clean, non-congested networks. The NFS server sends the client a file handle after the client is authorized to access the shared volume. This file handle is an opaque object stored on the server's side and is passed along with RPC requests from the client. The NFS server can be restarted without affecting the clients and the cookie remains intact. However, because UDP is stateless, if the server goes down unexpectedly, UDP clients continue to saturate the network with requests for the server. For this reason, TCP is the preferred protocol when connecting to an NFS server.

Because protocol support has been incorporated into the v4 protocol, NFSv4 has no interaction with the portmap, rpc.lockd, and rpc.statd daemons. NFSv4 listens on the well-known TCP port 2049, which eliminates the need for portmap interaction. The mounting and locking protocols have been incorporated into the V4 protocol which eliminates the need for interaction with rpc.lockd and rpc.statd. The rpc.mountd daemon is still required on the server, but is not involved in any over-the-wire operations.

참고

TCP is the default transport protocol for NFS under Red Hat Enterprise Linux. UDP can be used for compatibility purposes as needed, but is not recommended for wide usage.

All the RPC/NFS daemon have a -p command line option that can set the port, making firewall configuration easier.

After the client is granted access by TCP wrappers, the NFS server refers to its configuration file, /etc/exports, to determine whether the client is allowed to access any of the exported file systems. Once access is granted, all file and directory operations are available to the user.

중요

In order for NFS to work with a default installation of Red Hat Enterprise Linux with a firewall enabled, IPTables with the default TCP port 2049 must be configured. Without proper IPTables configuration, NFS does not function properly.

The NFS initialization script and rpc.nfsd process now allow binding to any specified port during system start up. However, this can be error prone if the port is unavailable or conflicts with another daemon.

20.1.1. Required Services

Red Hat Enterprise Linux uses a combination of kernel-level support and daemon processes to provide NFS file sharing. All NFS versions rely on Remote Procedure Calls (RPC) between clients and servers. RPC services under Linux are controlled by the portmap service. To share or mount NFS file systems, the following services work together, depending on which version of NFS is implemented:

- nfs — (/sbin/service nfs start) starts the NFS server and the appropriate RPC processes to service requests for shared NFS file systems.

- nfslock — (/sbin/service nfslock start) is a mandatory service that starts the appropriate RPC processes to allow NFS clients to lock files on the server.

- portmap — accepts port reservations from local RPC services. These ports are then made available (or advertised) so the corresponding remote RPC services access them. portmap responds to requests for RPC services and sets up connections to the requested RPC service.

The following RPC processes facilitate NFS services:

- rpc.mountd — This process receives mount requests from NFS clients and verifies the requested file system is currently exported. This process is started automatically by the nfs service and does not require user configuration.

- rpc.nfsd — Allows explicit NFS versions and protocols the server advertises to be defined. It works with the Linux kernel to meet the dynamic demands of NFS clients, such as providing server threads each time an NFS client connects. This process corresponds to the nfs service.

- rpc.lockd — allows NFS clients to lock files on the server. If rpc.lockd is not started, file locking will fail. rpc.lockd implements the Network Lock Manager (NLM) protocol. This process corresponds to the nfslock service. This is not used with NFSv4.

- rpc.statd — This process implements the Network Status Monitor (NSM) RPC protocol which notifies NFS clients when an NFS server is restarted without being gracefully brought down. This process is started automatically by the nfslock service and does not require user configuration. This is not used with NFSv4.

- rpc.rquotad — This process provides user quota information for remote users. This process is started automatically by the nfs service and does not require user configuration.

- rpc.idmapd — This process provides NFSv4 client and server upcalls which map between on-the-wire NFSv4 names (which are strings in the form of user@domain) and local UIDs and GIDs. For idmapd to function with NFSv4, the /etc/idmapd.conf must be configured. This service is required for use with NFSv4.

To use NFS on your system, make sure you have the nfs-utils, nfs-utils-lib, and portmap packages installed.

20.2. NFS Client Configuration

NFS shares are mounted on the client side using the mount command. The format of the command is as follows:

```
mount -t <nfs-type> -o <options> <host>:</remote/export> </local/directory>
```

Replace <nfs-type> with either nfs for NFSv2 or NFSv3 servers, or nfs4 for NFSv4 servers. Replace <options> with a comma separated list of options for the NFS file system (refer to 20.4 절. "Common NFS Mount Options" for details). Replace <host> with the remote host, </remote/export> with the remote directory being mounted, and </local/directory> with the local directory where the remote file system is to be mounted.

Refer to the mount man page for more details.

If accessing an NFS share by manually issuing the mount command, the file system must be remounted manually after the system is rebooted. Red Hat Enterprise Linux offers two methods for mounting remote file systems automatically at boot time: the /etc/fstab file or the autofs service.

20.2.1. Mounting NFS File Systems using /etc/fstab

An alternate way to mount an NFS share from another machine is to add a line to the /etc/fstab file. The /etc/fstab file is referenced by the netfs service at boot time, so lines referencing NFS shares have the same effect as manually typing the mount command during the boot process. Each line in this file must state the hostname of the NFS server, the directory on the server being exported, and the directory on the local machine where the NFS share is to be mounted. You must be root to modify the /etc/fstab file.

The general syntax for a line in /etc/fstab is as follows:

```
<server>:</remote/export> </local/directory> <nfs-type> <options> 0 0
```

Replace <server> with the hostname, IP address, or fully qualified domain name of the server exporting the file system. Replace </remote/export> with the path to the exported directory, and </local/directory> with the local file system on which the exported directory is mounted. Replace <nfs-type> with either nfs for NFSv2 or NFSv3 servers, or nfs4 for NFSv4 servers. Finally, replace <options> with a comma separated list of options for the NFS file system (see 20.4절. "Common NFS Mount Options" for details). Note that the mount point must exist before /etc/fstab is read, otherwise the mount fails.

The following is a sample /etc/fstab line to mount an NFS export:

```
server:/usr/local/pub    /pub    nfs    defaults 0 0
```

After adding this line to /etc/fstab on the client system, type the command mount /pub at a shell prompt, and the mount point /pub is mounted from the server. The mount point /pub must exist on the client machine before this command can be executed.

For more information about the /etc/fstab configuration file and its contents, refer to the fstab manual page.

20.3. autofs

One drawback to using /etc/fstab is that, regardless of how infrequently a user accesses the NFS mounted file system, the system must dedicate resources to keep the mounted file system in place. This is not a problem with one or two mounts, but when the system is maintaining mounts to many systems at one time, overall system performance can be affected. An alternative to /etc/fstab is to use the kernel-based automount utility. An automounter consists of two components. One is a kernel module that implements a file system, while the other is a user-space daemon that performs all of the other functions. The automount utility can mount and unmount NFS file systems automatically (on demand mounting) therefore saving system resources. The automount utility can be used to mount other file systems including AFS, SMBFS, CIFS and local file systems.

autofs uses /etc/auto.master (master map) as its default primary configuration file. This can be changed to use another supported network source and name using the autofs configuration (in /etc/sysconfig/autofs) in conjunction with the Name Service Switch mechanism. An instance of the version 4 daemon was run for each mount point configured in the master map and so it could be run manually from the command line for any given mount point. This is not possible with version 5 because it uses a single daemon to manage all configured mount points, so all automounts must be configured in the master map. This is in line with the usual requirements of other industry standard automounters. Mount point, hostname, exported directory, and options can all be specified in a set of files (or other supported network sources) rather than configuring them manually for each host. Please ensure that you have the autofs package installed if you wish to use this service.

20.3.1. What's new in autofs version 5?

Direct map support

Autofs direct maps provide a mechanism to automatically mount file systems at arbitrary points in the file system hierarchy. A direct map is denoted by a mount point of "/-" in the master map. Entries in a direct map contain an absolute path name as a key (instead of the relative path names used in indirect maps).

Lazy mount and unmount support

Multimount map entries describe a hierarchy of mount points under a single key. A good example of this is the "-hosts" map, commonly used for automounting all exports from a host under "/net/<host>" as a multi-mount map entry. When using the "-hosts" map, an 'ls' of "/net/<host>" will mount autofs trigger mounts for each export from <host> and mount and expire them as they are accessed. This can greatly reduce the number of active mounts needed when accessing a server with a large number of exports.

Enhanced LDAP support

The Lightweight Directory Access Protocol, or LDAP, support in autofs version 5 has been enhanced in several ways with respect to autofs version 4. The autofs configuration file (/etc/sysconfig/autofs) provides a mechanism to specify the autofs schema that a site implements, thus precluding the need to determine this via trial and error in the application itself. In addition, authenticated binds to the LDAP server are now supported, using most mechanisms supported by the common LDAP server implementations. A new configuration file has been added for this support: /etc/autofs_ldap_auth.conf. The default configuration file is self-documenting, and uses an XML format.

Proper use of the Name Service Switch (nsswitch) configuration.

The Name Service Switch configuration file exists to provide a means of determining from where specific configuration data comes. The reason for this configuration is to allow administrators the flexibility of using the back-end database of choice, while maintaining a uniform software interface to access the data. While the version 4 automounter is becoming increasingly better at handling the name service switch configuration, it is still not complete. Autofs version 5, on the other hand, is a complete implementation. See the manual page for nsswitch.conf for more information on the supported syntax of this file. Please note that not all nss databases are valid map sources and the parser will reject ones that are invalid. Valid sources are files, yp, nis, nisplus, ldap and hesiod.

Multiple master map entries per autofs mount point

One thing that is frequently used but not yet mentioned is the handling of multiple master map entries for the direct mount point "/-". The map keys for each entry are merged and behave as one map.

An example is seen in the connectathon test maps for the direct mounts below:

```
/- /tmp/auto_dcthon
/- /tmp/auto_test3_direct
/- /tmp/auto_test4_direct
```

20.3.2. autofs Configuration

The primary configuration file for the automounter is /etc/auto.master, also referred to as the master map which may be changed as described in the introduction section above. The master map lists autofs-controlled mount points on the system, and their corresponding configuration files or network sources known as automount maps. The format of the master map is as follows:

```
<mount-point> <map-name> <options>
```

where:

- mount-point is the autofs mount point such as /home.
- map-name is the name of a map source which contains a list of mount points, and the file system location from which those mount points should be mounted. The syntax for a map entry is described below.
- options if supplied, will apply to all entries in the given map provided they don't themselves have options specified. This behavior is different from autofs version 4 where the options where cumulative. This has been changed to meet our primary goal of mixed environment compatibility.

The following is a sample /etc/auto.master file:

```
~]$ cat /etc/auto.master
/home /etc/auto.misc
```

The general format of maps is similar to the master map, however the "options" appear between the mount point and the location instead of at the end of the entry as in the master map:

```
<mount-point>   [<options>]   <location>
```

where:

- <mount-point> is the autofs mount point. This can be a single directory name for an indirect mount or the full path of the mount point for direct mounts. Each direct and indirect map entry key (<mount-point> above) may be followed by a space separated list of offset directories (sub directory names each beginning with a "/") making them what is known as a mutli-mount entry.

- <options> if supplied, are the mount options for the map entries that do not specify their own options.

- <location> is the file system location such as a local file system path (preceded with the Sun map format escape character ":" for map names beginning with "/"), an NFS file system or other valid file system location.

The following is a sample map file:

```
~]$ cat /etc/auto.misc
payroll -fstype=nfs personnel:/dev/hda3
sales -fstype=ext3 :/dev/hda4
```

The first column in a map file indicates the autofs mount point (sales and payroll from the server called personnel). The second column indicates the options for the autofs mount while the third column indicates the source of the mount. Following the above configuration, the autofs mount points will be /home/payroll and /home/sales. The -fstype= option is often omitted and is generally not needed for correct operation.

The automounter will create the directories if they do not exist. If the directories exist before the automounter was started, the automounter will not remove them when it exits. You can start or restart the automount daemon by issuing the following command:

```
service autofs start
```

or

```
service autofs restart
```

Using the above configuration, if a process requires access to an autofs unmounted directory such as /home/payroll/2006/July.sxc, the automount daemon automatically mounts the directory. If a timeout is specified, the directory will automatically be unmounted if the directory is not accessed for the timeout period.

You can view the status of the automount daemon by issuing the following command in your terminal:

```
/sbin/service/autofs status
```

20.3.3. autofs Common Tasks

20.3.3.1. Overriding or augmenting site configuration files

It can be useful to override site defaults for a specific mount point on a client system. For example, assuming that the automounter maps are stored in NIS and the /etc/nsswitch.conf file has the following directive:

```
automount:    files nis
```

and the NIS auto.master map file contains the following:

```
/home auto.home
```

Also assume the NIS auto.home map contains the following:

```
beth      fileserver.example.com:/export/home/beth
joe       fileserver.example.com:/export/home/joe
*         fileserver.example.com:/export/home/&
```

and the file map /etc/auto.home does not exist.

For the above example, lets assume that the client system needs to mount home directories from a different server. In this case, the client will need to use the following /etc/auto.master map:

```
/home /etc/auto.home2
+auto.master
```

And the /etc/auto.home2 map contains the entry:

```
*    labserver.example.com:/export/home/&
```

Because only the first occurrence of a mount point is processed, /home will contain the contents of /etc/auto.home2 instead of the NIS auto.home map.

Alternatively, if you just want to augment the site-wide auto.home map with a few entries, create a /etc/auto.home file map, and in it put your new entries and at the end, include the NIS auto.home map. Then the /etc/auto.home file map might look similar to:

```
mydir someserver:/export/mydir
+auto.home
```

Given the NIS auto.home map listed above, an ls of /home would now give:

```
~]$ ls /home
beth   joe   mydir
```

This last example works as expected because autofs knows not to include the contents of a file map of the same name as the one it is reading and so moves on to the next map source in the nsswitch configuration.

20.3.3.2. Using LDAP to Store Automounter Maps

LDAP client libraries must be installed on all systems which are to retrieve automounter maps from LDAP. On RHEL 5, the openldap package should be installed automatically as a dependency of the automounter. To configure LDAP access, modify /etc/openldap/ldap.conf. Ensure that BASE and URI are set appropriately for your site. Please also ensure that the schema is set in the configuration.

The most recently established schema for storing automount maps in LDAP is described by rfc2307bis. To use this schema it is necessary to set it in the autofs configuration (/etc/sysconfig/autofs) by removing the comment characters from the schema definition. For example:

```
DEFAULT_MAP_OBJECT_CLASS="automountMap"
DEFAULT_ENTRY_OBJECT_CLASS="automount"
DEFAULT_MAP_ATTRIBUTE="automountMapName"
DEFAULT_ENTRY_ATTRIBUTE="automountKey"
DEFAULT_VALUE_ATTRIBUTE="automountInformation"
```

Ensure that these are the only schema entries not commented in the configuration. Please also note that the automountKey replaces the cn attribute in the rfc2307bis schema. An LDIF of a sample configuration is described below:

```
# extended LDIF
#
# LDAPv3
# base <> with scope subtree
# filter: (&(objectclass=automountMap)(automountMapName=auto.master))
# requesting: ALL
#

# auto.master, example.com
dn: automountMapName=auto.master,dc=example,dc=com
objectClass: top
objectClass: automountMap
automountMapName: auto.master

# extended LDIF
#
# LDAPv3
# base <automountMapName=auto.master,dc=example,dc=com> with scope subtree
# filter: (objectclass=automount)
# requesting: ALL
#

# /home, auto.master, example.com
dn: automountMapName=auto.master,dc=example,dc=com
objectClass: automount
cn: /home

automountKey: /home
automountInformation: auto.home

# extended LDIF
#
# LDAPv3
# base <> with scope subtree
# filter: (&(objectclass=automountMap)(automountMapName=auto.home))
# requesting: ALL
#

# auto.home, example.com
dn: automountMapName=auto.home,dc=example,dc=com
objectClass: automountMap
automountMapName: auto.home

# extended LDIF
#
# LDAPv3
# base <automountMapName=auto.home,dc=example,dc=com> with scope subtree
# filter: (objectclass=automount)
# requesting: ALL
#

# foo, auto.home, example.com
dn: automountKey=foo,automountMapName=auto.home,dc=example,dc=com
objectClass: automount
automountKey: foo
automountInformation: filer.example.com:/export/foo

# /, auto.home, example.com
dn: automountKey=/,automountMapName=auto.home,dc=example,dc=com
objectClass: automount
automountKey: /
automountInformation: filer.example.com:/export/&
```

20.3.3.3. Adapting Autofs v4 Maps To Autofs v5

v4 Multi-map entries

Autofs version 4 introduced the notion of a multi-map entry in the master map. A multi-map entry is of the form:

```
<mount-point> <maptype1> <mapname1> <options1> -- <maptype2> <mapname2> <options2> -- ...
```

Any number of maps can be combined into a single map in this manner. This feature is no longer present in v5. This is because Version 5 supports included maps which can be used to attain the same results. Consider the following multi-map example:

```
/home file /etc/auto.home -- nis auto.home
```

This can be replaced by the following configuration for v5:

/etc/nsswitch.conf must list:

```
automount: files nis
```

/etc/auto.master should contain:

```
/home   auto.home
```

/etc/auto.home should contain:

```
<entries for the home directory>
+auto.home
```

In this way, the entries from /etc/auto.home and the nis auto.home map are combined.

Multiple master maps

In autofs version 4, it is possible to merge the contents of master maps from each source, such as files, nis, hesiod, and LDAP. The version 4 automounter looks for a master map for each of the sources listed in /etc/nsswitch.conf. The map is read if it exists and its contents are merged into one large auto.master map.

In version 5, this is no longer the behaviour. Only the first master map found from the list of sources in nsswitch.conf is consulted. If it is desirable to merge the contents of multiple master maps, included maps can be used. Consider the following example:

```
/etc/nsswitch.conf:
automount: files nis
```

```
/etc/auto.master:
/home   /etc/auto.home
+auto.master
```

The above configuration will merge the contents of the file-based auto.master and the NIS-based auto.master. However, because included map entries are only allowed in file maps, there is no way to include both an NIS auto.master and an LDAP auto.master.

This limitation can be overcome by creating a master maps that have a different name in the source. In the example above if we had an LDAP master map named auto.master.ldap we could also add "+auto.master.ldap" to the file based master map and provided that "ldap" is listed as a source in our nsswitch configuration it would also be included.

20.4. Common NFS Mount Options

Beyond mounting a file system via NFS on a remote host, other options can be specified at the time of the mount to make it easier to use. These options can be used with manual mount commands, /etc/fstab settings, and autofs.

The following are options commonly used for NFS mounts:

- hard or soft — Specifies whether the program using a file via an NFS connection should stop and wait (hard) for the server to come back online, if the host serving the exported file system is unavailable, or if it should report an error (soft).

  If hard is specified, the user cannot terminate the process waiting for the NFS communication to resume unless the intr option is also specified.

  If soft is specified, the user can set an additional timeo=<value> option, where <value> specifies the number of seconds to pass before the error is reported.

> **참고**
>
> Using soft mounts is not recommended as they can generate I/O errors in very congested networks or when using a very busy server.

- intr — Allows NFS requests to be interrupted if the server goes down or cannot be reached.

- nfsvers=2 or nfsvers=3 — Specifies which version of the NFS protocol to use. This is useful for hosts that run multiple NFS servers. If no version is specified, NFS uses the highest supported version by the kernel and mount command. This option is not supported with NFSv4 and should not be used.

- noacl — Turns off all ACL processing. This may be needed when interfacing with older versions of Red Hat Enterprise Linux, Red Hat Linux, or Solaris, since the most recent ACL technology is not compatible with older systems.

- nolock — Disables file locking. This setting is occasionally required when connecting to older NFS servers.

- noexec — Prevents execution of binaries on mounted file systems. This is useful if the system is mounting a non-Linux file system via NFS containing incompatible binaries.

- nosuid — Disables set-user-identifier or set-group-identifier bits. This prevents remote users from gaining higher privileges by running a setuid program.

- port=num — Specifies the numeric value of the NFS server port. If num is 0 (the default), then mount queries the remote host's portmapper for the port number to use. If the remote host's NFS daemon is not registered with its portmapper, the standard NFS port number of TCP 2049 is used instead.

- rsize=num and wsize=num — These settings speed up NFS communication for reads (rsize) and writes (wsize) by setting a larger data block size, in bytes, to be transferred at one time. Be careful when changing these values; some older Linux kernels and network cards do not work well with larger block sizes. For NFSv2 or NFSv3, the default values for both parameters is set to 8192. For NFSv4, the default values for both parameters is set to 32768.

- sec=mode — Specifies the type of security to utilize when authenticating an NFS connection.

  sec=sys is the default setting, which uses local UNIX UIDs and GIDs by means of AUTH_SYS to authenticate NFS operations.

  sec=krb5 uses Kerberos V5 instead of local UNIX UIDs and GIDs to authenticate users.

  sec=krb5i uses Kerberos V5 for user authentication and performs integrity checking of NFS operations using secure checksums to prevent data tampering.

  sec=krb5p uses Kerberos V5 for user authentication, integrity checking, and encrypts NFS traffic to prevent traffic sniffing. This is the most secure setting, but it also has the most performance overhead involved.

- tcp — Specifies for the NFS mount to use the TCP protocol.

- udp — Specifies for the NFS mount to use the UDP protocol.

Many more options are listed on the mount and nfs man pages.

20.5. Starting and Stopping NFS

To run an NFS server, the portmap service must be running. To verify that portmap is active, type the following command as root:

```
service portmap status
```

If the portmap service is running, then the nfs service can be started. To start an NFS server, as root type:

```
service nfs start
```

참고

nfslock also has to be started for both the NFS client and server to function properly. To start NFS locking as root type: /sbin/service nfslock start. If NFS is set to start at boot, please ensure that nfslock also starts by running chkconfig --list nfslock. If nfslock is not set to on, this implies that you will need to manually run the /sbin/service nfslock start each time the computer starts. To set nfslock to automatically start on boot, type the following command in a terminal chkconfig nfslock on.

To stop the server, as root, type:

```
service nfs stop
```

The restart option is a shorthand way of stopping and then starting NFS. This is the most efficient way to make configuration changes take effect after editing the configuration file for NFS.

To restart the server, as root, type:

```
service nfs restart
```

The condrestart (conditional restart) option only starts nfs if it is currently running. This option is useful for scripts, because it does not start the daemon if it is not running.

To conditionally restart the server, as root, type:

```
service nfs condrestart
```

To reload the NFS server configuration file without restarting the service, as root, type:

```
service nfs reload
```

By default, the nfs service does not start automatically at boot time. To configure the NFS to start up at boot time, use an initscript utility, such as /sbin/chkconfig, /usr/sbin/ntsysv, or the Services Configuration Tool program. Refer to 17장. 서비스로의 접근 통제 for more information regarding these tools.

20.6. NFS Server Configuration

There are three ways to configure an NFS server under Red Hat Enterprise Linux: using the NFS Server Configuration Tool (system-config-nfs), manually editing its configuration file (/etc/exports), or using the /usr/sbin/exportfs command.

To use the NFS Server Configuration Tool, you must be running X Windows, have root privileges, and have the system-config-nfs RPM package installed. To start the application, click on System > Administration > Server Settings > NFS. You can also type the command system-config-nfs in a terminal. The NFS Server Configuration tool window is illustrated below.

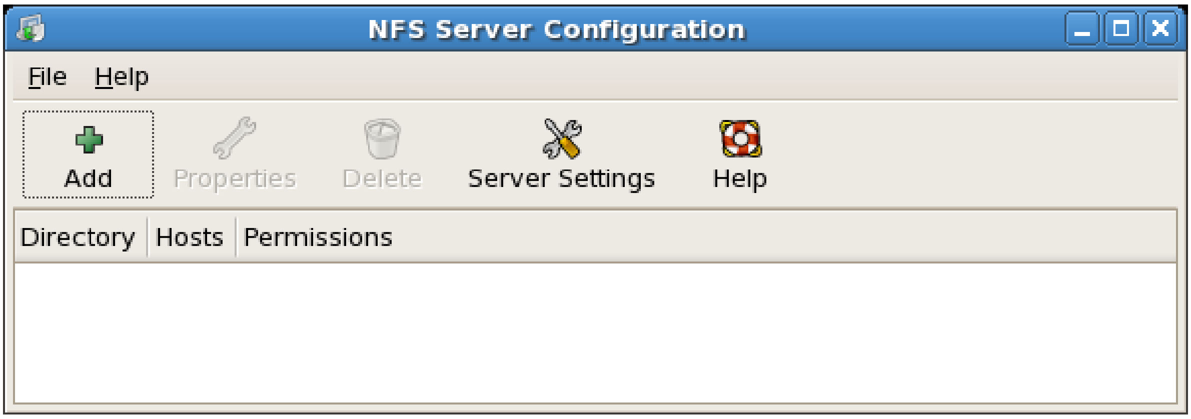

그림 20.1. NFS Server Configuration Tool

Based on certain firewall settings, you may need to configure the NFS daemon processes to use specific networking ports. The NFS server settings allows you to specify the ports for each process instead of using the random ports assigned by the portmapper. You can set the NFS Server settings by clicking on the Server Settings button. The figure below illustrates the NFS Server Settings window.

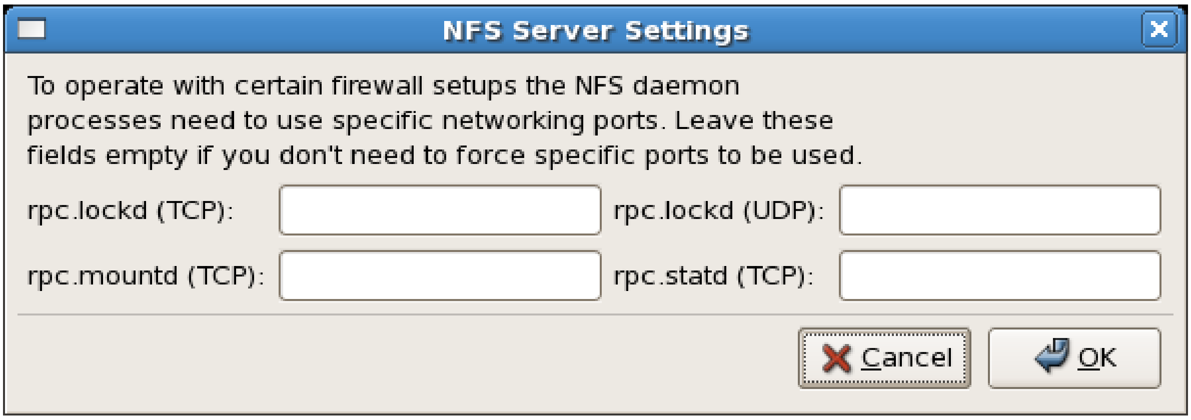

To operate with certain firewall setups the NFS daemon processes need to use specific networking ports. Leave these fields empty if you don't need to force specific ports to be used.

rpc.lockd (TCP):

rpc.lockd (UDP):

rpc.mountd (TCP):

rpc.statd (TCP):

그림 20.2. NFS Server Settings

20.6.1. Exporting or Sharing NFS File Systems

Sharing or serving files from an NFS server is known as exporting the directories. The NFS Server Configuration Tool can be used to configure a system as an NFS server.

To add an NFS share, click the Add button. The dialog box shown in 그림 20.3. “공유 추가” appears.

The Basic tab requires the following information:

- Directory — Specify the directory to share, such as /tmp.
- Host(s) — Specify the host(s) with which to share the directory. Refer to 20.6.4절. “호스트명 형식” for an explanation of possible formats.
- Basic permissions — Specify whether the directory should have read-only or read/write permissions.

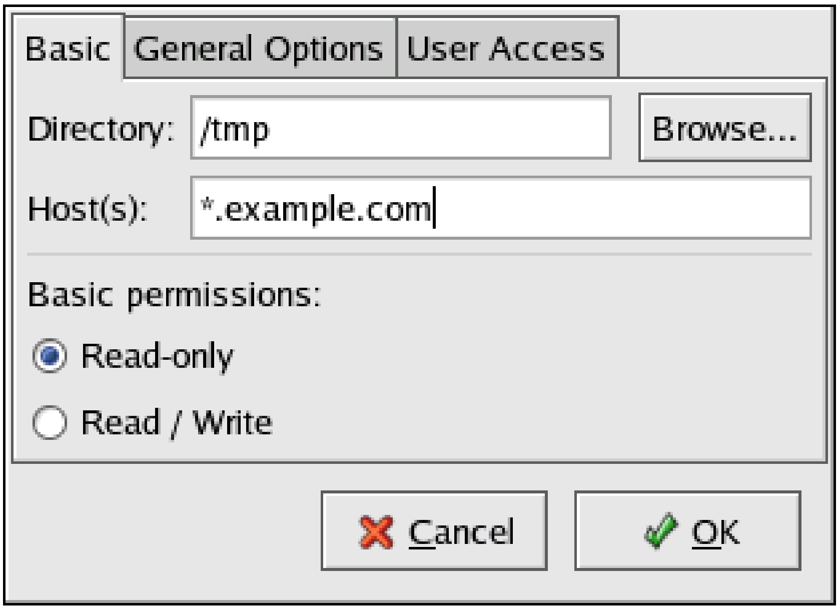

Directory: /tmp

Host(s): *.example.com

Basic permissions:

● Read-only

○ Read / Write

그림 20.3. 공유 추가

The General Options tab allows the following options to be configured:

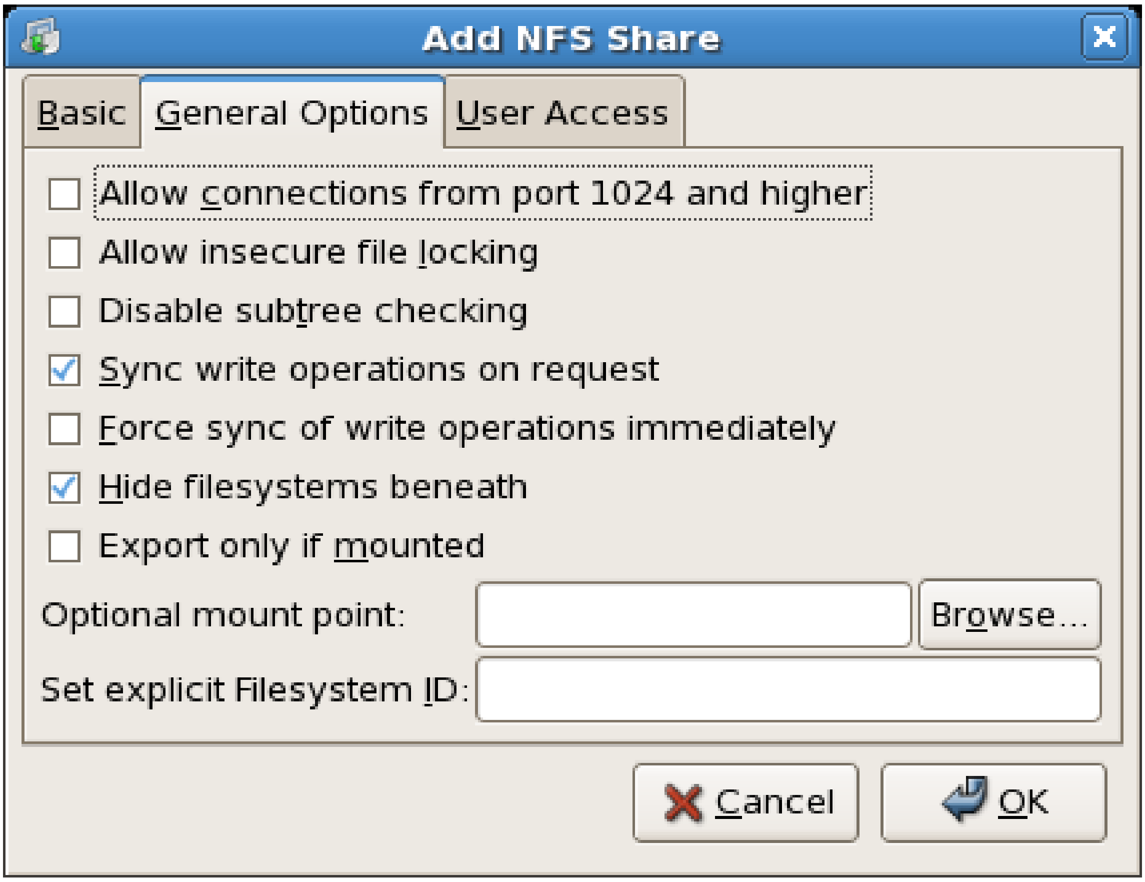

□ Allow connections from port 1024 and higher

□ Allow insecure file locking

□ Disable subtree checking

☑ Sync write operations on request

□ Force sync of write operations immediately

☑ Hide filesystems beneath

□ Export only if mounted

Optional mount point:

Set explicit Filesystem ID:

그림 20.4. NFS General Options

- Allow connections from port 1024 and higher — Services started on port numbers less than 1024 must be started as root. Select this option to allow the NFS service to be started by a user other than root. This option corresponds to insecure.

- Allow insecure file locking — Do not require a lock request. This option corresponds to insecure_locks.

- Disable subtree checking — If a subdirectory of a file system is exported, but the entire file system is not exported, the server checks to see if the requested file is in the subdirectory exported. This check is called subtree checking. Select this option to disable subtree checking. If the entire file system is exported, selecting to disable subtree checking can increase the transfer rate. This option corresponds to no_subtree_check.
- Sync write operations on request — Enabled by default, this option does not allow the server to reply to requests before the changes made by the request are written to the disk. This option corresponds to sync. If this is not selected, the async option is used.
  - Force sync of write operations immediately — Do not delay writing to disk. This option corresponds to no_wdelay.
- Hide filesystems beneath turns the nohide option on or off. When the nohide option is off, nested directories are revealed. The clients can therefore navigate through a filesystem from the parent without noticing any changes.
- Export only if mounted sets the mountpoint option which allows a directory to be exported only if it has been mounted.
- Optional Mount Point specifies the path to an optional mount point. Click on the Browse to navigate to the preferred mount point or type the path if known.
- Set explicit Filesystem ID: sets the fsid=X option. This is mainly used in a clustered setup. Using a consistent filesystem ID in all clusters avoids having stale NFS filehandles.

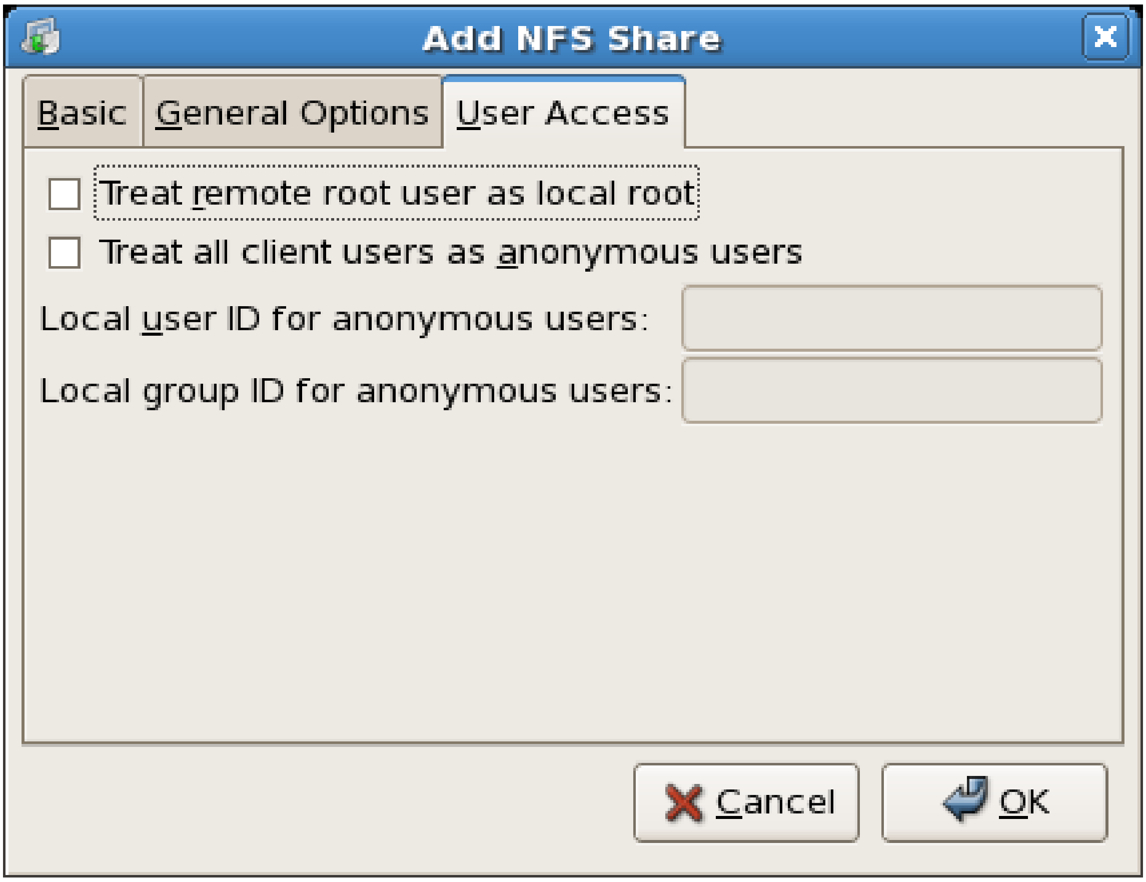

□ Treat remote root user as local root

□ Treat all client users as anonymous users

Local user ID for anonymous users:

Local group ID for anonymous users:

그림 20.5. NFS User Access

The User Access tab allows the following options to be configured:

- Treat remote root user as local root — By default, the user and group IDs of the root user are both 0. Root squashing maps the user ID 0 and the group ID 0 to the user and group IDs of anonymous so that root on the client does not have root privileges on the NFS server. If this option is selected, root is not mapped to anonymous, and root on a client has root privileges to exported directories. Selecting this option can greatly decrease the security of the system. Do not select it unless it is absolutely necessary. This option corresponds to no_root_squash.

- Treat all client users as anonymous users — If this option is selected, all user and group IDs are mapped to the anonymous user. This option corresponds to all_squash.

  - Specify local user ID for anonymous users — If Treat all client users as anonymous users is selected, this option lets you specify a user ID for the anonymous user. This option corresponds to anonuid.

  - Specify local group ID for anonymous users — If Treat all client users as anonymous users is selected, this option lets you specify a group ID for the anonymous user. This option corresponds to anongid.

To edit an existing NFS share, select the share from the list, and click the Properties button. To delete an existing NFS share, select the share from the list, and click the Delete button.

After clicking OK to add, edit, or delete an NFS share from the list, the changes take place immediately — the server daemon is restarted and the old configuration file is saved as /etc/exports.bak. The new configuration is written to /etc/exports.

The NFS Server Configuration Tool reads and writes directly to the /etc/exports configuration file. Thus, the file can be modified manually after using the tool, and the tool can be used after modifying the file manually (provided the file was modified with correct syntax).

The next this section discusses manually editing /etc/exports and using the /usr/sbin/exportfs command to export NFS file systems.

20.6.2. 명령행 설정

텍스트 편집기로 설정 파일을 편집하는 것을 선호하시거나 X 윈도우 시스템이 설치되어 있지 않은 경우에는, 설정 파일을 직접 수정하실 수 있습니다.

The /etc/exports file controls what directories the NFS server exports. Its format is as follows:

```
directory hostname(options)
```

The only option that needs to be specified is one of sync or async (sync is recommended). If sync is specified, the server does not reply to requests before the changes made by the request are written to the disk.

For example,

```
/misc/export speedy.example.com(sync)
```

would allow users from speedy.example.com to mount /misc/export with the default read-only permissions, but,

```
/misc/export speedy.example.com(rw,sync)
```

would allow users from speedy.example.com to mount /misc/export with read/write privileges.

Refer to 20.6.4절. "호스트명 형식" for an explanation of possible hostname formats.

경고

Be careful with spaces in the /etc/exports file. If there are no spaces between the hostname and the options in parentheses, the options apply only to the hostname. If there is a space between the hostname and the options, the options apply to the rest of the world. For example, examine the following lines:

```
/misc/export speedy.example.com(rw,sync) /misc/export speedy.example.com (rw,sync)
```

The first line grants users from speedy.example.com read-write access and denies all other users. The second line grants users from speedy.example.com read-only access (the default) and allows the rest of the world read-write access.

Each time you change /etc/exports, you must inform the NFS daemon of the change, or reload the configuration file with the following command:

```
service nfs reload
```

20.6.3. Running NFS Behind a Firewall

Because NFS requires portmap, which dynamically assigns ports for RPC services and can cause problems for configuring firewall rules, you can edit the /etc/sysconfig/nfs configuration file to control which ports the required RPC services run on. Refer to and read 30.1.22절. "/etc/sysconfig/nfs" for instructions on how to configure a firewall to allow NFS.

20.6.4. 호스트명 형식

호스트는 다음과 같은 형식으로 지정 가능합니다:

- Single machine — A fully qualified domain name (that can be resolved by the server), hostname (that can be resolved by the server), or an IP address.

- Series of machines specified with wildcards — Use the * or ? character to specify a string match. Wildcards are not to be used with IP addresses; however, they may accidentally work if reverse DNS lookups fail. When specifying wildcards in fully qualified domain names, dots (.) are not included in the wildcard. For example, *.example.com includes one.example.com but does not include one.two.example.com.

- IP 네트워크 — a.b.c.d/z를 사용합니다. 여기서 a.b.c.d는 네트워크이고 z는 넷마스크의 비트 수를 나타냅니다 (예, 192.168.0.0/24). a.b.c.d/netmask 형식도 사용 가능합니다. 여기서 a.b.c.d는 네트워크이고 netmask는 넷마스크를 의미합니다. (예,192.168.100.8/255.255.255.0)

- 넷그룹 — @group-name 형식을 사용합니다. 여기서 group-name 부분은 NIS 그룹명입니다.

20.7. The /etc/exports Configuration File

The /etc/exports file controls which file systems are exported to remote hosts and specifies options. Blank lines are ignored, comments can be made by starting a line with the hash mark (#), and long lines can be wrapped with a backslash (\). Each exported file system should be on its own individual line, and any lists of authorized hosts placed after an exported file system must be separated by space characters. Options for each of the hosts must be placed in parentheses directly after the host identifier, without any spaces separating the host and the first parenthesis. Valid host types are gss/krb5, gss/krb5i, and gss/krb5p.

A line for an exported file system has the following structure:

```
<export> <host1>(<options>) <hostN>(<options>)...
```

In this structure, replace <export> with the directory being exported, replace <host1> with the host or network to which the export is being shared, and replace <options> with the options for that host or network. Additional hosts can be specified in a space separated list.

The following methods can be used to specify host names:

- single host — Where one particular host is specified with a fully qualified domain name, hostname, or IP address.

- wildcards — Where a * or ? character is used to take into account a grouping of fully qualified domain names that match a particular string of letters. Wildcards should not be used with IP addresses; however, it is possible for them to work accidentally if reverse DNS lookups fail.

  Be careful when using wildcards with fully qualified domain names, as they tend to be more exact than expected. For example, the use of *.example.com as a wildcard allows sales.example.com to access an exported file system, but not bob.sales.example.com. To match both possibilities both *.example.com and *.*.example.com must be specified.

- IP networks — Allows the matching of hosts based on their IP addresses within a larger network. For example, 192.168.0.0/28 allows the first 16 IP addresses, from 192.168.0.0 to 192.168.0.15, to access the exported file system, but not 192.168.0.16 and higher.

- netgroups — Permits an NIS netgroup name, written as @<group-name>, to be used. This effectively puts the NIS server in charge of access control for this exported file system, where users can be added and removed from an NIS group without affecting /etc/exports.

In its simplest form, the /etc/exports file only specifies the exported directory and the hosts permitted to access it, as in the following example:

```
/exported/directory bob.example.com
```

In the example, bob.example.com can mount /exported/directory/. Because no options are specified in this example, the following default NFS options take effect:

- ro — Mounts of the exported file system are read-only. Remote hosts are not able to make changes to the data shared on the file system. To allow hosts to make changes to the file system, the read/write (rw) option must be specified.

- wdelay — Causes the NFS server to delay writing to the disk if it suspects another write request is imminent. This can improve performance by reducing the number of times the disk must be accessed by separate write commands, reducing write overhead. The no_wdelay option turns off this feature, but is only available when using the sync option.

- root_squash — Prevents root users connected remotely from having root privileges and assigns them the user ID for the user nfsnobody. This effectively "squashes" the power of the remote root user to the lowest local user, preventing unauthorized alteration of files on the remote server. Alternatively, the no_root_squash option turns off root squashing. To squash every remote user, including root, use the all_squash option. To specify the user and group IDs to use with remote users from a particular host, use the anonuid and anongid options, respectively. In this case, a special user account can be created for remote NFS users to share and specify (anonuid=<uid-value>,anongid=<gid-value>), where <uid-value> is the user ID number and <gid-value> is the group ID number.

중요

By default, access control lists (ACLs) are supported by NFS under Red Hat Enterprise Linux. To disable this feature, specify the no_acl option when exporting the file system.

Each default for every exported file system must be explicitly overridden. For example, if the rw option is not specified, then the exported file system is shared as read-only. The following is a sample line from /etc/exports which overrides two default options:

```
/another/exported/directory 192.168.0.3(rw,sync)
```

In this example 192.168.0.3 can mount /another/exported/directory/ read/write and all transfers to disk are committed to the disk before the write request by the client is completed.

Additionally, other options are available where no default value is specified. These include the ability to disable sub-tree checking, allow access from insecure ports, and allow insecure file locks (necessary for certain early NFS client implementations). Refer to the exports man page for details on these lesser used options.

주의

The format of the /etc/exports file is very precise, particularly in regards to use of the space character. Remember to always separate exported file systems from hosts and hosts from one another with a space character. However, there should be no other space characters in the file except on comment lines.

For example, the following two lines do not mean the same thing:

```
/home bob.example.com(rw)
/home bob.example.com (rw)
```

The first line allows only users from bob.example.com read/write access to the /home directory. The second line allows users from bob.example.com to mount the directory as read-only (the default), while the rest of the world can mount it read/write.

20.7.1. The exportfs Command

Every file system being exported to remote users via NFS, as well as the access level for those file systems, are listed in the /etc/exports file. When the nfs service starts, the /usr/sbin/exportfs command launches and reads this file, passes control to rpc.mountd (if NFSv2 or NFSv3) for the actual mounting process, then to rpc.nfsd where the file systems are then available to remote users.

When issued manually, the /usr/sbin/exportfs command allows the root user to selectively export or unexport directories without restarting the NFS service. When given the proper options, the /usr/sbin/exportfs command writes the exported file systems to /var/lib/nfs/xtab. Since rpc.mountd refers to the xtab file when deciding access privileges to a file system, changes to the list of exported file systems take effect immediately.

The following is a list of commonly used options available for /usr/sbin/exportfs:

- -r — Causes all directories listed in /etc/exports to be exported by constructing a new export list in /etc/lib/nfs/xtab. This option effectively refreshes the export list with any changes that have been made to /etc/exports.
- -a — Causes all directories to be exported or unexported, depending on what other options are passed to /usr/sbin/exportfs. If no other options are specified, /usr/sbin/exportfs exports all file systems specified in /etc/exports.
- -o file-systems — Specifies directories to be exported that are not listed in /etc/exports. Replace file-systems with additional file systems to be exported. These file systems must be formatted in the same way they are specified in /etc/exports. Refer to 20.7절. "The /etc/exports Configuration File" for more information on /etc/exports syntax. This option is often used to test an exported file system before adding it permanently to the list of file systems to be exported.
- -i — Ignores /etc/exports; only options given from the command line are used to define exported file systems.
- -u — Unexports all shared directories. The command /usr/sbin/exportfs -ua suspends NFS file sharing while keeping all NFS daemons up. To re-enable NFS sharing, type exportfs -r.
- -v — Verbose operation, where the file systems being exported or unexported are displayed in greater detail when the exportfs command is executed.

If no options are passed to the /usr/sbin/exportfs command, it displays a list of currently exported file systems.

For more information about the /usr/sbin/exportfs command, refer to the exportfs man page.

20.7.1.1. Using exportfs with NFSv4

The exportfs command is used in maintaining the NFS table of exported file systems. When typed in a terminal with no arguments, the exportfs command shows all the exported directories.

Since NFSv4 no longer utilizes the MOUNT protocol, which was used with the NFSv2 and NFSv3 protocols, the mounting of file systems has changed.

An NFSv4 client now has the ability to see all of the exports served by the NFSv4 server as a single file system, called the NFSv4 pseudo-file system. On Red Hat Enterprise Linux, the pseudo-file system is identified as a single, real file system, identified at export with the fsid=0 option.

For example, the following commands could be executed on an NFSv4 server:

```
mkdir /exports
mkdir /exports/opt
mkdir /exports/etc
mount --bind /usr/local/opt /exports/opt
mount --bind /usr/local/etc /exports/etc
exportfs -o fsid=0,insecure,no_subtree_check gss/krb5p:/exports
exportfs -o rw,nohide,insecure,no_subtree_check gss/krb5p:/exports/opt
exportfs -o rw,nohide,insecure,no_subtree_check gss/krb5p:/exports/etc
```

In this example, clients are provided with multiple file systems to mount, by using the --bind option which creates unbreakable links.

Because of the pseudo-file systems feature, NFS version 2, 3 and 4 export configurations are not always compatible. For example, given the following directory tree:

```
/home
/home/sam
/home/john
/home/joe
```

and the export:

```
/home *(rw,fsid=0,sync)
```

Using NFS version 2,3 and 4 the following would work:

```
mount server:/home /mnt/home
ls /mnt/home/joe
```

Using v4 the following would work:

```
mount -t nfs4 server:/ /mnt/home
ls /mnt/home/joe
```

The difference being "server:/home" and "server:/". To make the exports configurations compatible for all version, one needs to export (read only) the root filesystem with an fsid=0. The fsid=0 signals the NFS server that this export is the root.

```
/ *(ro,fsid=0)
/home *(rw,sync,nohide)
```

Now with these exports, both "mount server:/home /mnt/home" and "mount -t nfs server:/home /mnt/home" will work as expected.

20.8. Securing NFS

NFS is well suited for sharing entire file systems with a large number of known hosts in a transparent manner. However, with ease of use comes a variety of potential security problems.

The following points should be considered when exporting NFS file systems on a server or mounting them on a client. Doing so minimizes NFS security risks and better protects data on the server.

20.8.1. Host Access

Depending on which version of NFS you plan to implement, depends on your existing network environment, and your security concerns. The following sections explain the differences between implementing security measures with NFSv2, NFSv3, and NFSv4. If at all possible, use of NFSv4 is recommended over other versions of NFS.

20.8.1.1. Using NFSv2 or NFSv3

NFS controls who can mount an exported file system based on the host making the mount request, not the user that actually uses the file system. Hosts must be given explicit rights to mount the exported file system. Access control is not possible for users, other than through file and directory permissions. In other words, once a file system is exported via NFS, any user on any remote host connected to the NFS server can access the shared data. To limit the potential risks, administrators often allow read-only access or squash user permissions to a common user and group ID. Unfortunately, these solutions prevent the NFS share from being used in the way it was originally intended.

Additionally, if an attacker gains control of the DNS server used by the system exporting the NFS file system, the system associated with a particular hostname or fully qualified domain name can be pointed to an unauthorized machine. At this point, the unauthorized machine is the system permitted to mount the NFS share, since no username or password information is exchanged to provide additional security for the NFS mount.

Wildcards should be used sparingly when exporting directories via NFS as it is possible for the scope of the wildcard to encompass more systems than intended.

It is also possible to restrict access to the portmap service via TCP wrappers. Access to ports used by portmap, rpc.mountd, and rpc.nfsd can also be limited by creating firewall rules with iptables.

For more information on securing NFS and portmap, refer to 46.9절. "IPTables" .

20.8.1.2. Using NFSv4

The release of NFSv4 brought a revolution to authentication and security to NFS exports. NFSv4 mandates the implementation of the RPCSEC_GSS kernel module, the Kerberos version 5 GSS-API mechanism, SPKM-3, and LIPKEY. With NFSv4, the mandatory security mechanisms are oriented towards authenticating individual users, and not client machines as used in NFSv2 and NFSv3.

참고

It is assumed that a Kerberos ticket-granting server (KDC) is installed and configured correctly, prior to configuring an NFSv4 server. Kerberos is a network authentication system which allows clients and servers to authenticate to each other through use of symmetric encryption and a trusted third party, the KDC.

NFSv4 includes ACL support based on the Microsoft Windows NT model, not the POSIX model, because of its features and because it is widely deployed. NFSv2 and NFSv3 do not have support for native ACL attributes.

Another important security feature of NFSv4 is the removal of the use of the MOUNT protocol for mounting file systems. This protocol presented possible security holes because of the way that it handled file handles.

For more information on the RPCSEC_GSS framework, including how rpc.svcgssd and rpc.gssd inter operate, refer to http://www.citi.umich.edu/projects/nfsv4/gssd/.

20.8.2. File Permissions

Once the NFS file system is mounted read/write by a remote host, the only protection each shared file has is its permissions. If two users that share the same user ID value mount the same NFS file system, they can modify each others files. Additionally, anyone logged in as root on the client system can use the su - command to become a user who could access particular files via the NFS share.

By default, access control lists (ACLs) are supported by NFS under Red Hat Enterprise Linux. It is not recommended that this feature be disabled.

The default behavior when exporting a file system via NFS is to use root squashing. This sets the user ID of anyone accessing the NFS share as the root user on their local machine to a value of the server's nfsnobody account. Never turn off root squashing.

If exporting an NFS share as read-only, consider using the all_squash option, which makes every user accessing the exported file system take the user ID of the nfsnobody user.

20.9. NFS and portmap

> 참고
>
> The following section only applies to NFSv2 or NFSv3 implementations that require the portmap service for backward compatibility.

The portmapper maps RPC services to the ports they are listening on. RPC processes notify portmap when they start, registering the ports they are listening on and the RPC program numbers they expect to serve. The client system then contacts portmap on the server with a particular RPC program number. The portmap service redirects the client to the proper port number so it can communicate with the requested service.

Because RPC-based services rely on portmap to make all connections with incoming client requests, portmap must be available before any of these services start.

The portmap service uses TCP wrappers for access control, and access control rules for portmap affect all RPC-based services. Alternatively, it is possible to specify access control rules for each of the NFS RPC daemons. The man pages for rpc.mountd and rpc.statd contain information regarding the precise syntax for these rules.

20.9.1. Troubleshooting NFS and portmap

Because portmap provides coordination between RPC services and the port numbers used to communicate with them, it is useful to view the status of current RPC services using portmap when troubleshooting. The rpcinfo command shows each RPC-based service with port numbers, an RPC program number, a version number, and an IP protocol type (TCP or UDP).

To make sure the proper NFS RPC-based services are enabled for portmap, issue the following command as root:

```
rpcinfo -p
```

The following is sample output from this command:

```
program vers proto   port
100000    2   tcp    111  portmapper
100000    2   udp    111  portmapper
100021    1   udp  32774  nlockmgr
100021    3   udp  32774  nlockmgr
100021    4   udp  32774  nlockmgr
100021    1   tcp  34437  nlockmgr
100021    3   tcp  34437  nlockmgr
100021    4   tcp  34437  nlockmgr
100011    1   udp    819  rquotad
100011    2   udp    819  rquotad
100011    1   tcp    822  rquotad
100011    2   tcp    822  rquotad
100003    2   udp   2049  nfs
100003    3   udp   2049  nfs
100003    2   tcp   2049  nfs
100003    3   tcp   2049  nfs
100005    1   udp    836  mountd
100005    1   tcp    839  mountd
100005    2   udp    836  mountd
100005    2   tcp    839  mountd
100005    3   udp    836  mountd
100005    3   tcp    839  mountd
```

If one of the NFS services does not start up correctly, portmap is unable to map RPC requests from clients for that service to the correct port. In many cases, if NFS is not present in rpcinfo output, restarting NFS causes the service to correctly register with portmap and begin working. For instructions on starting NFS, refer to 20.5절. "Starting and Stopping NFS" .

Other useful options are available for the rpcinfo command. Refer to the rpcinfo man page for more information.

20.10. Using NFS over TCP

The default transport protocol for NFSv4 is TCP; however, the Red Hat Enterprise Linux 5 kernel includes support for NFS over UDP. To use NFS over UDP, include the -o udp option to mount when mounting the NFS-exported file system on the client system.

There are three ways to configure an NFS file system export. On demand via the command line (client side), automatically via the /etc/fstab file (client side), and automatically via autofs configuration files, such as /etc/auto.master and /etc/auto.misc (server side with NIS).

For example, on demand via the command line (client side):

```
mount -o udp shadowman.example.com:/misc/export /misc/local
```

When the NFS mount is specified in /etc/fstab (client side):

```
server:/usr/local/pub    /pub   nfs    rsize=8192,wsize=8192,timeo=14,intr,udp
```

When the NFS mount is specified in an autofs configuration file for a NIS server, available for NIS enabled workstations:

```
myproject   -rw,soft,intr,rsize=8192,wsize=8192,udp  penguin.example.net:/proj52
```

Since the default is TCP, if the -o udp option is not specified, the NFS-exported file system is accessed via TCP.

TCP를 사용하는 장점은 다음과 같습니다:

- Improved connection durability, thus less NFS stale file handles messages.

- UDP는 완료된 패킷만 확인하는 반면, TCP는 모든 패킷을 확인하기 때문에 작업 부하가 많은 네트워크 상에서 성능이 향상됩니다.

- TCP has better congestion control than UDP. On a very congested network, UDP packets are the first packets that are dropped. This means that if NFS is writing data (in 8K chunks) all of that 8K must be retransmitted over UDP. Because of TCP's reliability, only parts of that 8K data are transmitted at a time.

- Error detection. When a TCP connection breaks (due to the server being unavailable) the client stops sending data and restarts the connection process once the server becomes available. With UDP, since it's connection-less, the client continues to pound the network with data until the server reestablishes a connection.

주된 장점은 TCP 프로토콜의 오버헤드로 인한 성능 저하가 매우 작다는 것입니다.

20.11. 추가 자료

Administering an NFS server can be a challenge. Many options, including quite a few not mentioned in this chapter, are available for exporting or mounting NFS shares. Consult the following sources for more information.

20.11.1. 설치된 문서 자료

- /usr/share/doc/nfs-utils-<version-number>/ — Replace <version-number> with the version number of the NFS package installed. This directory contains a wealth of information about the NFS implementation for Linux, including a look at various NFS configurations and their impact on file transfer performance.

- man mount — Contains a comprehensive look at mount options for both NFS server and client configurations.

- man fstab — Gives details for the format of the /etc/fstab file used to mount file systems at boot-time.

- man nfs — Provides details on NFS-specific file system export and mount options.

- man exports — Shows common options used in the /etc/exports file when exporting NFS file systems.

20.11.2. 유용한 웹사이트

- http://nfs.sourceforge.net/ — The home of the Linux NFS project and a great place for project status updates.
- http://www.citi.umich.edu/projects/nfsv4/linux/ — An NFSv4 for Linux 2.6 kernel resource.
- http://www.nfsv4.org[1] — The home of NFS version 4 and all related standards.
- http://www.vanemery.com/Linux/NFSv4/NFSv4-no-rpcsec.html — Describes the details of NFSv4 with Fedora Core 2, which includes the 2.6 kernel.
- http://www.sane.nl/events/sane2000/papers/pawlowski.pdf — An excellent whitepaper on the features and enhancements of the NFS Version 4 protocol.
- http://wiki.autofs.net — The Autofs wiki, discussions, documentation and enhancements.

20.11.3. 관련 서적

- Managing NFS and NIS by Hal Stern, Mike Eisler, and Ricardo Labiaga; O'Reilly & Associates — Makes an excellent reference guide for the many different NFS export and mount options available.
- NFS Illustrated by Brent Callaghan; Addison-Wesley Publishing Company — Provides comparisons of NFS to other network file systems and shows, in detail, how NFS communication occurs.

[1] http://www.nfsv4.org/

Samba

Samba is an open source implementation of the Server Message Block (SMB) protocol. It allows the networking of Microsoft Windows®, Linux, UNIX, and other operating systems together, enabling access to Windows-based file and printer shares. Samba's use of SMB allows it to appear as a Windows server to Windows clients.

21.1. Introduction to Samba

The third major release of Samba, version 3.0.0, introduced numerous improvements from prior versions, including:

- The ability to join an Active Directory domain by means of LDAP and Kerberos
- Built in Unicode support for internationalization
- Support for Microsoft Windows XP Professional client connections to Samba servers without needing local registry hacking
- Two new documents developed by the Samba.org team, which include a 400+ page reference manual, and a 300+ page implementation and integration manual. For more information about these published titles, refer to 21.12.2절. "Related Books" .

21.1.1. Samba Features

Samba is a powerful and versatile server application. Even seasoned system administrators must know its abilities and limitations before attempting installation and configuration.

What Samba can do:

- Serve directory trees and printers to Linux, UNIX, and Windows clients
- Assist in network browsing (with or without NetBIOS)
- Authenticate Windows domain logins
- Provide Windows Internet Name Service (WINS) name server resolution
- Act as a Windows NT®-style Primary Domain Controller (PDC)
- Act as a Backup Domain Controller (BDC) for a Samba-based PDC
- Act as an Active Directory domain member server
- Join a Windows NT/2000/2003 PDC

What Samba cannot do:

- Act as a BDC for a Windows PDC (and vice versa)
- Act as an Active Directory domain controller

21.2. Samba Daemons and Related Services

The following is a brief introduction to the individual Samba daemons and services.

21.2.1. Samba Daemons

Samba is comprised of three daemons (smbd, nmbd, and winbindd). Two services (smb and windbind) control how the daemons are started, stopped, and other service-related features. Each daemon is listed in detail, as well as which specific service has control over it.

smbd

The smbd server daemon provides file sharing and printing services to Windows clients. In addition, it is responsible for user authentication, resource locking, and data sharing through the SMB protocol. The default ports on which the server listens for SMB traffic are TCP ports 139 and 445.

The smbd daemon is controlled by the smb service.

nmbd

The nmbd server daemon understands and replies to NetBIOS name service requests such as those produced by SMB/CIFS in Windows-based systems. These systems include Windows 95/98/ME, Windows NT, Windows 2000, Windows XP, and LanManager clients. It also participates in the browsing protocols that make up the Windows Network Neighborhood view. The default port that the server listens to for NMB traffic is UDP port 137.

The nmbd daemon is controlled by the smb service.

winbindd

The winbind service resolves user and group information on a server running Windows NT 2000 or Windows Server 2003. This makes Windows user / group information understandable by UNIX platforms. This is achieved by using Microsoft RPC calls, Pluggable Authentication Modules (PAM), and the Name Service Switch (NSS). This allows Windows NT domain users to appear and operate as UNIX users on a UNIX machine. Though bundled with the Samba distribution, the winbind service is controlled separately from the smb service.

The winbindd daemon is controlled by the winbind service and does not require the smb service to be started in order to operate. Winbindd is also used when Samba is an Active Directory member, and may also be used on a Samba domain controller (to implement nested groups and/or interdomain trust). Because winbind is a client-side service used to connect to Windows NT-based servers, further discussion of winbind is beyond the scope of this manual.

알림

You may refer to 21.11절. "Samba Distribution Programs" for a list of utilities included in the Samba distribution.

21.3. Connecting to a Samba Share

You can use Nautilus to view available Samba shares on your network. Select Places (on the Panel) > Network Servers to view a list of Samba workgroups on your network. You can also type smb: in the File > Open Location bar of Nautilus to view the workgroups.

As shown in 그림 21.1. "SMB Workgroups in Nautilus" , an icon appears for each available SMB workgroup on the network.

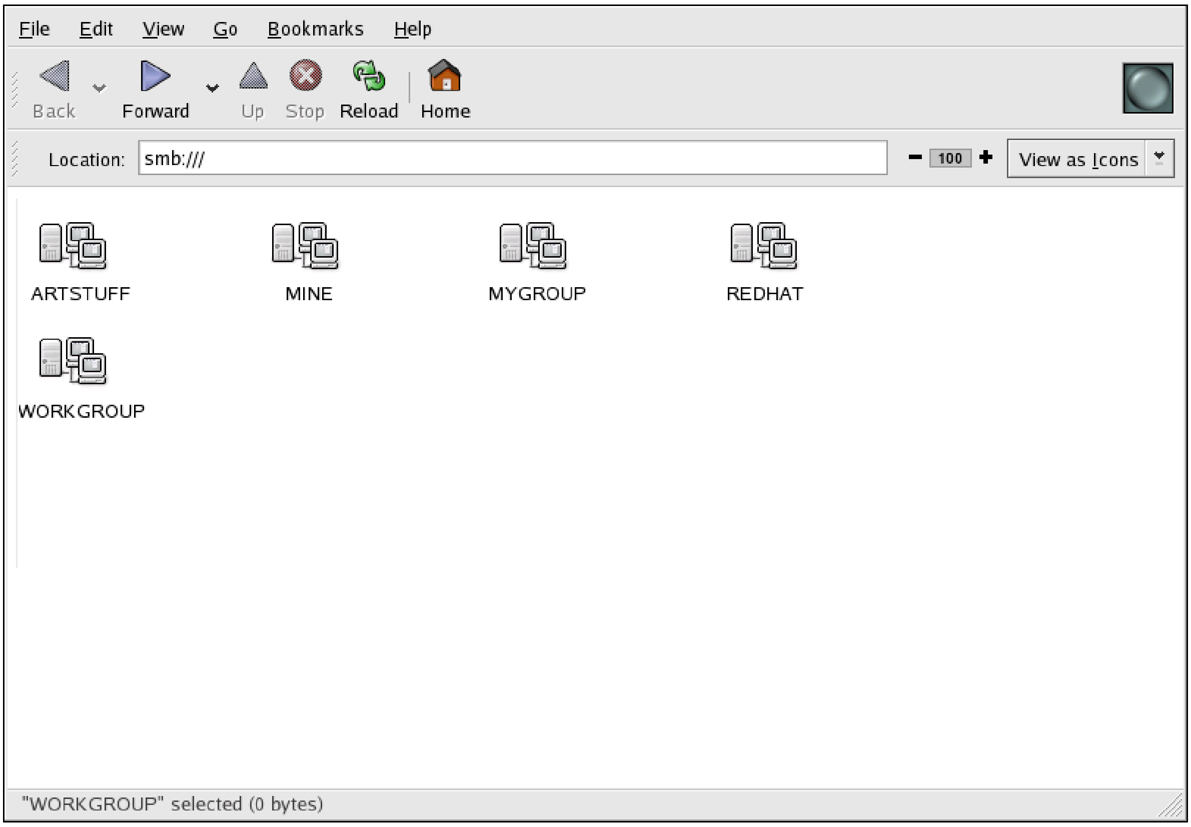

smb:///

"WORKGROUP" selected (0 bytes)

그림 21.1. SMB Workgroups in Nautilus

Double-click one of the workgroup icons to view a list of computers within the workgroup.

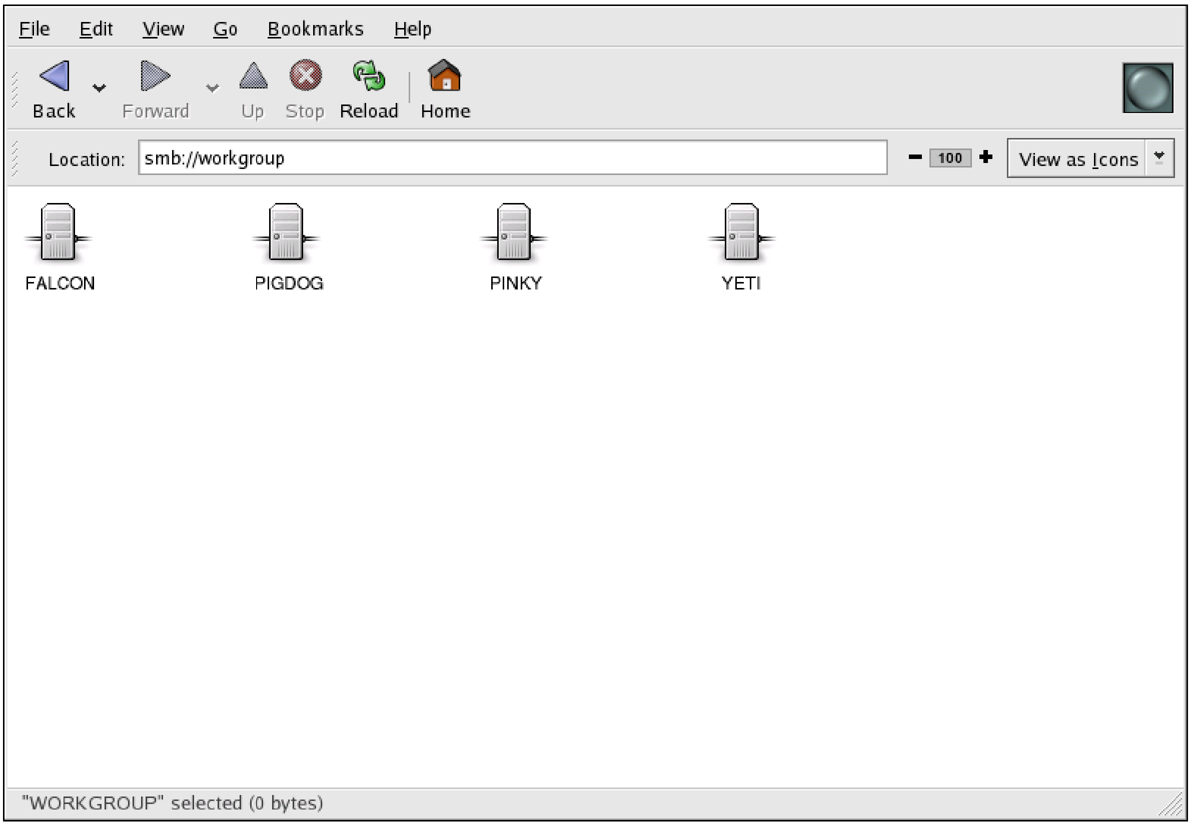

그림 21.2. SMB Machines in Nautilus

As you can see from 그림 21.2. "SMB Machines in Nautilus" , there is an icon for each machine within the workgroup. Double-click on an icon to view the Samba shares on the machine. If a username and password combination is required, you are prompted for them.

Alternately, you can also specify the Samba server and sharename in the Location: bar for Nautilus using the following syntax (replace <servername> and <sharename> with the appropriate values):

```
smb://<servername>/<sharename>
```

21.3.1. Command Line

To query the network for Samba servers, use the findsmb command. For each server found, it displays its IP address, NetBIOS name, workgroup name, operating system, and SMB server version.

To connect to a Samba share from a shell prompt, type the following command:

```
smbclient //<hostname>/<sharename> -U <username>
```

Replace <hostname> with the hostname or IP address of the Samba server you want to connect to, <sharename> with the name of the shared directory you want to browse, and <username> with the Samba username for the system. Enter the correct password or press Enter if no password is required for the user.

If you see the smb:\> prompt, you have successfully logged in. Once you are logged in, type help for a list of commands. If you wish to browse the contents of your home directory, replace sharename with your username. If the -U switch is not used, the username of the current user is passed to the Samba server.

To exit smbclient, type exit at the smb:\> prompt.

21.3.2. Mounting the Share

Sometimes it is useful to mount a Samba share to a directory so that the files in the directory can be treated as if they are part of the local file system.

To mount a Samba share to a directory, create create a directory to mount it to (if it does not already exist), and execute the following command as root:

```
mount -t cifs -o <username>,<password> //<servername>/<sharename> /mnt/point/
```

This command mounts <sharename> from <servername> in the local directory /mnt/point/. For more information about mounting a samba share, refer to man mount.cifs.

21.4. Configuring a Samba Server

The default configuration file (/etc/samba/smb.conf) allows users to view their home directories as a Samba share. It also shares all printers configured for the system as Samba shared printers. In other words, you can attach a printer to the system and print to it from the Windows machines on your network.

21.4.1. Graphical Configuration

To configure Samba using a graphical interface, use the Samba Server Configuration Tool. For command line configuration, skip to 21.4.2절. “Command Line Configuration” .

The Samba Server Configuration Tool is a graphical interface for managing Samba shares, users, and basic server settings. It modifies the configuration files in the /etc/samba/ directory. Any changes to these files not made using the application are preserved.

To use this application, you must be running the X Window System, have root privileges, and have the system-config-samba RPM package installed. To start the Samba Server Configuration Tool from the desktop, go to the System (on the Panel) > Administration > Server Settings > Samba or type the command system-config-samba at a shell prompt (for example, in an XTerm or a GNOME terminal).

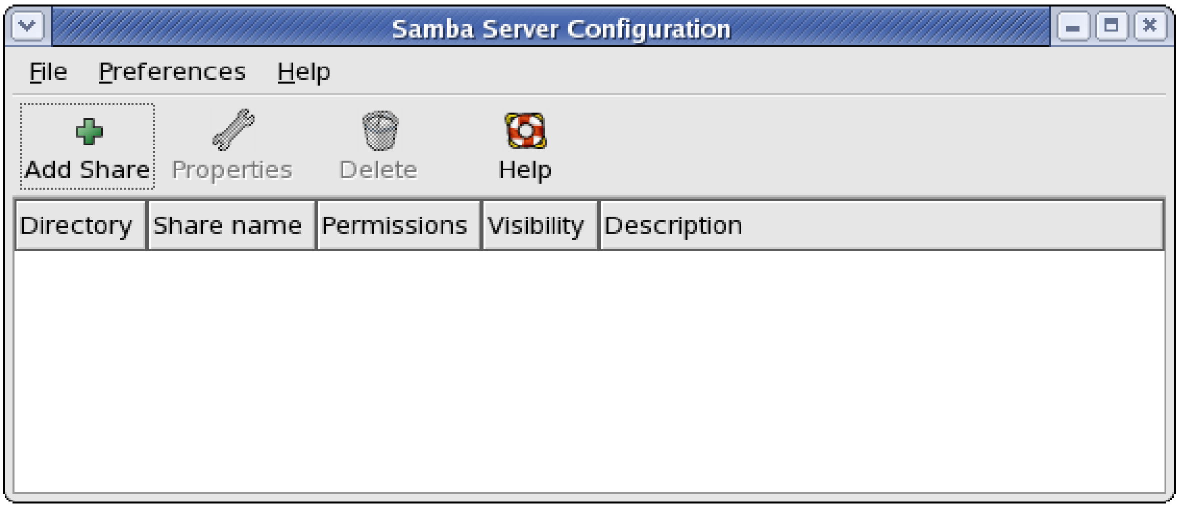

그림 21.3. Samba Server Configuration Tool

알림

The Samba Server Configuration Tool does not display shared printers or the default stanza that allows users to view their own home directories on the Samba server.

21.4.1.1. Configuring Server Settings

The first step in configuring a Samba server is to configure the basic settings for the server and a few security options. After starting the application, select Preferences > Server Settings from the pulldown menu. The Basic tab is displayed as shown in 그림 21.4. "Configuring Basic Server Settings" .

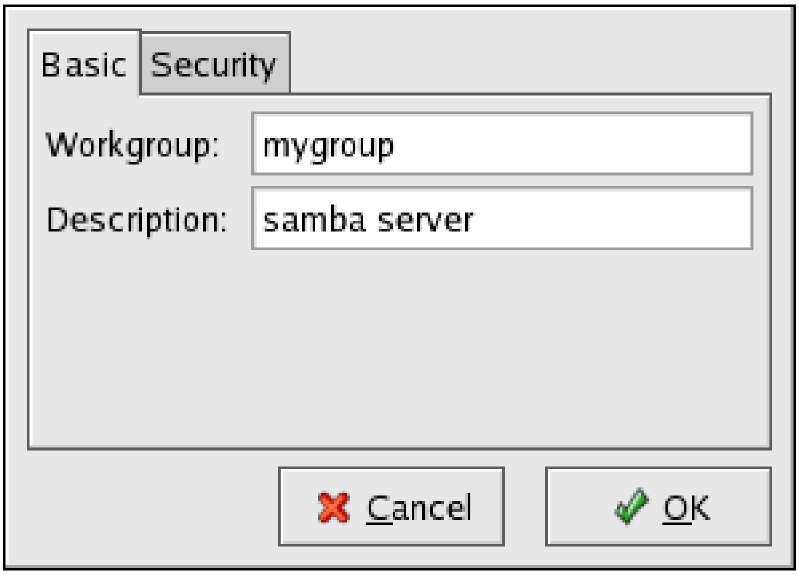

Workgroup: mygroup

Description: samba server

그림 21.4. Configuring Basic Server Settings

On the Basic tab, specify which workgroup the computer should be in as well as a brief description of the computer. They correspond to the workgroup and server string options in smb.conf.

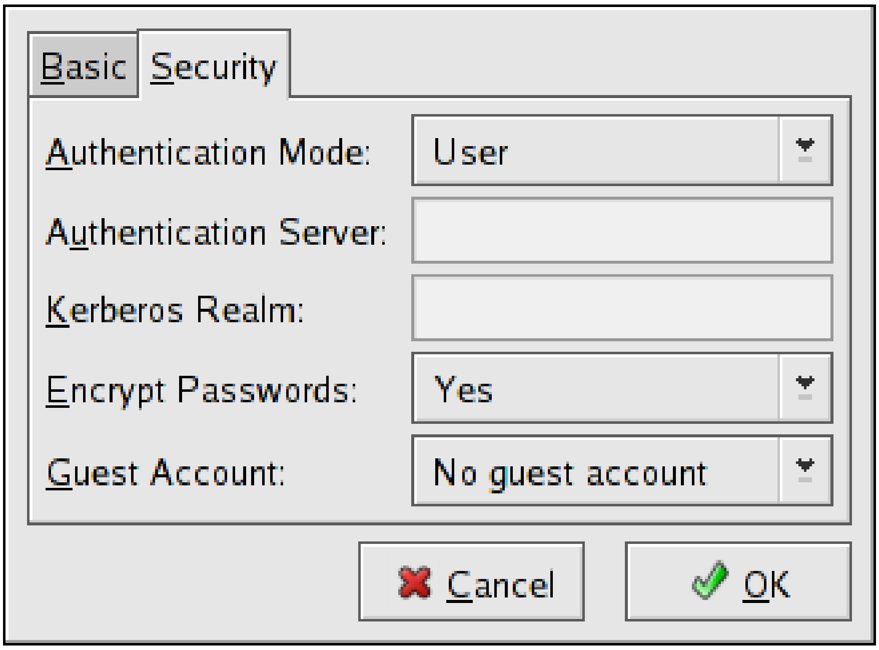

Authentication Mode: User

Authentication Server:

Kerberos Realm:

Encrypt Passwords: Yes

Guest Account: No guest account

그림 21.5. Configuring Security Server Settings

The Security tab contains the following options:

- Authentication Mode — This corresponds to the security option. Select one of the following types of authentication.

  - ADS — The Samba server acts as a domain member in an Active Directory Domain (ADS) realm. For this option, Kerberos must be installed and configured on the server, and Samba must become a member of the ADS realm using the net utility, which is part of the samba-client package. Refer to the net man page for details. This option does not configure Samba to be an ADS Controller. Specify the realm of the Kerberos server in the Kerberos Realm field.

    알림

    The Kerberos Realm field must be supplied in all uppercase letters, such as EXAMPLE.COM.

    Using a Samba server as a domain member in an ADS realm assumes proper configuration of Kerberos, including the /etc/krb5.conf file.

  - Domain — The Samba server relies on a Windows NT Primary or Backup Domain Controller to verify the user. The server passes the username and password to the Controller and waits for it to return. Specify the NetBIOS name of the Primary or Backup Domain Controller in the Authentication Server field.

    The Encrypted Passwords option must be set to Yes if this is selected.

  - Server — The Samba server tries to verify the username and password combination by passing them to another Samba server. If it can not, the server tries to verify using the user authentication mode. Specify the NetBIOS name of the other Samba server in the Authentication Server field.

  - Share — Samba users do not have to enter a username and password combination on a per Samba server basis. They are not prompted for a username and password until they try to connect to a specific shared directory from a Samba server.
  - User — (Default) Samba users must provide a valid username and password on a per Samba server basis. Select this option if you want the Windows Username option to work. Refer to 21.4.1.2절. "Managing Samba Users" for details.

- Encrypt Passwords — This option must be enabled if the clients are connecting from a system with Windows 98, Windows NT 4.0 with Service Pack 3, or other more recent versions of Microsoft Windows. The passwords are transferred between the server and the client in an encrypted format instead of as a plain-text word that can be intercepted. This corresponds to the encrypted passwords option. Refer to 21.4.3절. "Encrypted Passwords" for more information about encrypted Samba passwords.
- Guest Account — When users or guest users log into a Samba server, they must be mapped to a valid user on the server. Select one of the existing usernames on the system to be the guest Samba account. When guests log in to the Samba server, they have the same privileges as this user. This corresponds to the guest account option.

After clicking OK, the changes are written to the configuration file and the daemon is restarted; thus, the changes take effect immediately.

21.4.1.2. Managing Samba Users

The Samba Server Configuration Tool requires that an existing user account be active on the system acting as the Samba server before a Samba user can be added. The Samba user is associated with the existing user account.

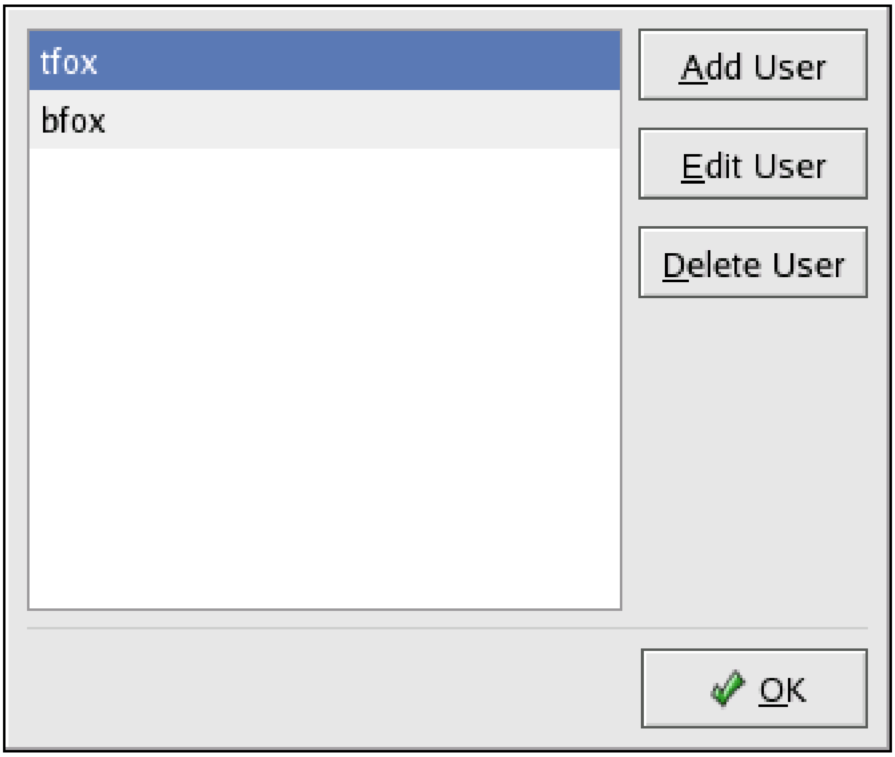

그림 21.6. Managing Samba Users

To add a Samba user, select Preferences > Samba Users from the pulldown menu, and click the Add User button. In the Create New Samba User window select a Unix Username from the list of existing users on the local system.

If the user has a different username on a Windows machine and needs to log into the Samba server from the Windows machine, specify that Windows username in the Windows Username field. The Authentication Mode on the Security tab of the Server Settings preferences must be set to User for this option to work.

Also, configure a Samba Password for the Samba User and confirm it by typing it again. Even if you opt to use encrypted passwords for Samba, it is recommended that the Samba passwords for all users are different from their system passwords.

To edit an existing user, select the user from the list, and click Edit User. To delete an existing Samba user, select the user, and click the Delete User button. Deleting a Samba user does not delete the associated system user account.

The users are modified immediately after clicking the OK button.

21.4.1.3. Adding a Share

To create a Samba share, click the Add button from the main Samba configuration window.

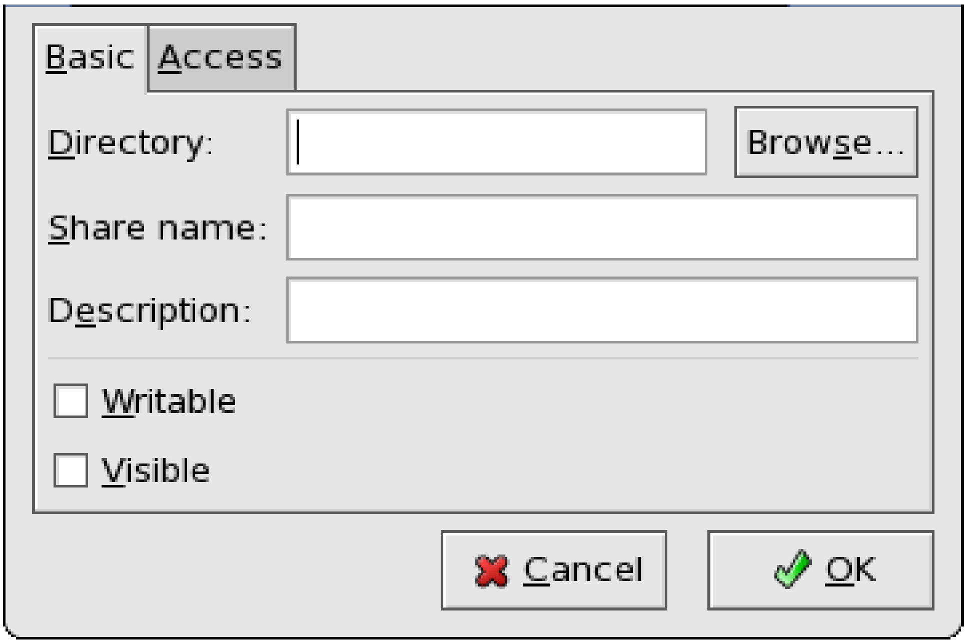

Directory:

Share name:

Description:

□ Writable

□ Visible

그림 21.7. Adding a Share

The Basic tab configures the following options:

- Directory — The directory to share via Samba. The directory must exist before it can be entered here.
- Share name — The actual name of the share that is seen from remote machines. By default, it is the same value as Directory, but can be configured.
- Descriptions — A brief description of the share.
- Writable — Enables users to read and write to the shared directory

- Visible — Grants read-only rights to users for the shared directory.

On the Access tab, select whether to allow only specified users to access the share or whether to allow all Samba users to access the share. If you select to allow access to specific users, select the users from the list of available Samba users.

The share is added immediately after clicking OK.

21.4.2. Command Line Configuration

Samba uses /etc/samba/smb.conf as its configuration file. If you change this configuration file, the changes do not take effect until you restart the Samba daemon with the command service smb restart.

To specify the Windows workgroup and a brief description of the Samba server, edit the following lines in your smb.conf file:

```
workgroup = WORKGROUPNAME
server string = BRIEF COMMENT ABOUT SERVER
```

Replace WORKGROUPNAME with the name of the Windows workgroup to which this machine should belong. The BRIEF COMMENT ABOUT SERVER is optional and is used as the Windows comment about the Samba system.

To create a Samba share directory on your Linux system, add the following section to your smb.conf file (after modifying it to reflect your needs and your system):

```
[sharename]
comment = Insert a comment here
path = /home/share/
valid users = tfox carole
public = no
writable = yes
printable = no
create mask = 0765
```

The above example allows the users tfox and carole to read and write to the directory /home/share, on the Samba server, from a Samba client.

21.4.3. Encrypted Passwords

Encrypted passwords are enabled by default because it is more secure to do so. To create a user with an encrypted password, use the command smbpasswd -a <username> .

21.5. Starting and Stopping Samba

To start a Samba server, type the following command in a shell prompt while logged in as root:

```
service smb start
```

Important

To set up a domain member server, you must first join the domain or Active Directory using the net join command before starting the smb service.

To stop the server, type the following command in a shell prompt while logged in as root:

```
service smb stop
```

The restart option is a quick way of stopping and then starting Samba. This is the most reliable way to make configuration changes take effect after editing the configuration file for Samba. Note that the restart option starts the daemon even if it was not running originally.

To restart the server, type the following command in a shell prompt while logged in as root:

```
service smb restart
```

The condrestart (conditional restart) option only starts smb on the condition that it is currently running. This option is useful for scripts, because it does not start the daemon if it is not running.

알림

When the smb.conf file is changed, Samba automatically reloads it after a few minutes. Issuing a manual restart or reload is just as effective.

To conditionally restart the server, type the following command as root:

```
service smb condrestart
```

A manual reload of the smb.conf file can be useful in case of a failed automatic reload by the smb service. To ensure that the Samba server configuration file is reloaded without restarting the service, type the following command as root:

```
service smb reload
```

By default, the smb service does not start automatically at boot time. To configure Samba to start at boot time, use an initscript utility, such as /sbin/chkconfig, /usr/sbin/ntsysv, or the Services Configuration Tool program. Refer to 17장. 서비스로의 접근 통제 for more information regarding these tools.

21.6. Samba Server Types and the smb.conf File

Samba configuration is straightforward. All modifications to Samba are done in the /etc/samba/smb.conf configuration file. Although the default smb.conf file is well documented, it does not address complex topics such as LDAP, Active Directory, and the numerous domain controller implementations.

The following sections describe the different ways a Samba server can be configured. Keep in mind your needs and the changes required to the smb.conf file for a successful configuration.

21.6.1. Stand-alone Server

A stand-alone server can be a workgroup server or a member of a workgroup environment. A stand-alone server is not a domain controller and does not participate in a domain in any way. The following examples include several anonymous share-level security configurations and one user-level security configuration. For more information on share-level and user-level security modes, refer to 21.7절. "Samba Security Modes" .

21.6.1.1. Anonymous Read-Only

The following smb.conf file shows a sample configuration needed to implement anonymous read-only file sharing. The security = share parameter makes a share anonymous. Note, security levels for a single Samba server cannot be mixed. The security directive is a global Samba parameter located in the [global] configuration section of the smb.conf file.

```
[global]
workgroup = DOCS
netbios name = DOCS_SRV
security = share
[data]
comment = Documentation Samba Server
path = /export
read only = Yes
guest only = Yes
```

21.6.1.2. Anonymous Read/Write

The following smb.conf file shows a sample configuration needed to implement anonymous read/write file sharing. To enable anonymous read/write file sharing, set the read only directive to no. The force user and force group directives are also added to enforce the ownership of any newly placed files specified in the share.

알림

Although having an anonymous read/write server is possible, it is not recommended. Any files placed in the share space, regardless of user, are assigned the user/group combination as specified by a generic user (force user) and group (force group) in the smb.conf file.

```
[global]
workgroup = DOCS
netbios name = DOCS_SRV
security = share
[data]
comment = Data
path = /export
force user = docsbot
force group = users
```

```
read only = No
guest ok = Yes
```

21.6.1.3. Anonymous Print Server

The following smb.conf file shows a sample configuration needed to implement an anonymous print server. Setting browseable to no as shown does not list the printer in Windows Network Neighborhood. Although hidden from browsing, configuring the printer explicitly is possible. By connecting to DOCS_SRV using NetBIOS, the client can have access to the printer if the client is also part of the DOCS workgroup. It is also assumed that the client has the correct local printer driver installed, as the use client driver directive is set to Yes. In this case, the Samba server has no responsibility for sharing printer drivers to the client.

```
[global]
workgroup = DOCS
netbios name = DOCS_SRV
security = share
printcap name = cups
disable spools= Yes
show add printer wizard = No
printing = cups
[printers]
comment = All Printers
path = /var/spool/samba
guest ok = Yes
printable = Yes
use client driver = Yes
browseable = Yes
```

21.6.1.4. Secure Read/Write File and Print Server

The following smb.conf file shows a sample configuration needed to implement a secure read/write print server. Setting the security directive to user forces Samba to authenticate client connections. Notice the [homes] share does not have a force user or force group directive as the [public] share does. The [homes] share uses the authenticated user details for any files created as opposed to the force user and force group in [public].

```
[global]
workgroup = DOCS
netbios name = DOCS_SRV
security = user
printcap name = cups
disable spools = Yes
show add printer wizard = No
printing = cups
[homes]
comment = Home Directories
valid users = %S
read only = No
browseable = No
[public]
comment = Data
path = /export
force user = docsbot
force group = users
guest ok = Yes
[printers]
comment = All Printers
```

```
path = /var/spool/samba
printer admin = john, ed, @admins
create mask = 0600
guest ok = Yes
printable = Yes
use client driver = Yes
browseable = Yes
```

21.6.2. Domain Member Server

A domain member, while similar to a stand-alone server, is logged into a domain controller (either Windows or Samba) and is subject to the domain's security rules. An example of a domain member server would be a departmental server running Samba that has a machine account on the Primary Domain Controller (PDC). All of the department's clients still authenticate with the PDC, and desktop profiles and all network policy files are included. The difference is that the departmental server has the ability to control printer and network shares.

21.6.2.1. Active Directory Domain Member Server

The following smb.conf file shows a sample configuration needed to implement an Active Directory domain member server. In this example, Samba authenticates users for services being run locally but is also a client of the Active Directory. Ensure that your kerberos realm parameter is shown in all caps (for example realm = EXAMPLE.COM). Since Windows 2000/2003 requires Kerberos for Active Directory authentication, the realm directive is required. If Active Directory and Kerberos are running on different servers, the password server directive may be required to help the distinction.

```
[global]
realm = EXAMPLE.COM
security = ADS
encrypt passwords = yes
# Optional. Use only if Samba cannot determine the Kerberos server automatically.
password server = kerberos.example.com
```

In order to join a member server to an Active Directory domain, the following steps must be completed:

- Configuration of the smb.conf file on the member server
- Configuration of Kerberos, including the /etc/krb5.conf file, on the member server
- Creation of the machine account on the Active Directory domain server
- Association of the member server to the Active Directory domain

To create the machine account and join the Windows 2000/2003 Active Directory, Kerberos must first be initialized for the member server wishing to join the Active Directory domain. To create an administrative Kerberos ticket, type the following command as root on the member server:

```
kinit administrator@EXAMPLE.COM
```

The kinit command is a Kerberos initialization script that references the Active Directory administrator account and Kerberos realm. Since Active Directory requires Kerberos tickets, kinit obtains and caches a Kerberos ticket-granting ticket for client/server authentication. For more information on Kerberos, the /etc/krb5.conf file, and the kinit command, refer to 46.6절. “Kerberos” .

To join an Active Directory server (windows1.example.com), type the following command as root on the member server:

```
net ads join -S windows1.example.com -U administrator%password
```

Since the machine windows1 was automatically found in the corresponding Kerberos realm (the kinit command succeeded), the net command connects to the Active Directory server using its required administrator account and password. This creates the appropriate machine account on the Active Directory and grants permissions to the Samba domain member server to join the domain.

알림

Since security = ads and not security = user is used, a local password backend such as smbpasswd is not needed. Older clients that do not support security = ads are authenticated as if security = domain had been set. This change does not affect functionality and allows local users not previously in the domain.

21.6.2.2. Windows NT4-based Domain Member Server

The following smb.conf file shows a sample configuration needed to implement a Windows NT4-based domain member server. Becoming a member server of an NT4-based domain is similar to connecting to an Active Directory. The main difference is NT4-based domains do not use Kerberos in their authentication method, making the smb.conf file simpler. In this instance, the Samba member server functions as a pass through to the NT4-based domain server.

```
[global]
workgroup = DOCS
netbios name = DOCS_SRV
security = domain
[homes]
comment = Home Directories
valid users = %S
read only = No
browseable = No
[public]
comment = Data
path = /export
force user = docsbot
force group = users
guest ok = Yes
```

Having Samba as a domain member server can be useful in many situations. There are times where the Samba server can have other uses besides file and printer sharing. It may be beneficial to make Samba a domain member server in instances where Linux-only applications are required for use in the domain environment. Administrators appreciate keeping track of all machines in the domain, even if not Windows-based. In the event the Windows-based server hardware is deprecated, it is quite easy to modify the smb.conf file to convert the server to a Samba-based PDC. If Windows NT-based servers are upgraded to Windows 2000/2003, the smb.conf file is easily modifiable to incorporate the infrastructure change to Active Directory if needed.

Important

After configuring the smb.conf file, join the domain before starting Samba by typing the following command as root:

```
net rpc join -U administrator%password
```

Note that the -S option, which specifies the domain server hostname, does not need to be stated in the net rpc join command. Samba uses the hostname specified by the workgroup directive in the smb.conf file instead of it being stated explicitly.

21.6.3. Domain Controller

A domain controller in Windows NT is functionally similar to a Network Information Service (NIS) server in a Linux environment. Domain controllers and NIS servers both host user/group information databases as well as related services. Domain controllers are mainly used for security, including the authentication of users accessing domain resources. The service that maintains the user/group database integrity is called the Security Account Manager (SAM). The SAM database is stored differently between Windows and Linux Samba-based systems, therefore SAM replication cannot be achieved and platforms cannot be mixed in a PDC/BDC environment.

In a Samba environment, there can be only one PDC and zero or more BDCs.

Important

Samba cannot exist in a mixed Samba/Windows domain controller environment (Samba cannot be a BDC of a Windows PDC or vice versa). Alternatively, Samba PDCs and BDCs can coexist.

21.6.3.1. Primary Domain Controller (PDC) using tdbsam

The simplest and most common implementation of a Samba PDC uses the tdbsam password database backend. Planned to replace the aging smbpasswd backend, tdbsam has numerous improvements that are explained in more detail in 21.8절. "Samba Account Information Databases" . The passdb backend directive controls which backend is to be used for the PDC.

```
[global]
workgroup = DOCS
netbios name = DOCS_SRV
passdb backend = tdbsam
security = user
add user script = /usr/sbin/useradd -m "%u"
delete user script = /usr/sbin/userdel -r "%u"
add group script = /usr/sbin/groupadd "%g"
delete group script = /usr/sbin/groupdel "%g"
add user to group script = /usr/sbin/usermod -G "%g" "%u"
add machine script = /usr/sbin/useradd -s /bin/false -d /dev/null  -g machines "%u"
# The following specifies the default logon script
# Per user logon scripts can be specified in the user
# account using pdbedit logon script = logon.bat
# This sets the default profile path.
# Set per user paths with pdbedit
logon drive = H:
domain logons = Yes
os level = 35
preferred master = Yes
domain master = Yes
[homes]
 comment = Home Directories
 valid users = %S
 read only = No
[netlogon]
 comment = Network Logon Service
 path = /var/lib/samba/netlogon/scripts
 browseable = No
 read only = No
# For profiles to work, create a user directory under the
# path shown.
mkdir -p /var/lib/samba/profiles/john
[Profiles]
 comment = Roaming Profile Share
 path = /var/lib/samba/profiles
 read only = No
 browseable = No
 guest ok = Yes
 profile acls = Yes
# Other resource shares ... ...
```

To provide a functional PDC system which uses the tdbsam follow these steps:

1. Use a configuration of the smb.conf file as shown in the example above.

2. Add the root user to the Samba password database.

   ```
   smbpasswd -a root
   Provide the password here.
   ```

3. Start the smb service.

4. Make sure all profile, user, and netlogon directories are created.

5. Add groups that users can be members of.

   ```
   groupadd -f users
   groupadd -f nobody
   groupadd -f ntadmins
   ```

6. Associate the UNIX groups with their respective Windows groups.

   ```
   net groupmap add ntgroup="Domain Users" unixgroup=users
   net groupmap add ntgroup="Domain Guests" unixgroup=nobody
   net groupmap add ntgroup="Domain Admins" unixgroup=ntadmins
   ```

7. Grant access rights to a user or a group. For example, to grant the right to add client machines to the domain on a Samba domain controller, to the members to the Domain Admins group, execute the following command:

   ```
   net rpc rights grant 'DOCS\Domain Admins' SetMachineAccountPrivilege -S PDC -U root
   ```

Keep in mind that Windows systems prefer to have a primary group which is mapped to a domain group such as Domain Users.

Windows groups and users use the same namespace thus not allowing the existence of a group and a user with the same name like in UNIX.

알림

If you need more than one domain controller or have more than 250 users, do not use a tdbsam authentication backend. LDAP is recommended in these cases.

21.6.3.2. Primary Domain Controller (PDC) with Active Directory

Although it is possible for Samba to be a member of an Active Directory, it is not possible for Samba to operate as an Active Directory domain controller.

21.7. Samba Security Modes

There are only two types of security modes for Samba, share-level and user-level, which are collectively known as security levels . Share-level security can only be implemented in one way, while user-level security can be implemented in one of four different ways. The different ways of implementing a security level are called security modes .

21.7.1. User-Level Security

User-level security is the default setting for Samba. Even if the security = user directive is not listed in the smb.conf file, it is used by Samba. If the server accepts the client's username/password, the client can then mount multiple shares without specifying a password for each instance. Samba can also accept session-based username/password requests. The client maintains multiple authentication contexts by using a unique UID for each logon.

In smb.conf, the security = user directive that sets user-level security is:

```
[GLOBAL]
...
security = user
...
```

The following sections describe other implementations of user-level security.

21.7.1.1. Domain Security Mode (User-Level Security)

In domain security mode, the Samba server has a machine account (domain security trust account) and causes all authentication requests to be passed through to the domain controllers. The Samba server is made into a domain member server by using the following directives in smb.conf:

```
[GLOBAL]
...
```

```
security = domain
workgroup = MARKETING
...
```

21.7.1.2. Active Directory Security Mode (User-Level Security)

If you have an Active Directory environment, it is possible to join the domain as a native Active Directory member. Even if a security policy restricts the use of NT-compatible authentication protocols, the Samba server can join an ADS using Kerberos. Samba in Active Directory member mode can accept Kerberos tickets.

In smb.conf, the following directives make Samba an Active Directory member server:

```
[GLOBAL]
...
security = ADS
realm = EXAMPLE.COM
password server = kerberos.example.com
...
```

21.7.1.3. Server Security Mode (User-Level Security)

Server security mode was previously used when Samba was not capable of acting as a domain member server.

> **알림**
>
> It is highly recommended to not use this mode since there are numerous security drawbacks.

In smb.conf, the following directives enable Samba to operate in server security mode:

```
[GLOBAL]
...
encrypt passwords = Yes
security = server
password server = "NetBIOS_of_Domain_Controller"
...
```

21.7.2. Share-Level Security

With share-level security, the server accepts only a password without an explicit username from the client. The server expects a password for each share, independent of the username. There have been recent reports that Microsoft Windows clients have compatibility issues with share-level security servers. Samba developers strongly discourage use of share-level security.

In smb.conf, the security = share directive that sets share-level security is:

```
[GLOBAL]
...
security = share
```

```
…
```

21.8. Samba Account Information Databases

The latest release of Samba offers many new features including new password database backends not previously available. Samba version 3.0.0 fully supports all databases used in previous versions of Samba. However, although supported, many backends may not be suitable for production use.

The following is a list different backends you can use with Samba. Other backends not listed here may also be available.

Plain Text

> Plain text backends are nothing more than the /etc/passwd type backends. With a plain text backend, all usernames and passwords are sent unencrypted between the client and the Samba server. This method is very unsecure and is not recommended for use by any means. It is possible that different Windows clients connecting to the Samba server with plain text passwords cannot support such an authentication method.

smbpasswd

> A popular backend used in previous Samba packages, the smbpasswd backend utilizes a plain ASCII text layout that includes the MS Windows LanMan and NT account, and encrypted password information. The smbpasswd backend lacks the storage of the Windows NT/2000/2003 SAM extended controls. The smbpasswd backend is not recommended because it does not scale well or hold any Windows information, such as RIDs for NT-based groups. The tdbsam backend solves these issues for use in a smaller database (250 users), but is still not an enterprise-class solution.

ldapsam_compat

> The ldapsam_compat backend allows continued OpenLDAP support for use with upgraded versions of Samba. This option normally used when migrating to Samba 3.0.

tdbsam

> The tdbsam backend provides an ideal database backend for local servers, servers that do not need built-in database replication, and servers that do not require the scalability or complexity of LDAP. The tdbsam backend includes all of the smbpasswd database information as well as the previously-excluded SAM information. The inclusion of the extended SAM data allows Samba to implement the same account and system access controls as seen with Windows NT/2000/2003-based systems.
>
> The tdbsam backend is recommended for 250 users at most. Larger organizations should require Active Directory or LDAP integration due to scalability and possible network infrastructure concerns.

ldapsam

> The ldapsam backend provides an optimal distributed account installation method for Samba. LDAP is optimal because of its ability to replicate its database to any number of servers using the OpenLDAP slurpd daemon. LDAP databases are light-weight and scalable, and as such are preferred by large enterprises.
>
> If you are upgrading from a previous version of Samba to 3.0, note that the /usr/share/doc/samba-<version>/LDAP/samba.schema has changed. This file contains the attribute syntax definitions and objectclass definitions that the ldapsam backend will need in order to function properly.

As such, if you are using the ldapsam backend for your Samba server, you will need to configure slapd to include this schema file. Refer to 26.5절. "The /etc/openldap/schema/ Directory" for directions on how to do this.

알림

You will need to have the openldap-server package installed if you want to use the ldapsam backend.

mysqlsam
: The mysqlsam backend uses a MySQL-based database backend. This is useful for sites that already implement MySQL. At present, mysqlsam is now packed in a module separate from Samba, and as such is not officially supported by Samba.

21.9. Samba Network Browsing

Network browsing enables Windows and Samba servers to appear in the Windows Network Neighborhood. Inside the Network Neighborhood, icons are represented as servers and if opened, the server's shares and printers that are available are displayed.

Network browsing capabilities require NetBIOS over TCP/IP. NetBIOS-based networking uses broadcast (UDP) messaging to accomplish browse list management. Without NetBIOS and WINS as the primary method for TCP/IP hostname resolution, other methods such as static files (/etc/hosts) or DNS, must be used.

A domain master browser collates the browse lists from local master browsers on all subnets so that browsing can occur between workgroups and subnets. Also, the domain master browser should preferably be the local master browser for its own subnet.

21.9.1. Domain Browsing

By default, a Windows server PDC for a domain is also the domain master browser for that domain. A Samba server must not be set up as a domain master server in this type of situation

For subnets that do not include the Windows server PDC, a Samba server can be implemented as a local master browser. Configuring the smb.conf for a local master browser (or no browsing at all) in a domain controller environment is the same as workgroup configuration.

21.9.2. WINS (Windows Internetworking Name Server)

Either a Samba server or a Windows NT server can function as a WINS server. When a WINS server is used with NetBIOS enabled, UDP unicasts can be routed which allows name resolution across networks. Without a WINS server, the UDP broadcast is limited to the local subnet and therefore cannot be routed to other subnets, workgroups, or domains. If WINS replication is necessary, do not use Samba as your primary WINS server, as Samba does not currently support WINS replication.

In a mixed NT/2000/2003 server and Samba environment, it is recommended that you use the Microsoft WINS capabilities. In a Samba-only environment, it is recommended that you use only one Samba server for WINS.

The following is an example of the smb.conf file in which the Samba server is serving as a WINS server:

```
[global]
wins support = Yes
```

Tip

All servers (including Samba) should connect to a WINS server to resolve NetBIOS names. Without WINS, browsing only occurs on the local subnet. Furthermore, even if a domain-wide list is somehow obtained, hosts cannot be resolved for the client without WINS.

21.10. Samba with CUPS Printing Support

Samba allows client machines to share printers connected to the Samba server. In addition, Samba also allows client machines to send documents built in Linux to Windows printer shares. Although there are other printing systems that function with Red Hat Enterprise Linux, CUPS (Common UNIX Print System) is the recommended printing system due to its close integration with Samba.

21.10.1. Simple smb.conf Settings

The following example shows a very basic smb.conf configuration for CUPS support:

```
[global]
load printers = Yes
printing = cups
printcap name = cups
[printers]
comment = All Printers
path = /var/spool/samba/print
printer = IBMInfoP
browseable = No
public = Yes
guest ok = Yes
writable = No
printable = Yes
printer admin = @ntadmins
[print$]
comment = Printer Drivers Share
path = /var/lib/samba/drivers
write list = ed, john
printer admin = ed, john
```

Other printing configurations are also possible. To add additional security and privacy for printing confidential documents, users can have their own print spooler not located in a public path. If a job fails, other users would not have access to the file.

The print$ share contains printer drivers for clients to access if not available locally. The print$ share is optional and may not be required depending on the organization.

Setting browseable to Yes enables the printer to be viewed in the Windows Network Neighborhood, provided the Samba server is set up correctly in the domain/workgroup.

21.11. Samba Distribution Programs

findsmb

findsmb <subnet_broadcast_address>

The findsmb program is a Perl script which reports information about SMB-aware systems on a specific subnet. If no subnet is specified the local subnet is used. Items displayed include IP address, NetBIOS name, workgroup or domain name, operating system, and version.

The following example shows the output of executing findsmb as any valid user on a system:

```
~]# findsmb
IP ADDR         NETBIOS NAME   WORKGROUP/OS/VERSION
------------------------------------------------------------------
10.1.59.25      VERVE           [MYGROUP] [Unix] [Samba 3.0.0-15]
10.1.59.26      STATION22       [MYGROUP] [Unix] [Samba 3.0.2-7.FC1]
10.1.56.45      TREK           +[WORKGROUP] [Windows 5.0] [Windows 2000 LAN Manager]
10.1.57.94      PIXEL           [MYGROUP] [Unix] [Samba 3.0.0-15]
10.1.57.137     MOBILE001       [WORKGROUP] [Windows 5.0] [Windows 2000 LAN Manager]
10.1.57.141     JAWS           +[KWIKIMART] [Unix] [Samba 2.2.7a-security-rollup-fix]
10.1.56.159     FRED           +[MYGROUP] [Unix] [Samba 3.0.0-14.3E]
10.1.59.192     LEGION         *[MYGROUP] [Unix] [Samba 2.2.7-security-rollup-fix]
10.1.56.205     NANCYN          +[MYGROUP] [Unix] [Samba 2.2.7a-security-rollup-fix]
```

net

net <protocol> <function> <misc_options> <target_options>

The net utility is similar to the net utility used for Windows and MS-DOS. The first argument is used to specify the protocol to use when executing a command. The <protocol> option can be ads, rap, or rpc for specifying the type of server connection. Active Directory uses ads, Win9x/NT3 uses rap, and Windows NT4/2000/2003 uses rpc. If the protocol is omitted, net automatically tries to determine it.

The following example displays a list the available shares for a host named wakko:

```
~]# net -l share -S wakko
Password:

Enumerating shared resources (exports) on remote server:

Share name   Type     Description
----------   ----     -----------
data         Disk     Wakko data share
tmp          Disk     Wakko tmp share
IPC$         IPC      IPC Service (Samba Server)
ADMIN$       IPC      IPC Service (Samba Server)
```

The following example displays a list of Samba users for a host named wakko:

```
~]# net -l user -S wakko
root password:
User name             Comment
-----------------------------
andriusb              Documentation
```

```
joe                      Marketing
lisa                     Sales
```

nmblookup

nmblookup <options> <netbios_name>

The nmblookup program resolves NetBIOS names into IP addresses. The program broadcasts its query on the local subnet until the target machine replies.

Here is an example:

```
~]# nmblookup trek
querying trek on 10.1.59.255
10.1.56.45 trek<00>
```

pdbedit

pdbedit <options>

The pdbedit program manages accounts located in the SAM database. All backends are supported including smbpasswd, LDAP, NIS+, and the tdb database library.

The following are examples of adding, deleting, and listing users:

```
~]# pdbedit -a kristin
new password:
retype new password:
Unix username:        kristin
NT username:
Account Flags:        [U          ]
User SID:             S-1-5-21-1210235352-3804200048-1474496110-2012
Primary Group SID:    S-1-5-21-1210235352-3804200048-1474496110-2077
Full Name: Home Directory:      \\wakko\kristin
HomeDir Drive:
Logon Script:
Profile Path:         \\wakko\kristin\profile
Domain:               WAKKO
Account desc:
Workstations: Munged
dial:
Logon time:           0
Logoff time:          Mon, 18 Jan 2038 22:14:07 GMT
Kickoff time:         Mon, 18 Jan 2038 22:14:07 GMT
Password last set:    Thu, 29 Jan 2004 08:29:28
GMT Password can change:  Thu, 29 Jan 2004 08:29:28 GMT
Password must change: Mon, 18 Jan 2038 22:14:07 GMT

~]# pdbedit -v -L kristin
Unix username:        kristin
NT username:
Account Flags:        [U          ]
User SID:             S-1-5-21-1210235352-3804200048-1474496110-2012
Primary Group SID:    S-1-5-21-1210235352-3804200048-1474496110-2077
Full Name:
Home Directory:       \\wakko\kristin
HomeDir Drive:
Logon Script:
Profile Path:         \\wakko\kristin\profile
```

```
Domain:                    WAKKO
Account desc:
Workstations: Munged
dial:
Logon time:               0
Logoff time:              Mon, 18 Jan 2038 22:14:07 GMT
Kickoff time:             Mon, 18 Jan 2038 22:14:07 GMT
Password last set:      Thu, 29 Jan 2004 08:29:28 GMT
Password can change:  Thu, 29 Jan 2004 08:29:28 GMT
Password must change: Mon, 18 Jan 2038 22:14:07 GMT

~]# pdbedit -L
andriusb:505:
joe:503:
lisa:504:
kristin:506:

~]# pdbedit -x joe
~]# pdbedit -L

andriusb:505: lisa:504: kristin:506:
```

rpcclient

rpcclient <server> <options>

The rpcclient program issues administrative commands using Microsoft RPCs, which provide access to the Windows administration graphical user interfaces (GUIs) for systems management. This is most often used by advanced users that understand the full complexity of Microsoft RPCs.

smbcacls

smbcacls <//server/share> <filename> <options>

The smbcacls program modifies Windows ACLs on files and directories shared by the Samba server.

smbclient

smbclient <//server/share> <password> <options>

The smbclient program is a versatile UNIX client which provides functionality similar to ftp.

smbcontrol

smbcontrol -i <options>

smbcontrol <options> <destination> <messagetype> <parameters>

The smbcontrol program sends control messages to running smbd or nmbd daemons. Executing smbcontrol -i runs commands interactively until a blank line or a 'q' is entered.

smbpasswd

smbpasswd <options> <username> <password>

The smbpasswd program manages encrypted passwords. This program can be run by a superuser to change any user's password as well as by an ordinary user to change their own Samba password.

smbspool

smbspool <job> <user> <title> <copies> <options> <filename>

The smbspool program is a CUPS-compatible printing interface to Samba. Although designed for use with CUPS printers, smbspool can work with non-CUPS printers as well.

smbstatus

smbstatus <options>

The smbstatus program displays the status of current connections to a Samba server.

smbtar

smbtar <options>

The smbtar program performs backup and restores of Windows-based share files and directories to a local tape archive. Though similar to the tar command, the two are not compatible.

testparm

testparm <options> <filename> <hostname IP_address>

The testparm program checks the syntax of the smb.conf file. If your smb.conf file is in the default location (/etc/samba/smb.conf) you do not need to specify the location. Specifying the hostname and IP address to the testparm program verifies that the hosts.allow and host.deny files are configured correctly. The testparm program also displays a summary of your smb.conf file and the server's role (stand-alone, domain, etc.) after testing. This is convenient when debugging as it excludes comments and concisely presents information for experienced administrators to read.

For example:

```
~]# testparm
Load smb config files from /etc/samba/smb.conf
Processing section "[homes]"
Processing section "[printers]"
Processing section "[tmp]"
Processing section "[html]"
Loaded services file OK.
Server role: ROLE_STANDALONE
Press enter to see a dump of your service definitions

<enter>
# Global parameters
[global]
 workgroup = MYGROUP
 server string = Samba Server
 security = SHARE
 log file = /var/log/samba/%m.log
 max log size = 50
 socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192
 dns proxy = No
[homes]
 comment = Home Directories
 read only = No
 browseable = No
[printers]
 comment = All Printers
 path = /var/spool/samba
 printable = Yes
 browseable = No
[tmp]
 comment = Wakko tmp
 path = /tmp
 guest only = Yes
[html]
 comment = Wakko www
 path = /var/www/html
 force user = andriusb
 force group = users
 read only = No
 guest only = Yes
```

wbinfo

wbinfo <options>

The wbinfo program displays information from the winbindd daemon. The winbindd daemon must be running for wbinfo to work.

21.12. Additional Resources

The following sections give you the means to explore Samba in greater detail.

21.12.1. 설치된 문서

- /usr/share/doc/samba-<version-number>/ — All additional files included with the Samba distribution. This includes all helper scripts, sample configuration files, and documentation.

  This directory also contains online versions of The Official Samba-3 HOWTO-Collection and Samba-3 by Example, both of which are cited below.

21.12.2. Related Books

- The Official Samba-3 HOWTO-Collection by John H. Terpstra and Jelmer R. Vernooij; Prentice Hall — The official Samba-3 documentation as issued by the Samba development team. This is more of a reference guide than a step-by-step guide.
- Samba-3 by Example by John H. Terpstra; Prentice Hall — This is another official release issued by the Samba development team which discusses detailed examples of OpenLDAP, DNS, DHCP, and printing configuration files. This has step-by-step related information that helps in real-world implementations.
- Using Samba, 2nd Edition by Jay T's, Robert Eckstein, and David Collier-Brown; O'Reilly — A good resource for novice to advanced users, which includes comprehensive reference material.

21.12.3. 유용한 웹사이트

- http://www.samba.org/ — Homepage for the Samba distribution and all official documentation created by the Samba development team. Many resources are available in HTML and PDF formats, while others are only available for purchase. Although many of these links are not Red Hat Enterprise Linux specific, some concepts may apply.

- http://samba.org/samba/archives.html [1] — Active email lists for the Samba community. Enabling digest mode is recommended due to high levels of list activity.

- Samba newsgroups — Samba threaded newsgroups, such as gmane.org, that use the NNTP protocol are also available. This an alternative to receiving mailing list emails.

[1] http://us1.samba.org/samba/archives.html

동적 호스트 설정 프로토콜 (DHCP)

Dynamic Host Configuration Protocol (DHCP) is a network protocol that automatically assigns TCP/IP information to client machines. Each DHCP client connects to the centrally located DHCP server, which returns that client's network configuration (including the IP address, gateway, and DNS servers).

22.1. DHCP를 사용하는 이유?

DHCP is useful for automatic configuration of client network interfaces. When configuring the client system, the administrator chooses DHCP instead of specifying an IP address, netmask, gateway, or DNS servers. The client retrieves this information from the DHCP server. DHCP is also useful if an administrator wants to change the IP addresses of a large number of systems. Instead of reconfiguring all the systems, he can just edit one DHCP configuration file on the server for the new set of IP addresses. If the DNS servers for an organization changes, the changes are made on the DHCP server, not on the DHCP clients. When the administrator restarts the network or reboots the clients, the changes will go into effect.

If an organization has a functional DHCP server properly connected to a network, laptops and other mobile computer users can move these devices from office to office.

22.2. DHCP 서버 설정

The dhcp package contains an ISC DHCP server. First, install the package as the superuser:

```
~]# yum install dhcp
```

Installing the dhcp package creates a file, /etc/dhcpd.conf, which is merely an empty configuration file:

```
~]# cat /etc/dhcpd.conf
#
# DHCP Server Configuration file.
#    see /usr/share/doc/dhcp*/dhcpd.conf.sample
```

The sample configuration file can be found at /usr/share/doc/dhcp-<version>/dhcpd.conf.sample. You should use this file to help you configure /etc/dhcpd.conf, which is explained in detail below.

DHCP also uses the file /var/lib/dhcpd/dhcpd.leases to store the client lease database. Refer to 22.2.2 절. “할당 (Lease) 데이터베이스” for more information.

22.2.1. 설정 파일

The first step in configuring a DHCP server is to create the configuration file that stores the network information for the clients. Use this file to declare options and global options for client systems.

The configuration file can contain extra tabs or blank lines for easier formatting. Keywords are case-insensitive and lines beginning with a hash mark (#) are considered comments.

Two DNS update schemes are currently implemented — the ad-hoc DNS update mode and the interim DHCP-DNS interaction draft update mode. If and when these two are accepted as part of the Internet

Engineering Task Force (IETF) standards process, there will be a third mode — the standard DNS update method. You must configure the DNS server for compatibility with these schemes. Version 3.0b2pl11 and previous versions used the ad-hoc mode; however, it has been deprecated. To keep the same behavior, add the following line to the top of the configuration file:

```
ddns-update-style ad-hoc;
```

추천된 모드를 사용하시려면, 설정 파일 처음 부분에 다음과 같은 줄을 추가해 주십시오:

```
ddns-update-style interim;
```

Refer to the dhcpd.conf man page for details about the different modes.

설정 파일에는 다음과 같은 두가지 유형의 문장 (statememt)이 사용됩니다:

- Parameters — State how to perform a task, whether to perform a task, or what network configuration options to send to the client.
- Declarations — Describe the topology of the network, describe the clients, provide addresses for the clients, or apply a group of parameters to a group of declarations.

The parameters that start with the keyword option are referred to as options. These options control DHCP options; whereas, parameters configure values that are not optional or control how the DHCP server behaves.

중 괄호 ({ }) 내에 포함된 부분 이전에 선언된 (옵션을 포함한) 매개 변수는 전역 매개 변수 (global parameter)로 취급됩니다. 전역 매개 변수는 자신 이하에 위치한 모든 부분에 적용됩니다.

중요

If the configuration file is changed, the changes do not take effect until the DHCP daemon is restarted with the command service dhcpd restart.

Tip

Instead of changing a DHCP configuration file and restarting the service each time, using the omshell command provides an interactive way to connect to, query, and change the configuration of a DHCP server. By using omshell, all changes can be made while the server is running. For more information on omshell, refer to the omshell man page.

In 예 22.1. "subnet 선언" , the routers, subnet-mask, domain-name, domain-name-servers, and time-offset options are used for any host statements declared below it.

Additionally, a subnet can be declared, a subnet declaration must be included for every subnet in the network. If it is not, the DHCP server fails to start.

In this example, there are global options for every DHCP client in the subnet and a range declared. Clients are assigned an IP address within the range.

예 22.1. subnet 선언

```
subnet 192.168.1.0 netmask 255.255.255.0 {
        option routers                  192.168.1.254;
        option subnet-mask              255.255.255.0;

        option domain-name              "example.com";
        option domain-name-servers      192.168.1.1;

        option time-offset              -18000;     # Eastern Standard Time

 range 192.168.1.10 192.168.1.100;
}
```

All subnets that share the same physical network should be declared within a shared-network declaration as shown in 예 22.2. “share-network (공유-네트워크) 선언”. Parameters within the shared-network, but outside the enclosed subnet declarations, are considered to be global parameters. The name of the shared-network must be a descriptive title for the network, such as using the title 'test-lab' to describe all the subnets in a test lab environment.

예 22.2. share-network (공유-네트워크) 선언

```
shared-network name {
    option domain-name              "test.redhat.com";
    option domain-name-servers      ns1.redhat.com, ns2.redhat.com;
    option routers                  192.168.0.254;
    more parameters for EXAMPLE shared-network
    subnet 192.168.1.0 netmask 255.255.252.0 {
        parameters for subnet
        range 192.168.1.1 192.168.1.254;
    }
    subnet 192.168.2.0 netmask 255.255.252.0 {
        parameters for subnet
        range 192.168.2.1 192.168.2.254;
    }
}
```

As demonstrated in 예 22.3. “group 선언”, the group declaration is used to apply global parameters to a group of declarations. For example, shared networks, subnets, and hosts can be grouped.

예 22.3. group 선언

```
group {
   option routers                  192.168.1.254;
   option subnet-mask              255.255.255.0;

   option domain-name              "example.com";
   option domain-name-servers      192.168.1.1;

   option time-offset              -18000;     # Eastern Standard Time

   host apex {
      option host-name "apex.example.com";
```

```
        hardware ethernet 00:A0:78:8E:9E:AA;
        fixed-address 192.168.1.4;
    }

    host raleigh {
        option host-name "raleigh.example.com";
        hardware ethernet 00:A1:DD:74:C3:F2;
        fixed-address 192.168.1.6;
    }
}
```

To configure a DHCP server that leases a dynamic IP address to a system within a subnet, modify 예 22.4. "Range 매개 변수" with your values. It declares a default lease time, maximum lease time, and network configuration values for the clients. This example assigns IP addresses in the range 192.168.1.10 and 192.168.1.100 to client systems.

예 22.4. Range 매개 변수

```
default-lease-time 600;
max-lease-time 7200;
option subnet-mask 255.255.255.0;
option broadcast-address 192.168.1.255;
option routers 192.168.1.254;
option domain-name-servers 192.168.1.1, 192.168.1.2;
option domain-name "example.com";

subnet 192.168.1.0 netmask 255.255.255.0 {
    range 192.168.1.10 192.168.1.100;
}
```

To assign an IP address to a client based on the MAC address of the network interface card, use the hardware ethernet parameter within a host declaration. As demonstrated in 예 22.5. "DHCP를 사용하는 정적 IP 주소", the host apex declaration specifies that the network interface card with the MAC address 00:A0:78:8E:9E:AA always receives the IP address 192.168.1.4.

Note that the optional parameter host-name can also be used to assign a host name to the client.

예 22.5. DHCP를 사용하는 정적 IP 주소

```
host apex {
    option host-name "apex.example.com";
    hardware ethernet 00:A0:78:8E:9E:AA;
    fixed-address 192.168.1.4;
}
```

Tip

The sample configuration file provided can be used as a starting point and custom configuration options can be added to it. To copy it to the proper location, use the following command:

```
cp /usr/share/doc/dhcp-<version-number>/dhcpd.conf.sample /etc/dhcpd.conf
```

(where <version-number> is the DHCP version number).

For a complete list of option statements and what they do, refer to the dhcp-options man page.

22.2.2. 할당 (Lease) 데이터베이스

On the DHCP server, the file /var/lib/dhcpd/dhcpd.leases stores the DHCP client lease database. Do not change this file. DHCP lease information for each recently assigned IP address is automatically stored in the lease database. The information includes the length of the lease, to whom the IP address has been assigned, the start and end dates for the lease, and the MAC address of the network interface card that was used to retrieve the lease.

All times in the lease database are in Coordinated Universal Time (UTC), not local time.

The lease database is recreated from time to time so that it is not too large. First, all known leases are saved in a temporary lease database. The dhcpd.leases file is renamed dhcpd.leases~ and the temporary lease database is written to dhcpd.leases.

The DHCP daemon could be killed or the system could crash after the lease database has been renamed to the backup file but before the new file has been written. If this happens, the dhcpd.leases file does not exist, but it is required to start the service. Do not create a new lease file. If you do, all old leases are lost which causes many problems. The correct solution is to rename the dhcpd.leases~ backup file to dhcpd.leases and then start the daemon.

22.2.3. 서버 시작과 중지

중요

When the DHCP server is started for the first time, it fails unless the dhcpd.leases file exists. Use the command touch /var/lib/dhcpd/dhcpd.leases to create the file if it does not exist.

If the same server is also running BIND as a DNS server, this step is not necessary, as starting the named service automatically checks for a dhcpd.leases file.

To start the DHCP service, use the command /sbin/service dhcpd start. To stop the DHCP server, use the command /sbin/service dhcpd stop.

By default, the DHCP service does not start at boot time. To configure the daemon to start automatically at boot time, refer to 17장. 서비스로의 접근 통제.

If more than one network interface is attached to the system, but the DHCP server should only be started on one of the interfaces, configure the DHCP server to start only on that device. In /etc/sysconfig/dhcpd, add the name of the interface to the list of DHCPDARGS:

```
# Command line options here
DHCPDARGS=eth0
```

두개의 네트워크 카드를 가진 방화벽 컴퓨터에 이 옵션이 유용합니다. 한개의 네트워크 카드는 인터넷 IP 주소를 검색하도록 DHCP 클라이언트로 설정 가능하며, 다른 네트워크 카드는 방화벽 뒤에서 내부 네트워크 용 DHCP 서버로 사용하실 수 있습니다. 내부 네트워크에 연결된 네트워크 카드만을 지정함으로서 다른 사용자가 인터넷을 통하여 데몬에 접속하지 못하게 되므로 시스템 보안이 더욱 강화됩니다.

Other command line options that can be specified in /etc/sysconfig/dhcpd include:

- -p <portnum> — Specifies the UDP port number on which dhcpd should listen. The default is port 67. The DHCP server transmits responses to the DHCP clients at a port number one greater than the UDP port specified. For example, if the default port 67 is used, the server listens on port 67 for requests and responses to the client on port 68. If a port is specified here and the DHCP relay agent is used, the same port on which the DHCP relay agent should listen must be specified. Refer to 22.2.4절. "DHCP 릴레이 에이전트 (Relay Agent)" for details.

- -f — Runs the daemon as a foreground process. This is mostly used for debugging.

- -d — Logs the DHCP server daemon to the standard error descriptor. This is mostly used for debugging. If this is not specified, the log is written to /var/log/messages.

- -cf <filename> — Specifies the location of the configuration file. The default location is /etc/dhcpd.conf.

- -lf <filename> — Specifies the location of the lease database file. If a lease database file already exists, it is very important that the same file be used every time the DHCP server is started. It is strongly recommended that this option only be used for debugging purposes on non-production machines. The default location is /var/lib/dhcpd/dhcpd.leases.

- -q — Do not print the entire copyright message when starting the daemon.

22.2.4. DHCP 릴레이 에이전트 (Relay Agent)

The DHCP Relay Agent (dhcrelay) allows for the relay of DHCP and BOOTP requests from a subnet with no DHCP server on it to one or more DHCP servers on other subnets.

DHCP 클라이언트가 정보를 요청하는 경우, DHCP 릴레이 에이전트는 그 요청을 DHCP 릴레이 에이전트가 시작될 때 지정된 DHCP 서버 목록으로 전송합니다. DHCP 서버가 응답을 보내오면, 원래 요청을 보낸 네트워크 상에서 그 응답을 브로드캐스트 (broadcast)하거나 유니캐스트 (unicast)합니다.

The DHCP Relay Agent listens for DHCP requests on all interfaces unless the interfaces are specified in /etc/sysconfig/dhcrelay with the INTERFACES directive.

To start the DHCP Relay Agent, use the command service dhcrelay start.

22.3. DHCP 클라이언트 설정

The first step for configuring a DHCP client is to make sure the kernel recognizes the network interface card. Most cards are recognized during the installation process and the system is configured to use the correct kernel module for the card. If a card is added after installation, Kudzu [1] will recognize it and prompt you for the proper kernel module (Be sure to check the Hardware Compatibility List at http://hardware.redhat.com/hcl/). If either the installation program or kudzu does not recognize the network card, you can load the correct kernel module (refer to 43장. General Parameters and Modules for details).

To configure a DHCP client manually, modify the /etc/sysconfig/network file to enable networking and the configuration file for each network device in the /etc/sysconfig/network-scripts directory. In this directory, each device should have a configuration file named ifcfg-eth0, where eth0 is the network device name.

The /etc/sysconfig/network file should contain the following line:

```
NETWORKING=yes
```

The NETWORKING variable must be set to yes if you want networking to start at boot time.

The /etc/sysconfig/network-scripts/ifcfg-eth0 file should contain the following lines:

```
DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
```

DHCP를 사용하도록 설정하기 위해서는 각 장치마다 설정 파일이 필요합니다.

Other options for the network script includes:

- DHCP_HOSTNAME — Only use this option if the DHCP server requires the client to specify a hostname before receiving an IP address. (The DHCP server daemon in Red Hat Enterprise Linux does not support this feature.)
- PEERDNS=<answer> , where <answer> is one of the following:
  - yes — Modify /etc/resolv.conf with information from the server. If using DHCP, then yes is the default.
  - no — Do not modify /etc/resolv.conf.
- SRCADDR=<address> , where <address> is the specified source IP address for outgoing packets.
- USERCTL=<answer> , where <answer> is one of the following:
  - yes — Non-root users are allowed to control this device.
  - no — Non-root users are not allowed to control this device.

If you prefer using a graphical interface, refer to 16장. 네트워크 설정 for instructions on using the Network Administration Tool to configure a network interface to use DHCP.

[1] Kudzu is a hardware probing tool run at system boot time to determine what hardware has been added or removed from the system.

Tip

For advanced configurations of client DHCP options such as protocol timing, lease requirements and requests, dynamic DNS support, aliases, as well as a wide variety of values to override, prepend, or append to client-side configurations, refer to the dhclient and dhclient.conf man pages.

22.4. Configuring a Multihomed DHCP Server

A multihomed DHCP server serves multiple networks, that is, multiple subnets. The examples in these sections detail how to configure a DHCP server to serve multiple networks, select which network interfaces to listen on, and how to define network settings for systems that move networks.

Before making any changes, back up the existing /etc/sysconfig/dhcpd and /etc/dhcpd.conf files.

The DHCP daemon listens on all network interfaces unless otherwise specified. Use the /etc/sysconfig/dhcpd file to specify which network interfaces the DHCP daemon listens on. The following /etc/sysconfig/dhcpd example specifies that the DHCP daemon listens on the eth0 and eth1 interfaces:

```
DHCPDARGS="eth0 eth1";
```

If a system has three network interfaces cards -- eth0, eth1, and eth2 -- and it is only desired that the DHCP daemon listens on eth0, then only specify eth0 in /etc/sysconfig/dhcpd:

```
DHCPDARGS="eth0";
```

The following is a basic /etc/dhcpd.conf file, for a server that has two network interfaces, eth0 in a 10.0.0.0/24 network, and eth1 in a 172.16.0.0/24 network. Multiple subnet declarations allow different settings to be defined for multiple networks:

```
ddns-update-style interim;
default-lease-time 600;
max-lease-time 7200;

subnet 10.0.0.0 netmask 255.255.255.0 {
 option subnet-mask 255.255.255.0;
 option routers 10.0.0.1;
 range 10.0.0.5 10.0.0.15;
}

subnet 172.16.0.0 netmask 255.255.255.0 {
 option subnet-mask 255.255.255.0;
 option routers 172.16.0.1;
 range 172.16.0.5 172.16.0.15;

}
```

subnet 10.0.0.0 netmask 255.255.255.0

> A subnet declaration is required for every network your DHCP server is serving. Multiple subnets require multiple subnet declarations. If the DHCP server does not have a network interface in a range of a subnet declaration, the DHCP server does not serve that network.

If there is only one subnet declaration, and no network interfaces are in the range of that subnet, the DHCP daemon fails to start, and an error such as the following is logged to /var/log/messages:

```
dhcpd: No subnet declaration for eth0 (0.0.0.0).
dhcpd: ** Ignoring requests on eth0.  If this is not what
dhcpd:    you want, please write a subnet declaration
dhcpd:    in your dhcpd.conf file for the network segment
dhcpd:    to which interface eth1 is attached. **
dhcpd:
dhcpd:
dhcpd: Not configured to listen on any interfaces!
```

option subnet-mask 255.255.255.0;

The option subnet-mask option defines a subnet mask, and overrides the netmask value in the subnet declaration. In simple cases, the subnet and netmask values are the same.

option routers 10.0.0.1;

The option routers option defines the default gateway for the subnet. This is required for systems to reach internal networks on a different subnet, as well as external networks.

range 10.0.0.5 10.0.0.15;

The range option specifies the pool of available IP addresses. Systems are assigned an address from the range of specified IP addresses.

For further information, refer to the dhcpd.conf(5) man page.

Alias Interfaces

Alias interfaces are not supported by DHCP. If an alias interface is the only interface, in the only subnet specified in /etc/dhcpd.conf, the DHCP daemon fails to start.

22.4.1. Host Configuration

Before making any changes, back up the existing /etc/sysconfig/dhcpd and /etc/dhcpd.conf files.

Configuring a single system for multiple networks

The following /etc/dhcpd.conf example creates two subnets, and configures an IP address for the same system, depending on which network it connects to:

```
ddns-update-style interim;
default-lease-time 600;
max-lease-time 7200;

subnet 10.0.0.0 netmask 255.255.255.0 {
 option subnet-mask 255.255.255.0;
 option routers 10.0.0.1;
 range 10.0.0.5 10.0.0.15;
}

subnet 172.16.0.0 netmask 255.255.255.0 {
 option subnet-mask 255.255.255.0;
 option routers 172.16.0.1;
 range 172.16.0.5 172.16.0.15;
}

host example0 {
 hardware ethernet 00:1A:6B:6A:2E:0B;
 fixed-address 10.0.0.20;
}

host example1 {
 hardware ethernet 00:1A:6B:6A:2E:0B;
 fixed-address 172.16.0.20;
}
```

host example0

The host declaration defines specific parameters for a single system, such as an IP address. To configure specific parameters for multiple hosts, use multiple host declarations.

Most DHCP clients ignore the name in host declarations, and as such, this name can anything, as long as it is unique to other host declarations. To configure the same system for multiple networks, use a different name for each host declaration, otherwise the DHCP daemon fails to start. Systems are identified by the hardware ethernet option, not the name in the host declaration.

hardware ethernet 00:1A:6B:6A:2E:0B;

The hardware ethernet option identifies the system. To find this address, run the ifconfig command on the desired system, and look for the HWaddr address.

fixed-address 10.0.0.20;

The fixed-address option assigns a valid IP address to the system specified by the hardware ethernet option. This address must be outside the IP address pool specified with the range option.

If option statements do not end with a semicolon, the DHCP daemon fails to start, and an error such as the following is logged to /var/log/messages:

```
/etc/dhcpd.conf line 20: semicolon expected.
dhcpd: }
dhcpd: ^
dhcpd: /etc/dhcpd.conf line 38: unexpected end of file
dhcpd:
dhcpd: ^
dhcpd: Configuration file errors encountered -- exiting
```

Configuring systems with multiple network interfaces

The following host declarations configure a single system, that has multiple network interfaces, so that each interface receives the same IP address. This configuration will not work if both network interfaces are connected to the same network at the same time:

```
host interface0 {
 hardware ethernet 00:1a:6b:6a:2e:0b;
 fixed-address 10.0.0.18;
}

host interface1 {
 hardware ethernet 00:1A:6B:6A:27:3A;
 fixed-address 10.0.0.18;
}
```

For this example, interface0 is the first network interface, and interface1 is the second interface. The different hardware ethernet options identify each interface.

If such a system connects to another network, add more host declarations, remembering to:

- assign a valid fixed-address for the network the host is connecting to.
- make the name in the host declaration unique.

When a name given in a host declaration is not unique, the DHCP daemon fails to start, and an error such as the following is logged to /var/log/messages:

```
dhcpd: /etc/dhcpd.conf line 31: host interface0: already exists
dhcpd: }
dhcpd: ^
dhcpd: Configuration file errors encountered -- exiting
```

This error was caused by having multiple host interface0 declarations defined in /etc/dhcpd.conf.

22.5. 추가 자료

For additional configuration options, refer to the following resources.

22.5.1. 설치된 문서 자료

- dhcpd man page — Describes how the DHCP daemon works.
- dhcpd.conf man page — Explains how to configure the DHCP configuration file; includes some examples.
- dhcpd.leases man page — Explains how to configure the DHCP leases file; includes some examples.
- dhcp-options man page — Explains the syntax for declaring DHCP options in dhcpd.conf; includes some examples.
- dhcrelay man page — Explains the DHCP Relay Agent and its configuration options.
- /usr/share/doc/dhcp-<version>/ — Contains sample files, README files, and release notes for current versions of the DHCP service.

Apache HTTP Server

The Apache HTTP Server is a robust, commercial-grade open source Web server developed by the Apache Software Foundation (http://www.apache.org/). Red Hat Enterprise Linux includes the Apache HTTP Server 2.2 as well as a number of server modules designed to enhance its functionality.

The default configuration file installed with the Apache HTTP Server works without alteration for most situations. This chapter outlines many of the directives found within its configuration file (/etc/httpd/conf/httpd.conf) to aid those who require a custom configuration or need to convert a configuration file from the older Apache HTTP Server 1.3 format.

Warning

If using the graphical HTTP Configuration Tool (system-config-httpd), do not hand edit the Apache HTTP Server's configuration file as the HTTP Configuration Tool regenerates this file whenever it is used.

23.1. Apache HTTP Server 2.2

There are important differences between the Apache HTTP Server 2.2 and version 2.0 (version 2.0 shipped with Red Hat Enterprise Linux 4 and earlier). This section reviews some of the features of Apache HTTP Server 2.2 and outlines important changes. If you are upgrading from version 1.3, you should also read the instructions on migrating from version 1.3 to version 2.0. For instructions on migrating a version 1.3 configuration file to the 2.0 format, refer to 23.2.2절. "Migrating Apache HTTP Server 1.3 Configuration Files to 2.0" .

23.1.1. Features of Apache HTTP Server 2.2

Apache HTTP Server 2.2 features the following improvements over version 2.0 :

- Improved caching modules (mod_cache, mod_disk_cache, mod_mem_cache).
- A new structure for authentication and authorization support, replacing the authentication modules provided in previous versions.
- Support for proxy load balancing (mod_proxy_balancer)
- support for handling large files (namely, greater than 2GB) on 32-bit platforms

The following changes have been made to the default httpd configuration:

- The mod_cern_meta and mod_asis modules are no longer loaded by default.
- The mod_ext_filter module is now loaded by default.

If upgrading from a previous release of Red Hat Enterprise Linux, the httpd configuration will need to be updated for httpd 2.2. For more information, refer to http://httpd.apache.org/docs/2.2/upgrading.html

23.2. Migrating Apache HTTP Server Configuration Files

23.2.1. Migrating Apache HTTP Server 2.0 Configuration Files

This section outlines migration from version 2.0 to 2.2. If you are migrating from version 1.3, please refer to 23.2.2절. "Migrating Apache HTTP Server 1.3 Configuration Files to 2.0" .

- Configuration files and startup scripts from version 2.0 need minor adjustments particularly in module names which may have changed. Third party modules which worked in version 2.0 can also work in version 2.2 but need to be recompiled before you load them. Key modules that need to be noted are authentication and authorization modules. For each of the modules which has been renamed the LoadModule[1] line will need to be updated.

- The mod_userdir module will only act on requests if you provide a UserDir directive indicating a directory name. If you wish to maintain the procedures used in version 2.0, add the directive UserDir public_html in your configuration file.

- To enable SSL, edit the httpd.conf file adding the necessary mod_ssl directives. Use apachectl start as apachectl startssl is unavailable in version 2.2. You can view an example of SSL configuration for httpd in conf/extra/httpd-ssl.conf.

- To test your configuration it is advisable to use service httpd configtest which will detect configuration errors.

More information on upgrading from version 2.0 to 2.2 can be found on http://httpd.apache.org/docs/2.2/upgrading.html.

23.2.2. Migrating Apache HTTP Server 1.3 Configuration Files to 2.0

This section details migrating an Apache HTTP Server 1.3 configuration file to be utilized by Apache HTTP Server 2.0.

If upgrading to Red Hat Enterprise Linux 5 from Red Hat Enterprise Linux 2.1, note that the new stock configuration file for the Apache HTTP Server 2.0 package is installed as /etc/httpd/conf/httpd.conf.rpmnew and the original version 1.3 httpd.conf is left untouched. It is entirely up to you whether to use the new configuration file and migrate the old settings to it, or use the existing file as a base and modify it to suit; however, some parts of the file have changed more than others and a mixed approach is generally the best. The stock configuration files for both version 1.3 and 2.0 are divided into three sections.

If the /etc/httpd/conf/httpd.conf file is a modified version of the newly installed default and a saved a copy of the original configuration file is available, it may be easiest to invoke the diff command, as in the following example (logged in as root):

```
diff -u httpd.conf.orig httpd.conf | less
```

This command highlights any modifications made. If a copy of the original file is not available, extract it from an RPM package using the rpm2cpio and cpio commands, as in the following example:

```
rpm2cpio apache-<version-number>.i386.rpm | cpio -i --make
```

In the above command, replace <version-number> with the version number for the apache package.

[1] http://httpd.apache.org/docs/2.2/mod/mod_so.html#loadmodule

Finally, it is useful to know that the Apache HTTP Server has a testing mode to check for configuration errors. To use access it, type the following command:

```
apachectl configtest
```

23.2.2.1. Global Environment Configuration

The global environment section of the configuration file contains directives which affect the overall operation of the Apache HTTP Server, such as the number of concurrent requests it can handle and the locations of the various files. This section requires a large number of changes and should be based on the Apache HTTP Server 2.0 configuration file, while migrating the old settings into it.

23.2.2.1.1. Interface and Port Binding

The BindAddress and Port directives no longer exist; their functionality is now provided by a more flexible Listen directive.

If Port 80 was set in the 1.3 version configuration file, change it to Listen 80 in the 2.0 configuration file. If Port was set to some value other than 80, then append the port number to the contents of the ServerName directive.

For example, the following is a sample Apache HTTP Server 1.3 directive:

```
Port 123
ServerName www.example.com
```

To migrate this setting to Apache HTTP Server 2.0, use the following structure:

```
Listen 123
ServerName www.example.com:123
```

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mpm_common.html#listen

- http://httpd.apache.org/docs-2.0/mod/core.html#servername

23.2.2.1.2. Server-Pool Size Regulation

When the Apache HTTP Server accepts requests, it dispatches child processes or threads to handle them. This group of child processes or threads is known as a server-pool. Under Apache HTTP Server 2.0, the responsibility for creating and maintaining these server-pools has been abstracted to a group of modules called Multi-Processing Modules (MPMs). Unlike other modules, only one module from the MPM group can be loaded by the Apache HTTP Server. There are three MPM modules that ship with 2.0: prefork, worker, and perchild. Currently only the prefork and worker MPMs are available, although the perchild MPM may be available at a later date.

The original Apache HTTP Server 1.3 behavior has been moved into the prefork MPM. The prefork MPM accepts the same directives as Apache HTTP Server 1.3, so the following directives may be migrated directly:

- StartServers
- MinSpareServers
- MaxSpareServers
- MaxClients
- MaxRequestsPerChild

The worker MPM implements a multi-process, multi-threaded server providing greater scalability. When using this MPM, requests are handled by threads, conserving system resources and allowing large numbers of requests to be served efficiently. Although some of the directives accepted by the worker MPM are the same as those accepted by the prefork MPM, the values for those directives should not be transferred directly from an Apache HTTP Server 1.3 installation. It is best to instead use the default values as a guide, then experiment to determine what values work best.

Important

To use the worker MPM, create the file /etc/sysconfig/httpd and add the following directive:

```
HTTPD=/usr/sbin/httpd.worker
```

For more on the topic of MPMs, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mpm.html

23.2.2.1.3. Dynamic Shared Object (DSO) Support

There are many changes required here, and it is highly recommended that anyone trying to modify an Apache HTTP Server 1.3 configuration to suit version 2.0 (as opposed to migrating the changes into the version 2.0 configuration) copy this section from the stock Apache HTTP Server 2.0 configuration file.

Those who do not want to copy the section from the stock Apache HTTP Server 2.0 configuration should note the following:

- The AddModule and ClearModuleList directives no longer exist. These directives where used to ensure that modules could be enabled in the correct order. The Apache HTTP Server 2.0 API allows modules to specify their ordering, eliminating the need for these two directives.
- The order of the LoadModule lines are no longer relevant in most cases.
- Many modules have been added, removed, renamed, split up, or incorporated into others.
- LoadModule lines for modules packaged in their own RPMs (mod_ssl, php, mod_perl, and the like) are no longer necessary as they can be found in their relevant files within the /etc/httpd/conf.d/ directory.
- The various HAVE_XXX definitions are no longer defined.

Important

If modifying the original file, note that it is of paramount importance that the httpd.conf contains the following directive:

```
Include conf.d/*.conf
```

Omission of this directive results in the failure of all modules packaged in their own RPMs (such as mod_perl, php, and mod_ssl).

23.2.2.1.4. Other Global Environment Changes

The following directives have been removed from Apache HTTP Server 2.0's configuration:

- ServerType — The Apache HTTP Server can only be run as ServerType standalone making this directive irrelevant.
- AccessConfig and ResourceConfig — These directives have been removed as they mirror the functionality of the Include directive. If the AccessConfig and ResourceConfig directives are set, replace them with Include directives.

  To ensure that the files are read in the order implied by the older directives, the Include directives should be placed at the end of the httpd.conf, with the one corresponding to ResourceConfig preceding the one corresponding to AccessConfig. If using the default values, include them explicitly as conf/srm.conf and conf/access.conf files.

23.2.2.2. Main Server Configuration

The main server configuration section of the configuration file sets up the main server, which responds to any requests that are not handled by a virtual host defined within a <VirtualHost> container. Values here also provide defaults for any <VirtualHost> containers defined.

The directives used in this section have changed little between Apache HTTP Server 1.3 and version 2.0. If the main server configuration is heavily customized, it may be easier to modify the existing configuration file to suit Apache HTTP Server 2.0. Users with only lightly customized main server sections should migrate their changes into the default 2.0 configuration.

23.2.2.2.1. UserDir Mapping

The UserDir directive is used to enable URLs such as http://example.com/~bob/ to map to a subdirectory within the home directory of the user bob, such as /home/bob/public_html/. A side-effect of this feature allows a potential attacker to determine whether a given username is present on the system. For this reason, the default configuration for Apache HTTP Server 2.0 disables this directive.

To enable UserDir mapping, change the directive in httpd.conf from:

```
UserDir disable
```

to the following:

```
UserDir public_html
```

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_userdir.html#userdir

23.2.2.2.2. Logging

The following logging directives have been removed:

- AgentLog
- RefererLog
- RefererIgnore

However, agent and referrer logs are still available using the CustomLog and LogFormat directives.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_log_config.html#customlog
- http://httpd.apache.org/docs-2.0/mod/mod_log_config.html#logformat

23.2.2.2.3. Directory Indexing

The deprecated FancyIndexing directive has now been removed. The same functionality is available through the FancyIndexing option within the IndexOptions directive.

The VersionSort option to the IndexOptions directive causes files containing version numbers to be sorted in a more natural way. For example, httpd-2.0.6.tar appears before httpd-2.0.36.tar in a directory index page.

The defaults for the ReadmeName and HeaderName directives have changed from README and HEADER to README.html and HEADER.html.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_autoindex.html#indexoptions
- http://httpd.apache.org/docs-2.0/mod/mod_autoindex.html#readmename
- http://httpd.apache.org/docs-2.0/mod/mod_autoindex.html#headername

23.2.2.2.4. Content Negotiation

The CacheNegotiatedDocs directive now takes the argument on or off. Existing instances of CacheNegotiatedDocs should be replaced with CacheNegotiatedDocs on.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_negotiation.html#cachenegotiateddocs

23.2.2.2.5. Error Documents

To use a hard-coded message with the ErrorDocument directive, the message should be enclosed in a pair of double quotation marks ", rather than just preceded by a double quotation mark as required in Apache HTTP Server 1.3.

For example, the following is a sample Apache HTTP Server 1.3 directive:

```
ErrorDocument 404 "The document was not found
```

To migrate an ErrorDocument setting to Apache HTTP Server 2.0, use the following structure:

```
ErrorDocument 404 "The document was not found"
```

Note the trailing double quote in the previous ErrorDocument directive example.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/core.html#errordocument

23.2.2.3. Virtual Host Configuration

The contents of all <VirtualHost> containers should be migrated in the same way as the main server section as described in 23.2.2.2절. "Main Server Configuration" .

Important

Note that SSL/TLS virtual host configuration has been moved out of the main server configuration file and into /etc/httpd/conf.d/ssl.conf.

- http://httpd.apache.org/docs-2.0/vhosts/

23.2.2.4. Modules and Apache HTTP Server 2.0

In Apache HTTP Server 2.0, the module system has been changed to allow modules to be chained together or combined in new and interesting ways. Common Gateway Interface (CGI) scripts, for example, can generate server-parsed HTML documents which can then be processed by mod_include. This opens up a tremendous number of possibilities with regards to how modules can be combined to achieve a specific goal.

The way this works is that each request is served by exactly one handler module followed by zero or more filter modules.

Under Apache HTTP Server 1.3, for example, a Perl script would be handled in its entirety by the Perl module (mod_perl). Under Apache HTTP Server 2.0, the request is initially handled by the core module — which serves static files — and is then filtered by mod_perl.

Exactly how to use this, and all other new features of Apache HTTP Server 2.0, is beyond the scope of this document; however, the change has ramifications if the PATH_INFO directive is used for a document which is handled by a module that is now implemented as a filter, as each contains trailing path information after the true file name. The core module, which initially handles the request, does not by default understand PATH_INFO and returns 404 Not Found errors for requests that contain such information. As an alternative, use the AcceptPathInfo directive to coerce the core module into accepting requests with PATH_INFO.

The following is an example of this directive:

```
AcceptPathInfo on
```

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/core.html#acceptpathinfo
- http://httpd.apache.org/docs-2.0/handler.html
- http://httpd.apache.org/docs-2.0/filter.html

23.2.2.4.1. The suexec Module

In Apache HTTP Server 2.0, the mod_suexec module uses the SuexecUserGroup directive, rather than the User and Group directives, which is used for configuring virtual hosts. The User and Group directives can still be used in general, but are deprecated for configuring virtual hosts.

For example, the following is a sample Apache HTTP Server 1.3 directive:

```
<VirtualHost vhost.example.com:80>
  User someone
  Group somegroup
</VirtualHost>
```

To migrate this setting to Apache HTTP Server 2.0, use the following structure:

```
<VirtualHost vhost.example.com:80>
  SuexecUserGroup someone somegroup
</VirtualHost>
```

23.2.2.4.2. The mod_ssl Module

The configuration for mod_ssl has been moved from the httpd.conf file into the /etc/httpd/conf.d/ssl.conf file. For this file to be loaded, and for mod_ssl to work, the statement Include conf.d/*.conf must be in the httpd.conf file as described in 23.2.2.1.3절. "Dynamic Shared Object (DSO) Support" .

ServerName directives in SSL virtual hosts must explicitly specify the port number.

For example, the following is a sample Apache HTTP Server 1.3 directive:

```
<VirtualHost _default_:443>
  # General setup for the virtual host
  ServerName ssl.example.name
  ...
</VirtualHost>
```

To migrate this setting to Apache HTTP Server 2.0, use the following structure:

```
<VirtualHost _default_:443>
  # General setup for the virtual host
  ServerName ssl.host.name:443
  ...
</VirtualHost>
```

It is also important to note that both the SSLLog and SSLLogLevel directives have been removed. The mod_ssl module now obeys the ErrorLog and LogLevel directives. Refer to ErrorLog and LogLevel for more information about these directives.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_ssl.html
- http://httpd.apache.org/docs-2.0/vhosts/

23.2.2.4.3. The mod_proxy Module

Proxy access control statements are now placed inside a <Proxy> block rather than a <Directory proxy:>.

The caching functionality of the old mod_proxy has been split out into the following three modules:

- mod_cache
- mod_disk_cache
- mod_mem_cache

These generally use directives similar to the older versions of the mod_proxy module, but it is advisable to verify each directive before migrating any cache settings.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_proxy.html

23.2.2.4.4. The mod_include Module

The mod_include module is now implemented as a filter and is therefore enabled differently. Refer to 23.2.2.4절. "Modules and Apache HTTP Server 2.0" for more about filters.

For example, the following is a sample Apache HTTP Server 1.3 directive:

```
AddType text/html .shtml
AddHandler server-parsed .shtml
```

To migrate this setting to Apache HTTP Server 2.0, use the following structure:

```
AddType text/html .shtml
```

```
AddOutputFilter INCLUDES .shtml
```

Note that the Options +Includes directive is still required for the <Directory> container or in a .htaccess file.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_include.html

23.2.2.4.5. The mod_auth_dbm and mod_auth_db Modules

Apache HTTP Server 1.3 supported two authentication modules, mod_auth_db and mod_auth_dbm, which used Berkeley Databases and DBM databases respectively. These modules have been combined into a single module named mod_auth_dbm in Apache HTTP Server 2.0, which can access several different database formats. To migrate from mod_auth_db, configuration files should be adjusted by replacing AuthDBUserFile and AuthDBGroupFile with the mod_auth_dbm equivalents, AuthDBMUserFile and AuthDBMGroupFile. Also, the directive AuthDBMType DB must be added to indicate the type of database file in use.

The following example shows a sample mod_auth_db configuration for Apache HTTP Server 1.3:

```
<Location /private/>
  AuthType Basic
  AuthName "My Private Files"
  AuthDBUserFile /var/www/authdb
  require valid-user
</Location>
```

To migrate this setting to version 2.0 of Apache HTTP Server, use the following structure:

```
<Location /private/>
  AuthType Basic
  AuthName "My Private Files"
  AuthDBMUserFile /var/www/authdb
  AuthDBMType DB
  require valid-user
</Location>
```

Note that the AuthDBMUserFile directive can also be used in .htaccess files.

The dbmmanage Perl script, used to manipulate username and password databases, has been replaced by htdbm in Apache HTTP Server 2.0. The htdbm program offers equivalent functionality and, like mod_auth_dbm, can operate a variety of database formats; the -T option can be used on the command line to specify the format to use.

표 23.1. "Migrating from dbmmanage to htdbm" shows how to migrate from a DBM-format database to htdbm format using dbmmanage.

표 23.1. Migrating from dbmmanage to htdbm

Action	dbmmanage command (1.3)	Equivalent htdbm command (2.0)
Add user to database (using given password)	dbmmanage authdb add username password	htdbm -b -TDB authdb username password
Add user to database (prompts for password)	dbmmanage authdb adduser username	htdbm -TDB authdb username

Action	dbmmanage command (1.3)	Equivalent htdbm command (2.0)
Remove user from database	dbmmanage authdb delete username	htdbm -x -TDB authdb username
List users in database	dbmmanage authdb view	htdbm -l -TDB authdb
Verify a password	dbmmanage authdb check username	htdbm -v -TDB authdb username

The -m and -s options work with both dbmmanage and htdbm, enabling the use of the MD5 or SHA1 algorithms for hashing passwords, respectively.

When creating a new database with htdbm, the -c option must be used.

For more on this topic, refer to the following documentation on the Apache Software Foundation's website:

- http://httpd.apache.org/docs-2.0/mod/mod_auth_dbm.html

23.2.2.4.6. The mod_perl Module

The configuration for mod_perl has been moved from httpd.conf into the file /etc/httpd/conf.d/perl.conf. For this file to be loaded, and hence for mod_perl to work, the statement Include conf.d/*.conf must be included in httpd.conf as described in 23.2.2.1.3절. "Dynamic Shared Object (DSO) Support" .

Occurrences of Apache:: in httpd.conf must be replaced with ModPerl::. Additionally, the manner in which handlers are registered has been changed.

This is a sample Apache HTTP Server 1.3 mod_perl configuration:

```
<Directory /var/www/perl>
  SetHandler perl-script
  PerlHandler Apache::Registry
  Options +ExecCGI
</Directory>
```

This is the equivalent mod_perl for Apache HTTP Server 2.0:

```
<Directory /var/www/perl>
  SetHandler perl-script
  PerlResponseHandler ModPerl::Registry
  Options +ExecCGI
</Directory>
```

Most modules for mod_perl 1.x should work without modification with mod_perl 2.x. XS modules require recompilation and may require minor Makefile modifications.

23.2.2.4.7. The mod_python Module

Configuration for mod_python has moved from httpd.conf to the /etc/httpd/conf.d/python.conf file. For this file to be loaded, and hence for mod_python to work, the statement Include conf.d/*.conf must be in httpd.conf as described in 23.2.2.1.3절. "Dynamic Shared Object (DSO) Support" .

23.2.2.4.8. PHP

The configuration for PHP has been moved from httpd.conf into the file /etc/httpd/conf.d/php.conf. For this file to be loaded, the statement Include conf.d/*.conf must be in httpd.conf as described in 23.2.2.1.3절. "Dynamic Shared Object (DSO) Support" .

Note

Any PHP configuration directives used in Apache HTTP Server 1.3 are now fully compatible, when migrating to Apache HTTP Server 2.0 on Red Hat Enterprise Linux 5.

In PHP version 4.2.0 and later the default set of predefined variables which are available in the global scope has changed. Individual input and server variables are, by default, no longer placed directly into the global scope. This change may cause scripts to break. Revert to the old behavior by setting register_globals to On in the file /etc/php.ini.

For more on this topic, refer to the following URL for details concerning the global scope changes:

- http://www.php.net/release_4_1_0.php

23.2.2.4.9. The mod_authz_ldap Module

Red Hat Enterprise Linux ships with the mod_authz_ldap module for the Apache HTTP Server. This module uses the short form of the distinguished name for a subject and the issuer of the client SSL certificate to determine the distinguished name of the user within an LDAP directory. It is also capable of authorizing users based on attributes of that user's LDAP directory entry, determining access to assets based on the user and group privileges of the asset, and denying access for users with expired passwords. The mod_ssl module is required when using the mod_authz_ldap module.

Important

The mod_authz_ldap module does not authenticate a user to an LDAP directory using an encrypted password hash. This functionality is provided by the experimental mod_auth_ldap module. Refer to the mod_auth_ldap module documentation online at http://httpd.apache.org/docs-2.0/mod/mod_auth_ldap.html for details on the status of this module.

The /etc/httpd/conf.d/authz_ldap.conf file configures the mod_authz_ldap module.

Refer to /usr/share/doc/mod_authz_ldap-<version>/index.html (replacing <version> with the version number of the package) or http://authzldap.othello.ch/ for more information on configuring the mod_authz_ldap third party module.

23.3. Starting and Stopping httpd

After installing the httpd package, review the Apache HTTP Server's documentation available online at http://httpd.apache.org/docs/2.2/.

The httpd RPM installs the /etc/init.d/httpd script, which can be accessed using the /sbin/service command.

Starting httpd using the apachectl control script sets the environmental variables in /etc/sysconfig/httpd and starts httpd. You can also set the environment variables using the init script.

To start the server using the apachectl control script as root type:

```
apachectl start
```

You can also start httpd using /sbin/service httpd start. This starts httpd but does not set the environment variables. If you are using the default Listen directive in httpd.conf, which is port 80, you will need to have root privileges to start the apache server.

To stop the server, as root type:

```
apachectl stop
```

You can also stop httpd using /sbin/service httpd stop. The restart option is a shorthand way of stopping and then starting the Apache HTTP Server.

You can restart the server as root by typing:

```
apachectl restart
```

or:

```
service httpd restart
```

Apache will display a message on the console or in the ErrorLog if it encounters an error while starting.

By default, the httpd service does not start automatically at boot time. If you would wish to have Apache startup at boot time, you will need to add a call to apachectl in your startup files within the rc.N directory. A typical file used is rc.local. As this starts Apache as root, it is recommended to properly configure your security and authentication before adding this call.

You can also configure the httpd service to start up at boot time, using an initscript utility, such as /sbin/chkconfig, /usr/sbin/ntsysv, or the Services Configuration Tool program.

You can also display the status of your httpd server by typing:

```
apachectl status
```

The status module mod_status however needs to be enabled in your httpd.conf configuration file for this to work. For more details on mod_status can be found on http://httpd.apache.org/docs/2.2/mod/mod_status.html.

Note

If running the Apache HTTP Server as a secure server, the secure server's password is required after the machine boots when using an encrypted private SSL key.

You can find more information on http://httpd.apache.org/docs/2.2/ssl

23.4. Apache HTTP Server Configuration

The HTTP Configuration Tool allows you to configure the /etc/httpd/conf/httpd.conf configuration file for the Apache HTTP Server. It does not use the old srm.conf or access.conf configuration files; leave them empty. Through the graphical interface, you can configure directives such as virtual hosts, logging attributes, and maximum number of connections. To start the HTTD Configuration Tool, click on System > Administration > Server Settings > HTTP.

Only modules provided with Red Hat Enterprise Linux can be configured with the HTTP Configuration Tool. If additional modules are installed, they can not be configured using this tool.

Caution

Do not edit the /etc/httpd/conf/httpd.conf configuration file by hand if you wish to use this tool. The HTTP Configuration Tool generates this file after you save your changes and exit the program. If you want to add additional modules or configuration options that are not available in HTTP Configuration Tool, you cannot use this tool.

The general steps for configuring the Apache HTTP Server using the HTTP Configuration Tool are as follows:

1. Configure the basic settings under the Main tab.
2. Click on the Virtual Hosts tab and configure the default settings.
3. Under the Virtual Hosts tab, configure the Default Virtual Host.
4. To serve more than one URL or virtual host, add any additional virtual hosts.
5. Configure the server settings under the Server tab.
6. Configure the connections settings under the Performance Tuning tab.
7. Copy all necessary files to the DocumentRoot and cgi-bin directories.
8. Exit the application and select to save your settings.

23.4.1. Basic Settings

Use the Main tab to configure the basic server settings.

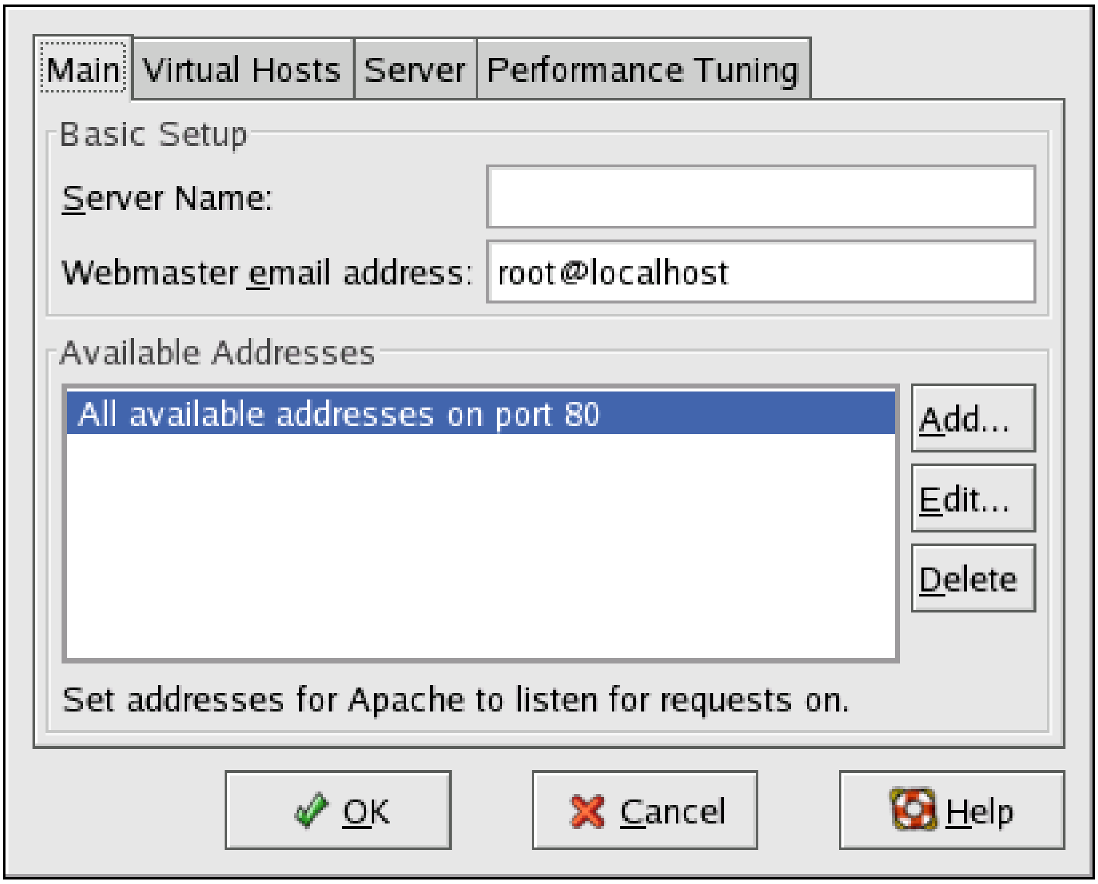

Server Name:

Webmaster email address: root@localhost

All available addresses on port 80

Set addresses for Apache to listen for requests on.

그림 23.1. Basic Settings

Enter a fully qualified domain name that you have the right to use in the Server Name text area. This option corresponds to the ServerName[2] directive in httpd.conf. The ServerName directive sets the hostname of the Web server. It is used when creating redirection URLs. If you do not define a server name, the Web server attempts to resolve it from the IP address of the system. The server name does not have to be the domain name resolved from the IP address of the server. For example, you might set the server name to www.example.com while the server's real DNS name is foo.example.com.

Enter the email address of the person who maintains the Web server in the Webmaster email address text area. This option corresponds to the ServerAdmin[3] directive in httpd.conf. If you configure the server's error pages to contain an email address, this email address is used so that users can report a problem to the server's administrator. The default value is root@localhost.

Use the Available Addresses area to define the ports on which the server accepts incoming requests. This option corresponds to the Listen[4] directive in httpd.conf. By default, Red Hat configures the Apache HTTP Server to listen to port 80 for non-secure Web communications.

Click the Add button to define additional ports on which to accept requests. A window as shown in 그림 23.2. "Available Addresses" appears. Either choose the Listen to all addresses option to

[2] http://httpd.apache.org/docs/2.2/mod/core.html#servername

[3] http://httpd.apache.org/docs/2.2/mod/core.html#serveradmin

[4] http://httpd.apache.org/docs/2.2/mod/mpm_common.html#listen

listen to all IP addresses on the defined port or specify a particular IP address over which the server accepts connections in the Address field. Only specify one IP address per port number. To specify more than one IP address with the same port number, create an entry for each IP address. If at all possible, use an IP address instead of a domain name to prevent a DNS lookup failure. Refer to http://httpd.apache.org/docs/2.2/dns-caveats.html for more information about Issues Regarding DNS and Apache.

Entering an asterisk (*) in the Address field is the same as choosing Listen to all addresses. Clicking the Edit button in the Available Addresses frame shows the same window as the Add button except with the fields populated for the selected entry. To delete an entry, select it and click the Delete button.

Tip

If you set the server to listen to a port under 1024, you must be root to start it. For port 1024 and above, httpd can be started as a regular user.

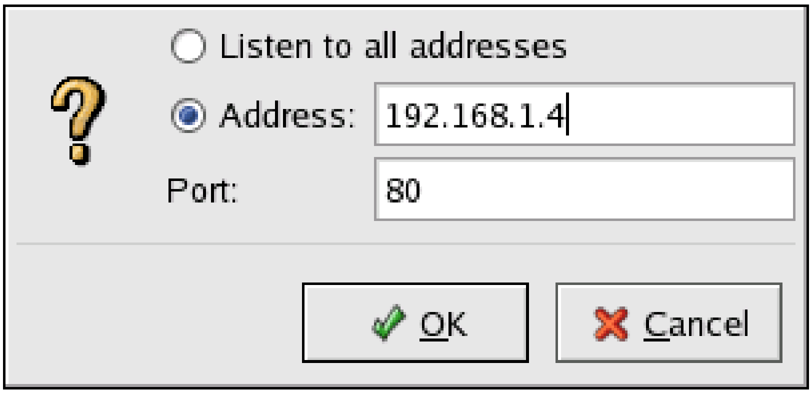

○ Listen to all addresses

● Address: 192.168.1.4

Port: 80

그림 23.2. Available Addresses

23.4.2. Default Settings

After defining the Server Name, Webmaster email address, and Available Addresses, click the Virtual Hosts tab. The figure below illustrates the Virtual Hosts tab.

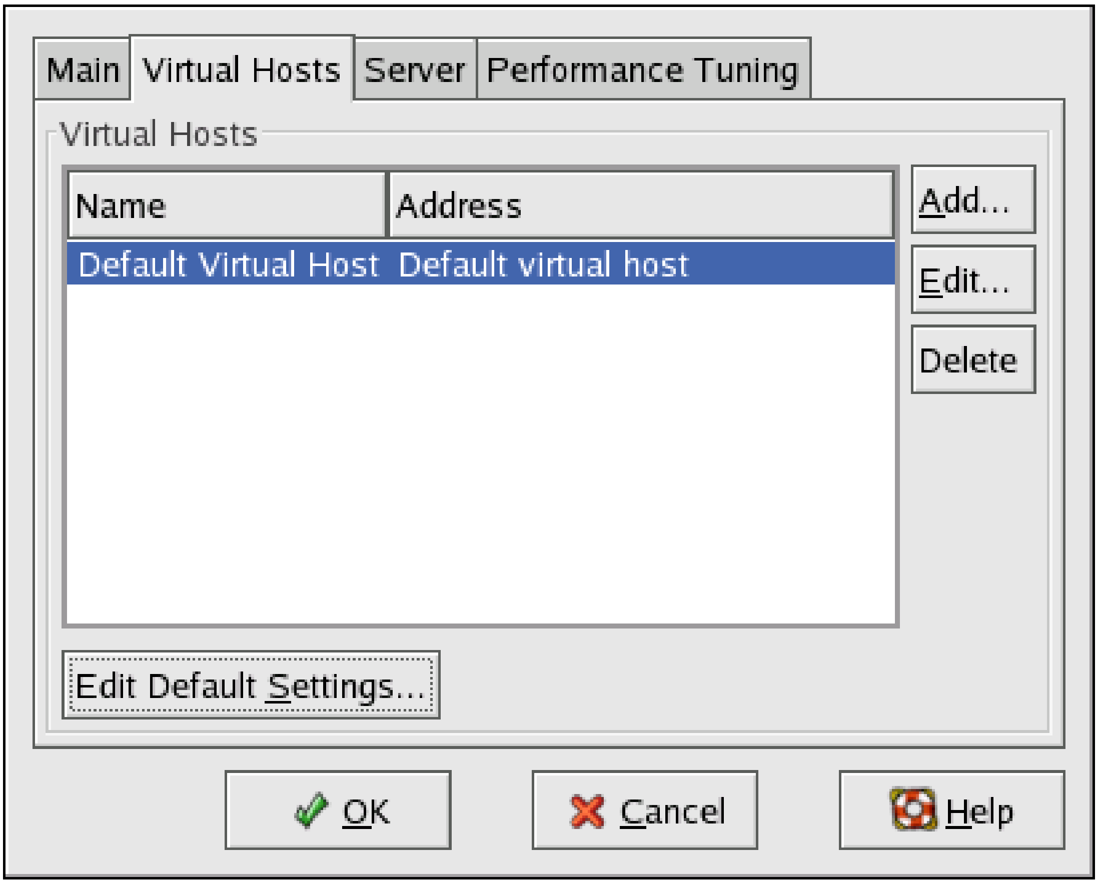

Name	Address
Default Virtual Host	Default virtual host

그림 23.3. Virtual Hosts Tab

Clicking on Edit will display the Virtual Host Properties window from which you can set your preferred settings. To add new settings, click on the Add button which will also display the Virtual Host Properties window. Clicking on the Edit Default Settings button, displays the Virtual Host Properties window without the General Options tab.

In the General Options tab, you can change the hostname, the document root directory and also set the webmaster's email address. In the Host information, you can set the Virtual Host's IP Address and Host Name. The figure below illustrates the General Options tab.

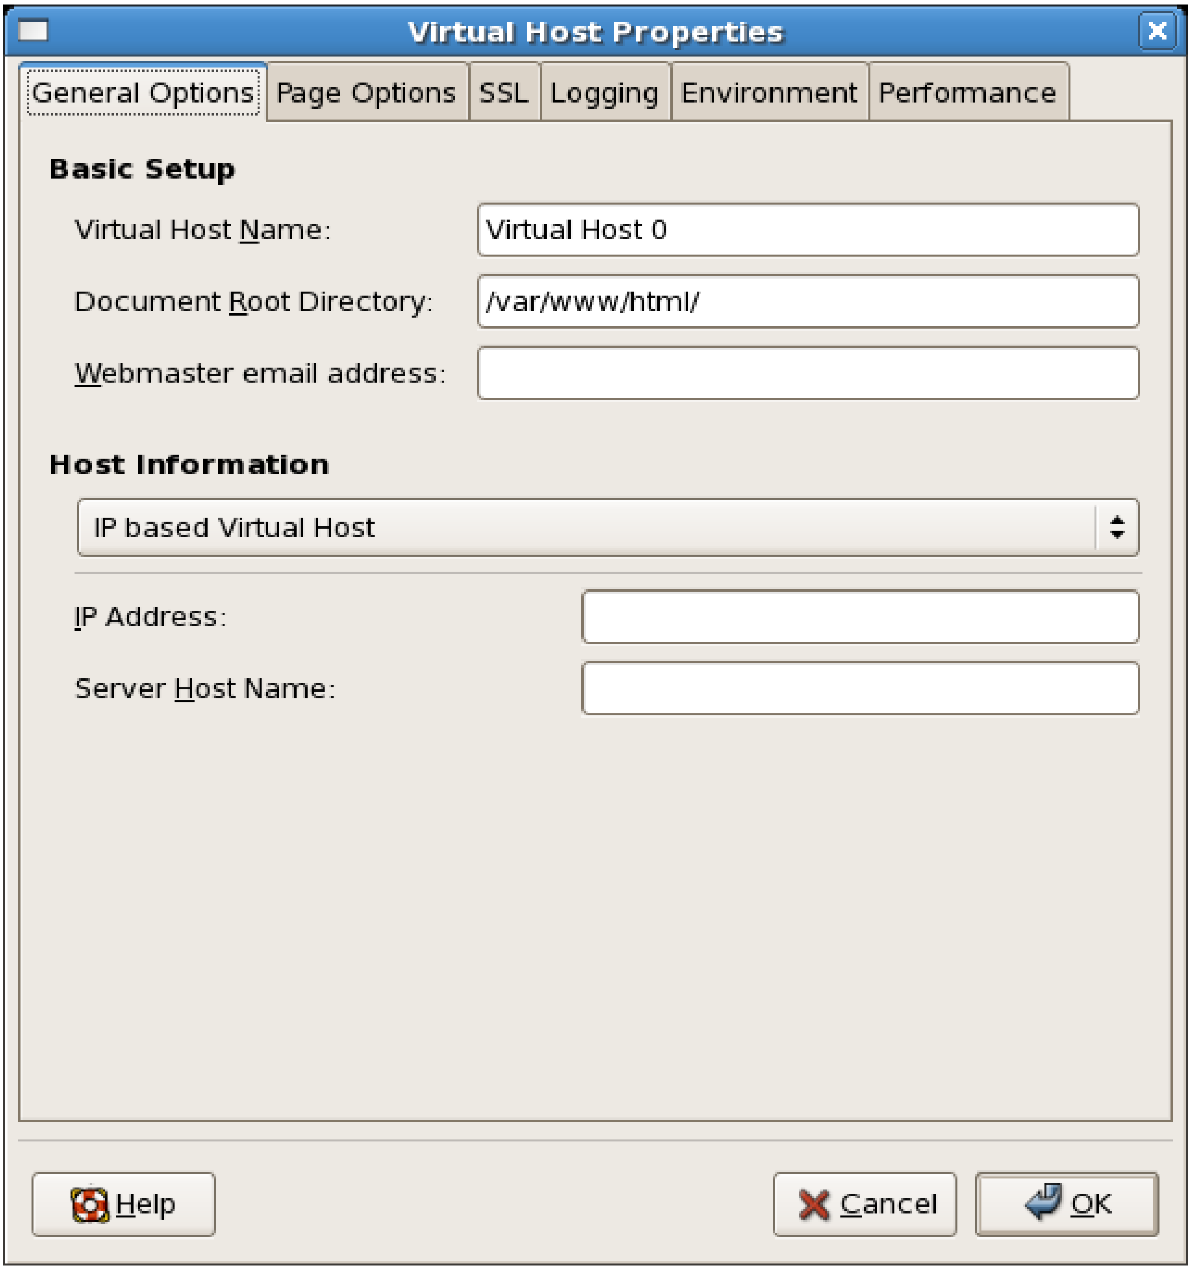

Basic Setup

Virtual Host Name: Virtual Host 0

Document Root Directory: /var/www/html/

Webmaster email address:

Host Information

IP based Virtual Host

IP Address:

Server Host Name:

그림 23.4. General Options

If you add a virtual host, the settings you configure for the virtual host take precedence for that virtual host. For a directive not defined within the virtual host settings, the default value is used.

23.4.2.1. Site Configuration

The figure below illustrates the Page Optionstab from which you can configure the Directory Page Search List and Error Pages. If you are unsure of these settings, do not modify them.

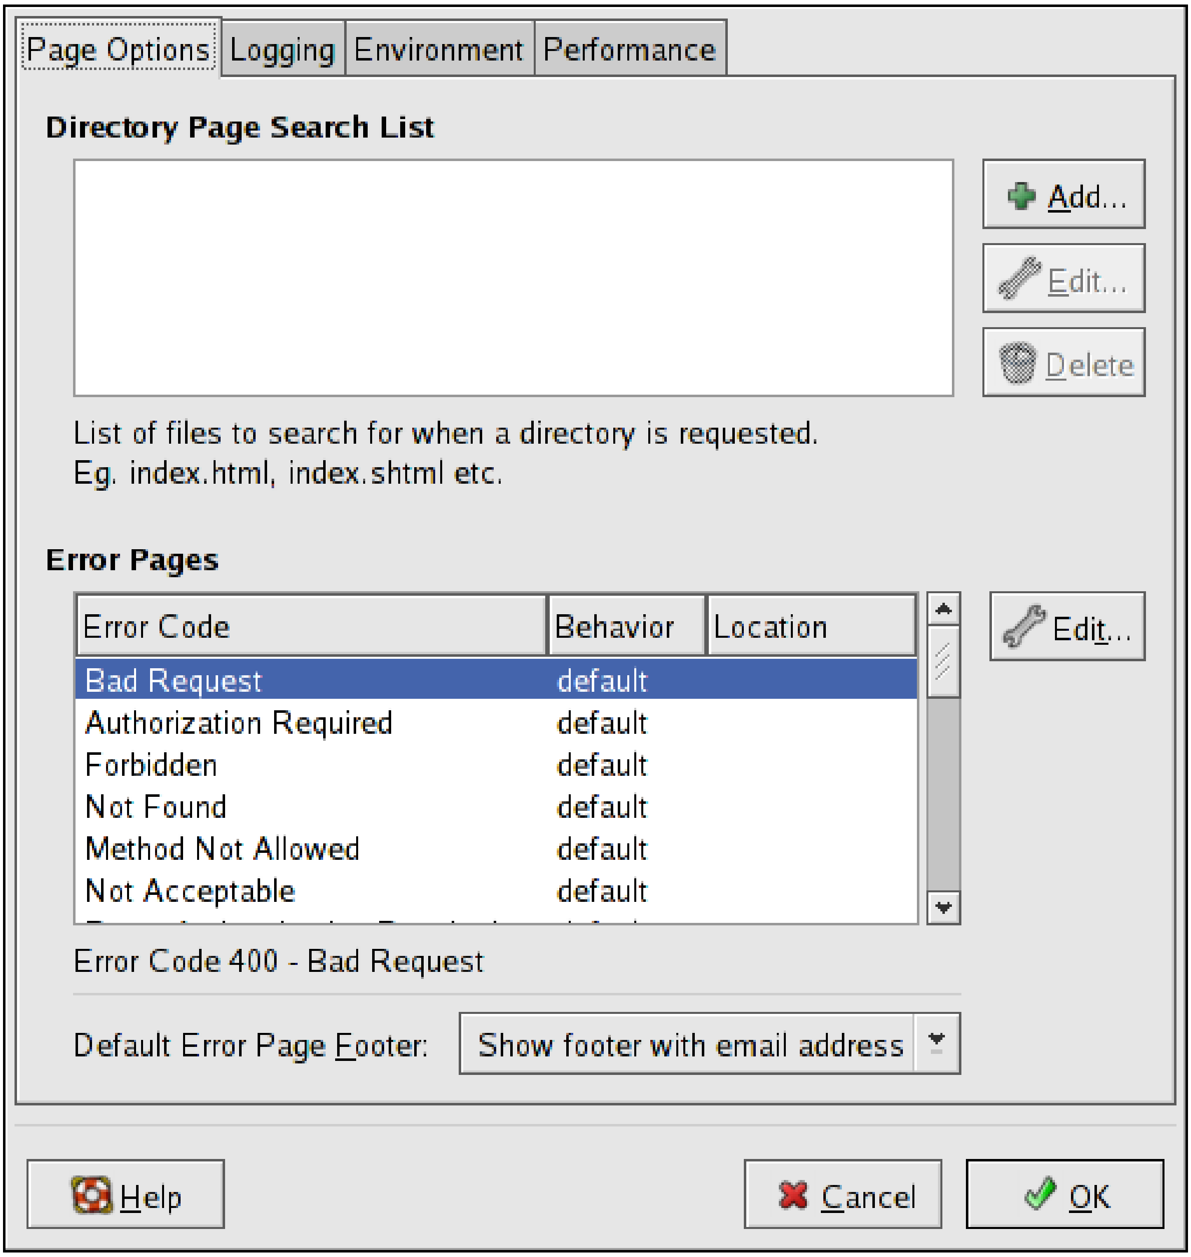

Directory Page Search List

List of files to search for when a directory is requested.
Eg. index.html, index.shtml etc.

Error Pages

Error Code	Behavior	Location
Bad Request	default	
Authorization Required	default	
Forbidden	default	
Not Found	default	
Method Not Allowed	default	
Not Acceptable	default	

Error Code 400 - Bad Request

Default Error Page Footer: Show footer with email address

그림 23.5. Site Configuration

The entries listed in the Directory Page Search List define the DirectoryIndex[5] directive. The DirectoryIndex is the default page served by the server when a user requests an index of a directory by specifying a forward slash (/) at the end of the directory name.

For example, when a user requests the page http://www.example.com/this_directory/, they are going to get either the DirectoryIndex page, if it exists, or a server-generated directory list. The server tries to find one of the files listed in the DirectoryIndex directive and returns the first one it finds. If it does not find any of these files and if Options Indexes is set for that directory, the server generates and returns a list, in HTML format, of the subdirectories and files in the directory.

[5] http://httpd.apache.org/docs/2.2/mod/mod_dir.html#directoryindex

Use the Error Code section to configure Apache HTTP Server to redirect the client to a local or external URL in the event of a problem or error. This option corresponds to the ErrorDocument[6] directive. If a problem or error occurs when a client tries to connect to the Apache HTTP Server, the default action is to display the short error message shown in the Error Code column. To override this default configuration, select the error code and click the Edit button. Choose Default to display the default short error message. Choose URL to redirect the client to an external URL and enter a complete URL, including the http://, in the Location field. Choose File to redirect the client to an internal URL and enter a file location under the document root for the Web server. The location must begin the a slash (/) and be relative to the Document Root.

For example, to redirect a 404 Not Found error code to a webpage that you created in a file called 404.html, copy 404.html to DocumentRoot/../error/404.html. In this case, DocumentRoot is the Document Root directory that you have defined (the default is /var/www/html/). If the Document Root is left as the default location, the file should be copied to /var/www/error/404.html. Then, choose File as the Behavior for 404 - Not Found error code and enter /error/404.html as the Location.

From the Default Error Page Footer menu, you can choose one of the following options:

- Show footer with email address — Display the default footer at the bottom of all error pages along with the email address of the website maintainer specified by the ServerAdmin[7] directive.

- Show footer — Display just the default footer at the bottom of error pages.

- No footer — Do not display a footer at the bottom of error pages.

23.4.2.2. SSL Support

The mod_ssl enables encryption of the HTTP protocol over SSL. SSL (Secure Sockets Layer) protocol is used for communication and encryption over TCP/IP networks. The SSL tab enables you to configure SSL for your server. To configure SSL you need to provide the path to your:

- Certificate file - equivalent to using the SSLCertificateFile directive which points the path to the PEM (Privacy Enhanced Mail)-encoded server certificate file.

- Key file - equivalent to using the SSLCertificateKeyFile directive which points the path to the PEM-encoded server private key file.

- Certificate chain file - equivalent to using the SSLCertificateChainFile directive which points the path to the certificate file containing all the server's chain of certificates.

- Certificate authority file - is an encrypted file used to confirm the authenticity or identity of parties communicating with the server.

You can find out more about configuration directives for SSL on http://httpd.apache.org/docs/2.2/mod/directives.html#S[8]. You also need to determine which SSL options to enable. These are equivalent to using the SSLOptions with the following options:

- FakeBasicAuth - enables standard authentication methods used by Apache. This means that the Client X509 certificate's Subject Distinguished Name (DN) is translated into a basic HTTP username.

[6] http://httpd.apache.org/docs/2.2/mod/core.html#errordocument
[7] http://httpd.apache.org/docs/2.2/mod/core.html#serveradmin
[8] http://httpd.apache.org/docs/2.2/mod/directives.html#S

- ExportCertData - creates CGI environment variables in SSL_SERVER_CERT, SSL_CLIENT_CERT and SSL_CLIENT_CERT_CHAIN_n where n is a number 0,1,2,3,4... These files are used for more certificate checks by CGI scripts.

- CompatEnvVars - enables backward compatibility for Apache SSL by adding CGI environment variables.

- StrictRequire - enables strict access which forces denial of access whenever the SSLRequireSSL and SSLRequire directives indicate access is forbidden.

- OptRenegotiate - enables avoidance of unnecessary handshakes by mod_ssl which also performs safe parameter checks. It is recommended to enable OptRenegotiate on a per directory basis.

More information on the above SSL options can be found on http://httpd.apache.org/docs/2.2/mod/mod_ssl.html#ssloptions[9]. The figure below illustrates the SSL tab and the options discussed above.

[9] http://httpd.apache.org/docs/2.2/mod/mod_ssl.html#ssloptions

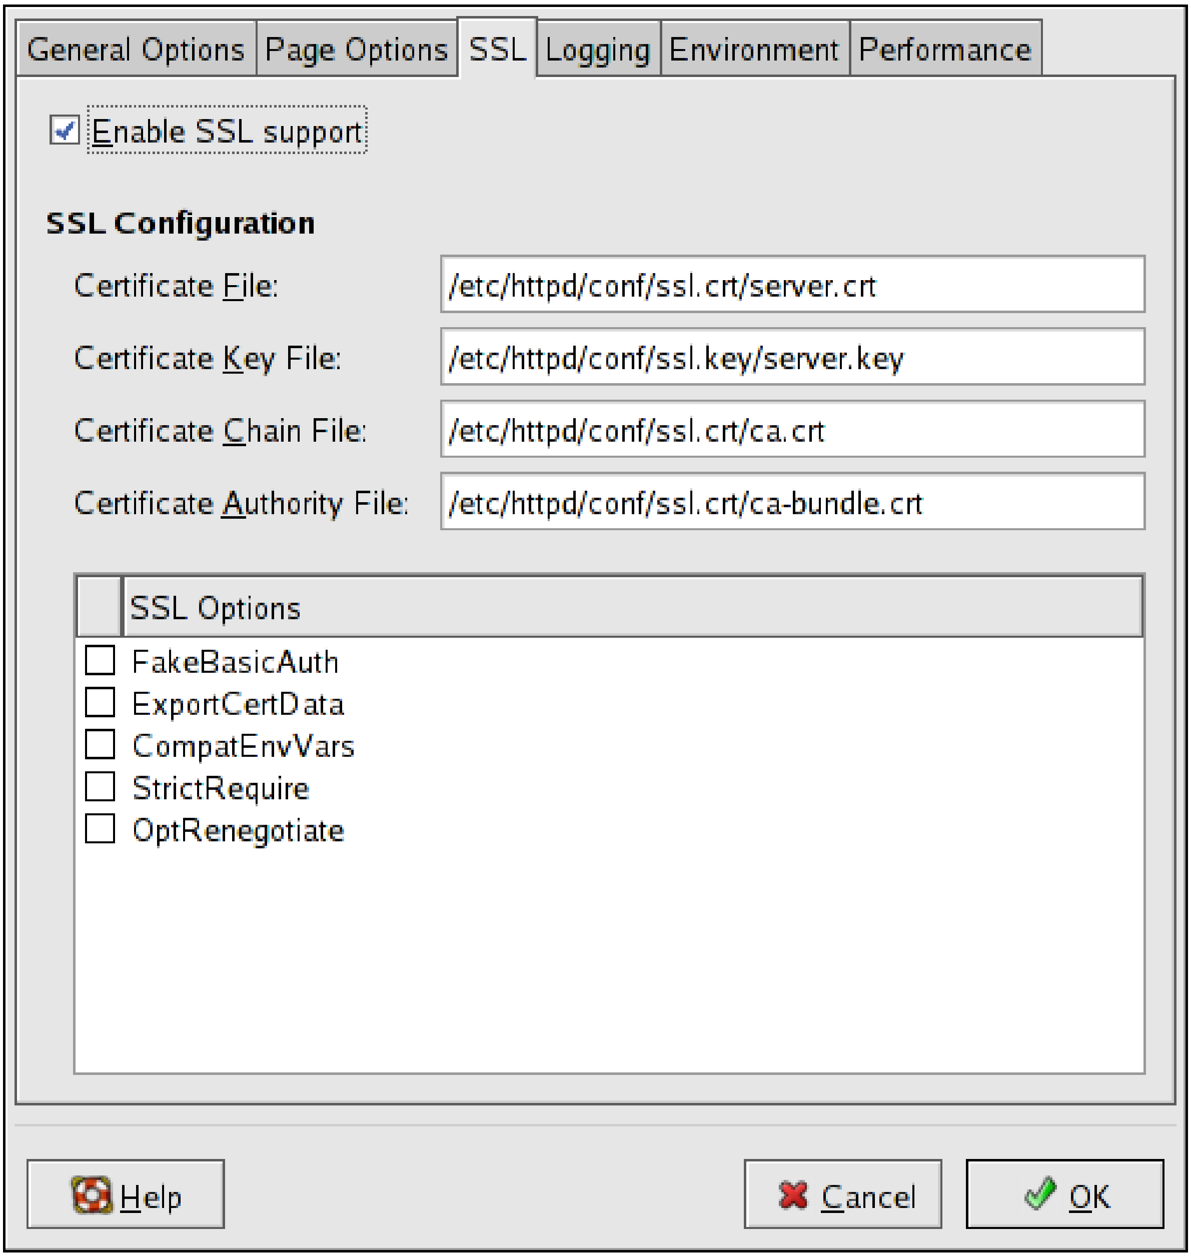

☑ Enable SSL support

SSL Configuration

Certificate File: /etc/httpd/conf/ssl.crt/server.crt

Certificate Key File: /etc/httpd/conf/ssl.key/server.key

Certificate Chain File: /etc/httpd/conf/ssl.crt/ca.crt

Certificate Authority File: /etc/httpd/conf/ssl.crt/ca-bundle.crt

SSL Options

- □ FakeBasicAuth
- □ ExportCertData
- □ CompatEnvVars
- □ StrictRequire
- □ OptRenegotiate

그림 23.6. SSL

23.4.2.3. Logging

Use the Logging tab to configure options for specific transfer and error logs.

By default, the server writes the transfer log to the /var/log/httpd/access_log file and the error log to the /var/log/httpd/error_log file.

The transfer log contains a list of all attempts to access the Web server. It records the IP address of the client that is attempting to connect, the date and time of the attempt, and the file on the Web server that it is trying to retrieve. Enter the name of the path and file in which to store this information. If the path and file name do not start with a slash (/), the path is relative to the server root directory as configured. This option corresponds to the TransferLog[10] directive.

[10] http://httpd.apache.org/docs/2.2/mod/mod_log_config.html#transferlog

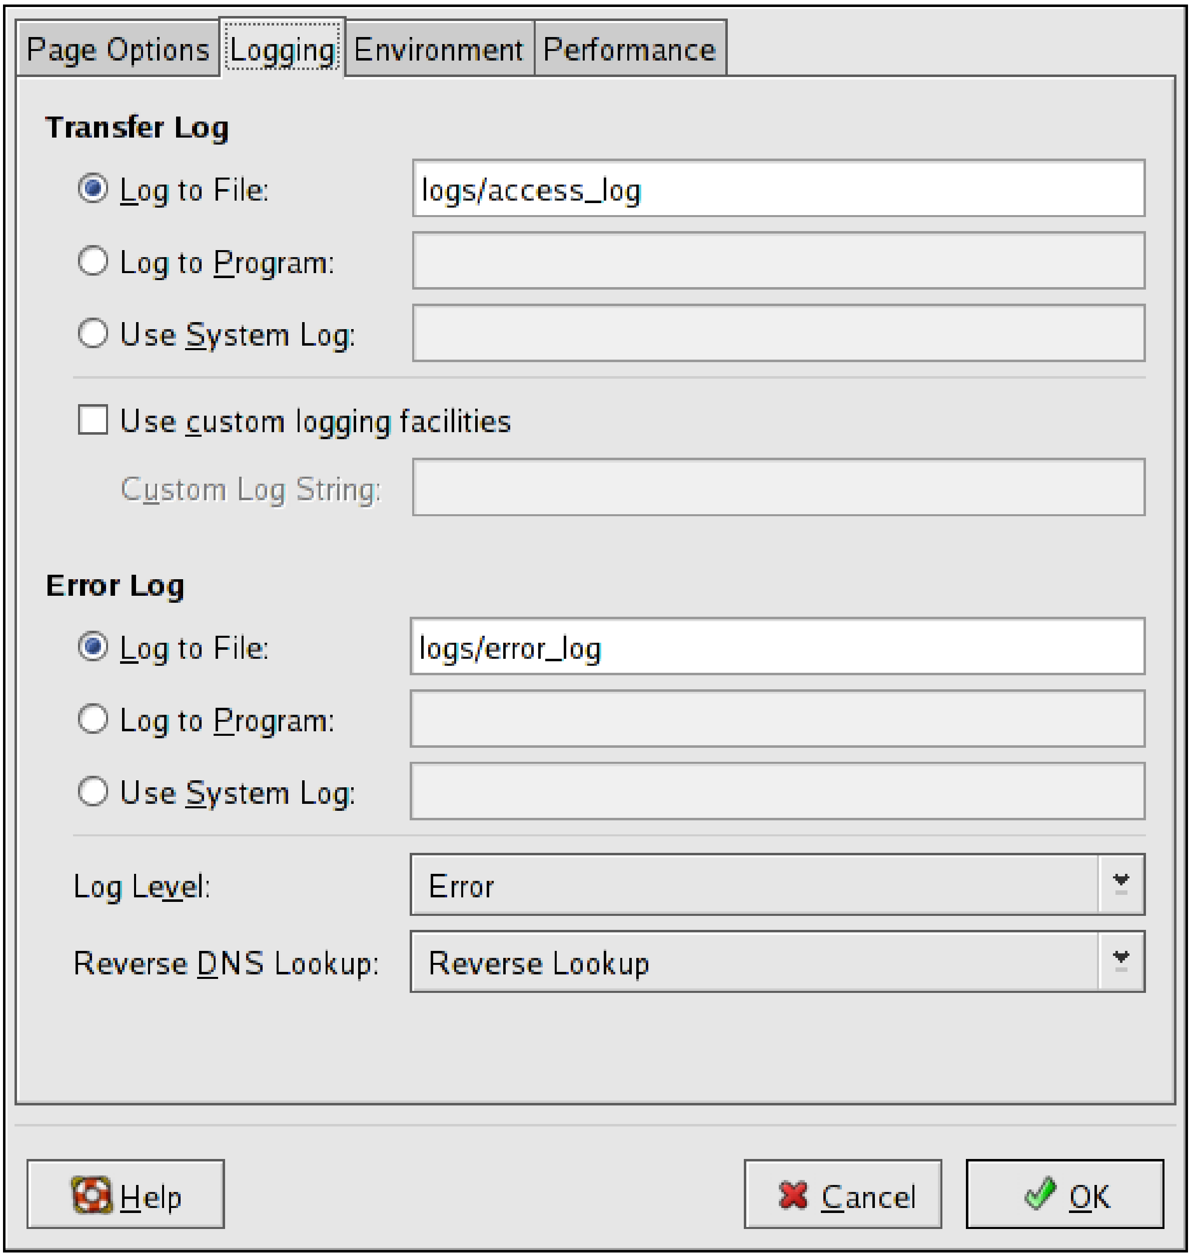

Transfer Log

● Log to File: logs/access_log

○ Log to Program:

○ Use System Log:

□ Use custom logging facilities

Custom Log String:

Error Log

● Log to File: logs/error_log

○ Log to Program:

○ Use System Log:

Log Level: Error

Reverse DNS Lookup: Reverse Lookup

그림 23.7. Logging

You can configure a custom log format by checking Use custom logging facilities and entering a custom log string in the Custom Log String field. This configures the LogFormat[11] directive. Refer to http://httpd.apache.org/docs/2.2/mod/mod_log_config.html#logformat[12] for details on the format of this directive.

The error log contains a list of any server errors that occur. Enter the name of the path and file in which to store this information. If the path and file name do not start with a slash (/), the path is relative to the server root directory as configured. This option corresponds to the ErrorLog[13] directive.

[11] http://httpd.apache.org/docs/2.2/mod/mod_log_config.html#logformat

[12] http://httpd.apache.org/docs-2.0/mod/mod_log_config.html#formats

[13] http://httpd.apache.org/docs/2.2/mod/core.html#errorlog

Use the Log Level menu to set the verbosity of the error messages in the error logs. It can be set (from least verbose to most verbose) to emerg, alert, crit, error, warn, notice, info or debug. This option corresponds to the LogLevel[14] directive.

The value chosen with the Reverse DNS Lookup menu defines the HostnameLookups[15] directive. Choosing No Reverse Lookup sets the value to off. Choosing Reverse Lookup sets the value to on. Choosing Double Reverse Lookup sets the value to double.

If you choose Reverse Lookup, your server automatically resolves the IP address for each connection which requests a document from your Web server. Resolving the IP address means that your server makes one or more connections to the DNS in order to find out the hostname that corresponds to a particular IP address.

If you choose Double Reverse Lookup, your server performs a double-reverse DNS. In other words, after a reverse lookup is performed, a forward lookup is performed on the result. At least one of the IP addresses in the forward lookup must match the address from the first reverse lookup.

Generally, you should leave this option set to No Reverse Lookup, because the DNS requests add a load to your server and may slow it down. If your server is busy, the effects of trying to perform these reverse lookups or double reverse lookups may be quite noticeable.

Reverse lookups and double reverse lookups are also an issue for the Internet as a whole. Each individual connection made to look up each hostname adds up. Therefore, for your own Web server's benefit, as well as for the Internet's benefit, you should leave this option set to No Reverse Lookup.

23.4.2.4. Environment Variables

Use the Environment tab to configure options for specific variables to set, pass, or unset for CGI scripts.

Sometimes it is necessary to modify environment variables for CGI scripts or server-side include (SSI) pages. The Apache HTTP Server can use the mod_env module to configure the environment variables which are passed to CGI scripts and SSI pages. Use the Environment Variables page to configure the directives for this module.

Use the Set for CGI Scripts section to set an environment variable that is passed to CGI scripts and SSI pages. For example, to set the environment variable MAXNUM to 50, click the Add button inside the Set for CGI Script section, as shown in 그림 23.8. “Environment Variables” , and type MAXNUM in the Environment Variable text field and 50 in the Value to set text field. Click OK to add it to the list. The Set for CGI Scripts section configures the SetEnv[16] directive.

Use the Pass to CGI Scripts section to pass the value of an environment variable when the server is first started to CGI scripts. To see this environment variable, type the command env at a shell prompt. Click the Add button inside the Pass to CGI Scripts section and enter the name of the environment variable in the resulting dialog box. Click OK to add it to the list. The Pass to CGI Scripts section configures the PassEnv [17] directive.

[14] http://httpd.apache.org/docs/2.2/mod/core.html#loglevel
[15] http://httpd.apache.org/docs/2.2/mod/core.html#hostnamelookups
[16] http://httpd.apache.org/docs/2.2/mod/mod_env.html#setenv
[17] http://httpd.apache.org/docs/2.2/mod/mod_env.html#passenv

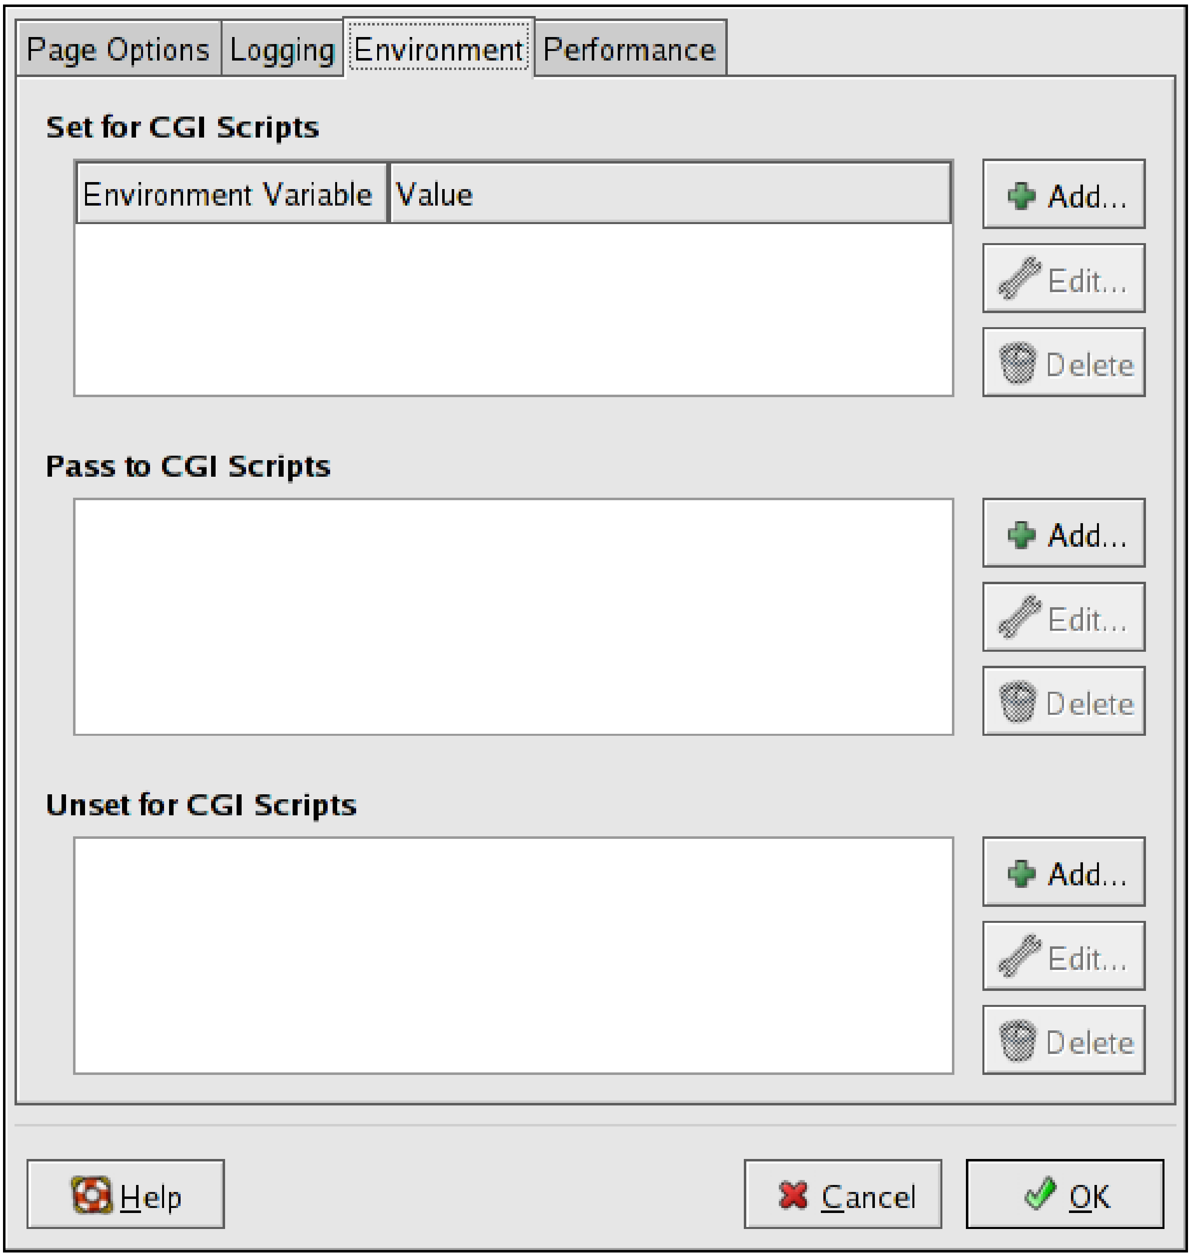

그림 23.8. Environment Variables

To remove an environment variable so that the value is not passed to CGI scripts and SSI pages, use the Unset for CGI Scripts section. Click Add in the Unset for CGI Scripts section, and enter the name of the environment variable to unset. Click OK to add it to the list. This corresponds to the UnsetEnv[18] directive.

To edit any of these environment values, select it from the list and click the corresponding Edit button. To delete any entry from the list, select it and click the corresponding Delete button.

To learn more about environment variables in the Apache HTTP Server, refer to the following: http://httpd.apache.org/docs/2.2/env.html

[18] http://httpd.apache.org/docs/2.2/mod/mod_env.html#unsetenv

23.4.2.5. Directories

Use the Directories page in the Performance tab to configure options for specific directories. This corresponds to the <Directory>[19] directive.

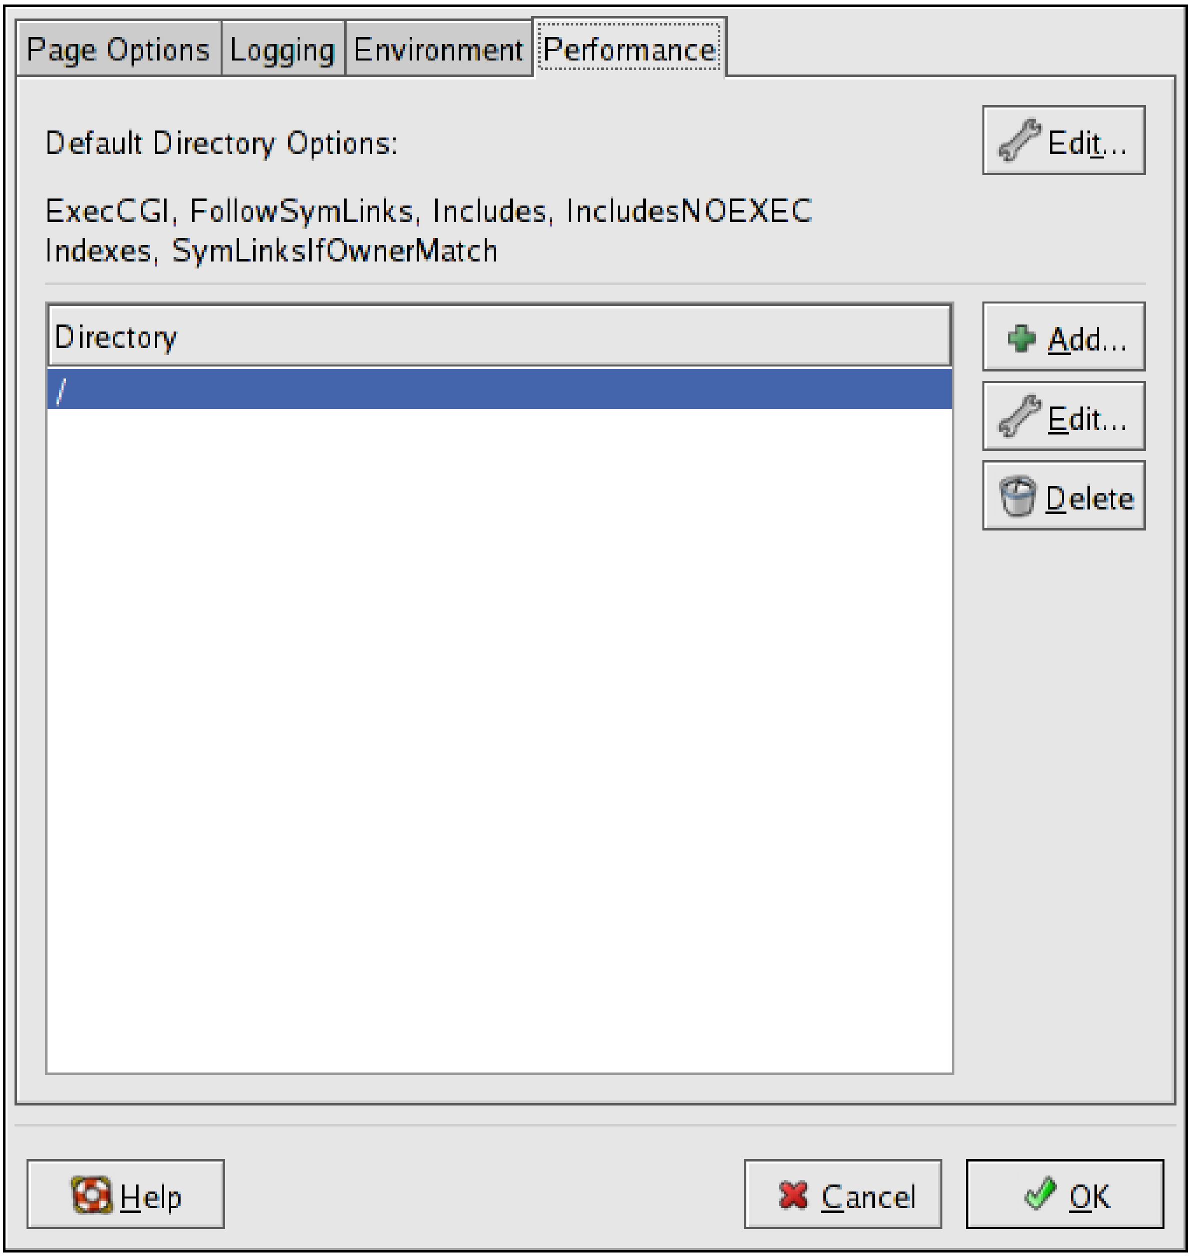

Default Directory Options:

ExecCGI, FollowSymLinks, Includes, IncludesNOEXEC
Indexes, SymLinksIfOwnerMatch

Directory
/

그림 23.9. Directories

Click the Edit button in the top right-hand corner to configure the Default Directory Options for all directories that are not specified in the Directory list below it. The options that you choose are listed as the Options[20] directive within the <Directory>[21] directive. You can configure the following options:

- ExecCGI — Allow execution of CGI scripts. CGI scripts are not executed if this option is not chosen.

[19] http://httpd.apache.org/docs/2.2/mod/core.html#directory
[20] http://httpd.apache.org/docs/2.2/mod/core.html#options
[21] http://httpd.apache.org/docs/2.2/mod/core.html#directory

- FollowSymLinks — Allow symbolic links to be followed.
- Includes — Allow server-side includes.
- IncludesNOEXEC — Allow server-side includes, but disable the #exec and #include commands in CGI scripts.
- Indexes — Display a formatted list of the directory's contents, if no DirectoryIndex (such as index.html) exists in the requested directory.
- Multiview — Support content-negotiated multiviews; this option is disabled by default.
- SymLinksIfOwnerMatch — Only follow symbolic links if the target file or directory has the same owner as the link.

To specify options for specific directories, click the Add button beside the Directory list box. A window as shown in 그림 23.10. "Directory Settings" appears. Enter the directory to configure in the Directory text field at the bottom of the window. Select the options in the right-hand list and configure the Order[22] directive with the left-hand side options. The Order directive controls the order in which allow and deny directives are evaluated. In the Allow hosts from and Deny hosts from text field, you can specify one of the following:

- Allow all hosts — Type all to allow access to all hosts.
- Partial domain name — Allow all hosts whose names match or end with the specified string.
- Full IP address — Allow access to a specific IP address.
- A subnet — Such as 192.168.1.0/255.255.255.0
- A network CIDR specification — such as 10.3.0.0/16

[22] http://httpd.apache.org/docs-2.0/mod/mod_access.html#order

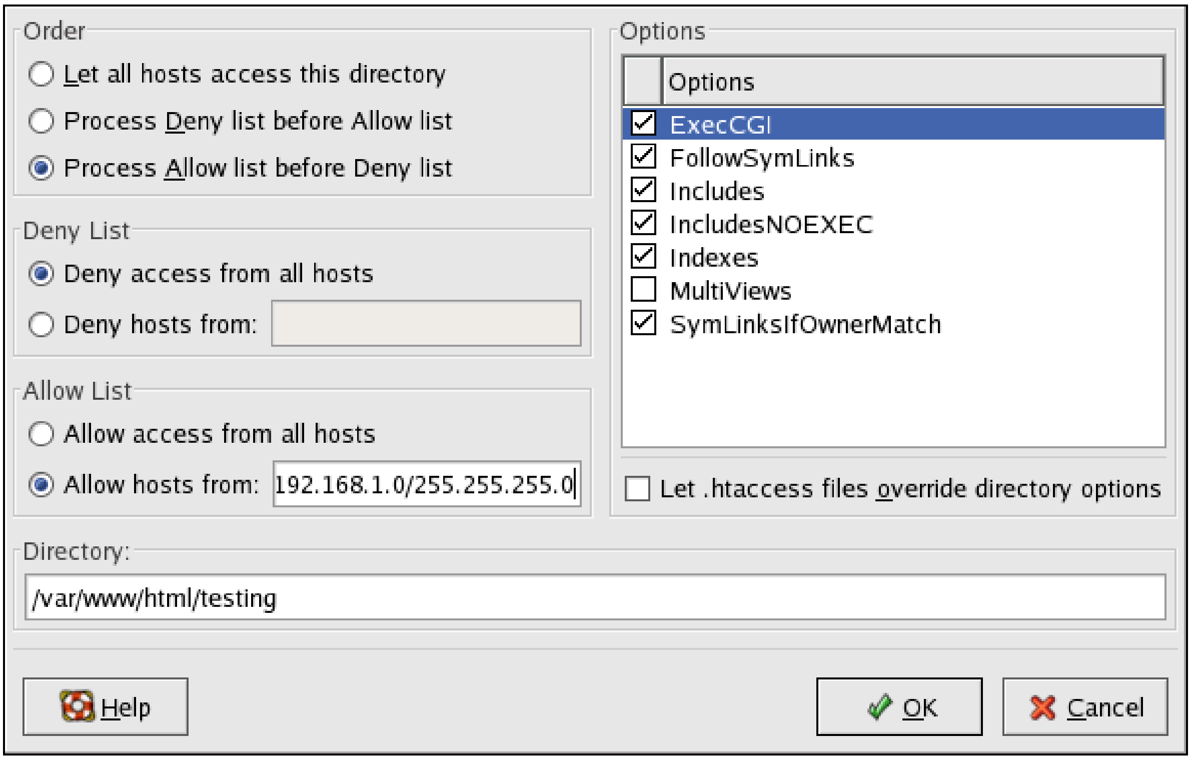

Order

○ Let all hosts access this directory

○ Process Deny list before Allow list

● Process Allow list before Deny list

Deny List

● Deny access from all hosts

○ Deny hosts from:

Allow List

○ Allow access from all hosts

● Allow hosts from: 192.168.1.0/255.255.255.0

Options

☑ ExecCGI

☑ FollowSymLinks

☑ Includes

☑ IncludesNOEXEC

☑ Indexes

□ MultiViews

☑ SymLinksIfOwnerMatch

□ Let .htaccess files override directory options

Directory: /var/www/html/testing

그림 23.10. Directory Settings

If you check the Let .htaccess files override directory options, the configuration directives in the .htaccess file take precedence.

23.5. Configuration Directives in httpd.conf

The Apache HTTP Server configuration file is /etc/httpd/conf/httpd.conf. The httpd.conf file is well-commented and mostly self-explanatory. The default configuration works for most situations; however, it is a good idea to become familiar some of the more important configuration options.

Warning

With the release of Apache HTTP Server 2.2, many configuration options have changed. If migrating from version 1.3 to 2.2, please firstly read 23.2.2절. "Migrating Apache HTTP Server 1.3 Configuration Files to 2.0" .

23.5.1. General Configuration Tips

If configuring the Apache HTTP Server, edit /etc/httpd/conf/httpd.conf and then either reload, restart, or stop and start the httpd process as outlined in 23.3절. "Starting and Stopping httpd" .

Before editing httpd.conf, make a copy the original file. Creating a backup makes it easier to recover from mistakes made while editing the configuration file.

If a mistake is made and the Web server does not work correctly, first review recently edited passages in httpd.conf to verify there are no typos.

Next look in the Web server's error log, /var/log/httpd/error_log. The error log may not be easy to interpret, depending on your level of expertise. However, the last entries in the error log should provide useful information.

The following subsections contain a list of short descriptions for many of the directives included in httpd.conf. These descriptions are not exhaustive. For more information, refer to the Apache documentation online at http://httpd.apache.org/docs/2.2/.

For more information about mod_ssl directives, refer to the documentation online at http://httpd.apache.org/docs/2.2/mod/mod_ssl.html.

AccessFileName

AccessFileName names the file which the server should use for access control information in each directory. The default is .htaccess.

Immediately after the AccessFileName directive, a set of Files tags apply access control to any file beginning with a .ht. These directives deny Web access to any .htaccess files (or other files which begin with .ht) for security reasons.

Action

Action specifies a MIME content type and CGI script pair, so that when a file of that media type is requested, a particular CGI script is executed.

AddDescription

When using FancyIndexing as an IndexOptions parameter, the AddDescription directive can be used to display user-specified descriptions for certain files or file types in a server generated directory listing. The AddDescription directive supports listing specific files, wildcard expressions, or file extensions.

AddEncoding

AddEncoding names file name extensions which should specify a particular encoding type. AddEncoding can also be used to instruct some browsers to uncompress certain files as they are downloaded.

AddHandler

AddHandler maps file extensions to specific handlers. For example, the cgi-script handler can be matched with the extension .cgi to automatically treat a file ending with .cgi as a CGI script. The following is a sample AddHandler directive for the .cgi extension.

```
AddHandler cgi-script .cgi
```

This directive enables CGIs outside of the cgi-bin to function in any directory on the server which has the ExecCGI option within the directories container. Refer to Directory for more information about setting the ExecCGI option for a directory.

In addition to CGI scripts, the AddHandler directive is used to process server-parsed HTML and image-map files.

AddIcon

AddIcon specifies which icon to show in server generated directory listings for files with certain extensions. For example, the Web server is set to show the icon binary.gif for files with .bin or .exe extensions.

AddIconByEncoding

This directive names icons which are displayed by files with MIME encoding in server generated directory listings. For example, by default, the Web server shows the compressed.gif icon next to MIME encoded x-compress and x-gzip files in server generated directory listings.

AddIconByType

This directive names icons which are displayed next to files with MIME types in server generated directory listings. For example, the server shows the icon text.gif next to files with a mime-type of text, in server generated directory listings.

AddLanguage

AddLanguage associates file name extensions with specific languages. This directive is useful for Apache HTTP Servers which serve content in multiple languages based on the client Web browser's language settings.

AddType

Use the AddType directive to define or override a default MIME type and file extension pairs. The following example directive tells the Apache HTTP Server to recognize the .tgz file extension:

```
AddType application/x-tar .tgz
```

Alias

The Alias setting allows directories outside the DocumentRoot directory to be accessible. Any URL ending in the alias automatically resolves to the alias' path. By default, one alias for an icons/ directory is already set up. An icons/ directory can be accessed by the Web server, but the directory is not in the DocumentRoot.

Allow

Allow specifies which client can access a given directory. The client can be all, a domain name, an IP address, a partial IP address, a network/netmask pair, and so on. The DocumentRoot directory is configured to Allow requests from all, meaning everyone has access.

AllowOverride

The AllowOverride directive sets whether any Options can be overridden by the declarations in an .htaccess file. By default, both the root directory and the DocumentRoot are set to allow no .htaccess overrides.

BrowserMatch

The BrowserMatch directive allows the server to define environment variables and take appropriate actions based on the User-Agent HTTP header field — which identifies the client's Web browser type. By default, the Web server uses BrowserMatch to deny connections to specific browsers with known problems and also to disable keepalives and HTTP header flushes for browsers that are known to have problems with those actions.

Cache Directives

A number of commented cache directives are supplied by the default Apache HTTP Server configuration file. In most cases, uncommenting these lines by removing the hash mark (#) from the beginning of the line is sufficient. The following, however, is a list of some of the more important cache-related directives.

- CacheEnable — Specifies whether the cache is a disk, memory, or file descriptor cache. By default CacheEnable configures a disk cache for URLs at or below /.
- CacheRoot — Specifies the name of the directory containing cached files. The default CacheRoot is the /var/httpd/proxy/ directory.
- CacheSize — Specifies how much space the cache can use in kilobytes. The default CacheSize is 5 KB.

The following is a list of some of the other common cache-related directives.

- CacheMaxExpire — Specifies how long HTML documents are retained (without a reload from the originating Web server) in the cache. The default is 24 hours (86400 seconds).
- CacheLastModifiedFactor — Specifies the creation of an expiry (expiration) date for a document which did not come from its originating server with its own expiry set. The default CacheLastModifiedFactor is set to 0.1, meaning that the expiry date for such documents equals one-tenth of the amount of time since the document was last modified.
- CacheDefaultExpire — Specifies the expiry time in hours for a document that was received using a protocol that does not support expiry times. The default is set to 1 hour (3600 seconds).
- NoProxy — Specifies a space-separated list of subnets, IP addresses, domains, or hosts whose content is not cached. This setting is most useful for Intranet sites.

CacheNegotiatedDocs

By default, the Web server asks proxy servers not to cache any documents which were negotiated on the basis of content (that is, they may change over time or because of the input from the requester). If CacheNegotiatedDocs is set to on, this function is disabled and proxy servers are allowed to cache such documents.

CustomLog

CustomLog identifies the log file and the log file format. By default, the access log is recorded to the /var/log/httpd/access_log file while errors are recorded in the /var/log/httpd/error_log file.

The default CustomLog format is the combined log file format, as illustrated here:

```
remotehost rfc931 user date "request" status bytes referrer user-agent
```

DefaultIcon

DefaultIcon specifies the icon displayed in server generated directory listings for files which have no other icon specified. The unknown.gif image file is the default.

DefaultType

DefaultType sets a default content type for the Web server to use for documents whose MIME types cannot be determined. The default is text/plain.

Deny

Deny works similar to Allow, except it specifies who is denied access. The DocumentRoot is not configured to Deny requests from anyone by default.

Directory

<Directory /path/to/directory> and </Directory> tags create a container used to enclose a group of configuration directives which apply only to a specific directory and its subdirectories. Any directive which is applicable to a directory may be used within Directory tags.

By default, very restrictive parameters are applied to the root directory (/), using the Options (refer to Options) and AllowOverride (refer to AllowOverride) directives. Under this configuration, any directory on the system which needs more permissive settings has to be explicitly given those settings.

In the default configuration, another Directory container is configured for the DocumentRoot which assigns less rigid parameters to the directory tree so that the Apache HTTP Server can access the files residing there.

The Directory container can be also be used to configure additional cgi-bin directories for server-side applications outside of the directory specified in the ScriptAlias directive (refer to ScriptAlias for more information).

To accomplish this, the Directory container must set the ExecCGI option for that directory.

For example, if CGI scripts are located in /home/my_cgi_directory, add the following Directory container to the httpd.conf file:

```
<Directory /home/my_cgi_directory>
  Options +ExecCGI
```

```
</Directory>
```

Next, the AddHandler directive must be uncommented to identify files with the .cgi extension as CGI scripts. Refer to AddHandler for instructions on setting AddHandler.

For this to work, permissions for CGI scripts, and the entire path to the scripts, must be set to 0755.

DirectoryIndex

The DirectoryIndex is the default page served by the server when a user requests an index of a directory by specifying a forward slash (/) at the end of the directory name.

When a user requests the page http://example/this_directory/, they get either the DirectoryIndex page, if it exists, or a server-generated directory list. The default for DirectoryIndex is index.html and the index.html.var type map. The server tries to find either of these files and returns the first one it finds. If it does not find one of these files and Options Indexes is set for that directory, the server generates and returns a listing, in HTML format, of the subdirectories and files within the directory, unless the directory listing feature is turned off.

DocumentRoot

DocumentRoot is the directory which contains most of the HTML files which are served in response to requests. The default DocumentRoot, for both the non-secure and secure Web servers, is the /var/www/html directory. For example, the server might receive a request for the following document:

```
http://example.com/foo.html
```

The server looks for the following file in the default directory:

```
/var/www/html/foo.html
```

To change the DocumentRoot so that it is not shared by the secure and the non-secure Web servers, refer to 23.7절. “Virtual Hosts” .

ErrorDocument

The ErrorDocument directive associates an HTTP response code with a message or a URL to be sent back to the client. By default, the Web server outputs a simple and usually cryptic error message when an error occurs. The ErrorDocument directive forces the Web server to instead output a customized message or page.

> **Important**
>
> To be valid, the message must be enclosed in a pair of double quotes ".

ErrorLog

ErrorLog specifies the file where server errors are logged. By default, this directive is set to /var/log/httpd/error_log.

ExtendedStatus

The ExtendedStatus directive controls whether Apache generates basic (off) or detailed server status information (on), when the server-status handler is called. The server-status handler is called using Location tags. More information on calling server-status is included in Location.

Group

Specifies the group name of the Apache HTTP Server processes.

This directive has been deprecated for the configuration of virtual hosts.

By default, Group is set to apache.

HeaderName

HeaderName names the file which, if it exists in the directory, is prepended to the start of server generated directory listings. Like ReadmeName, the server tries to include it as an HTML document if possible or in plain text if not.

HostnameLookups

HostnameLookups can be set to on, off, or double. If HostnameLookups is set to on, the server automatically resolves the IP address for each connection. Resolving the IP address means that the server makes one or more connections to a DNS server, adding processing overhead. If HostnameLookups is set to double, the server performs a double-reverse DNS look up adding even more processing overhead.

To conserve resources on the server, HostnameLookups is set to off by default.

If hostnames are required in server log files, consider running one of the many log analyzer tools that perform the DNS lookups more efficiently and in bulk when rotating the Web server log files.

IfDefine

The IfDefine tags surround configuration directives that are applied if the "test" stated in the IfDefine tag is true. The directives are ignored if the test is false.

The test in the IfDefine tags is a parameter name (for example, HAVE_PERL). If the parameter is defined, meaning that it is provided as an argument to the server's start-up command, then the test is true. In this case, when the Web server is started, the test is true and the directives contained in the IfDefine tags are applied.

IfModule

<IfModule> and </IfModule> tags create a conditional container which are only activated if the specified module is loaded. Directives within the IfModule container are processed under one of two conditions. The directives are processed if the module contained within the starting <IfModule> tag is loaded. Or, if an exclamation point ! appears before the module name, the directives are processed only if the module specified in the <IfModule> tag is not loaded.

For more information about Apache HTTP Server modules, refer to 23.6절. “Adding Modules” .

Include

Include allows other configuration files to be included at runtime.

The path to these configuration files can be absolute or relative to the ServerRoot.

Important

For the server to use individually packaged modules, such as mod_ssl, mod_perl, and php, the following directive must be included in Section 1: Global Environment of httpd.conf:

```
Include conf.d/*.conf
```

IndexIgnore

IndexIgnore lists file extensions, partial file names, wildcard expressions, or full file names. The Web server does not include any files which match any of those parameters in server generated directory listings.

IndexOptions

IndexOptions controls the appearance of server generated directing listings, by adding icons, file descriptions, and so on. If Options Indexes is set (refer to Options), the Web server generates a directory listing when the Web server receives an HTTP request for a directory without an index.

First, the Web server looks in the requested directory for a file matching the names listed in the DirectoryIndex directive (usually, index.html). If an index.html file is not found, Apache HTTP Server creates an HTML directory listing of the requested directory. The appearance of this directory listing is controlled, in part, by the IndexOptions directive.

The default configuration turns on FancyIndexing. This means that a user can re-sort a directory listing by clicking on column headers. Another click on the same header switches from ascending to descending order. FancyIndexing also shows different icons for different files, based upon file extensions.

The AddDescription option, when used in conjunction with FancyIndexing, presents a short description for the file in server generated directory listings.

IndexOptions has a number of other parameters which can be set to control the appearance of server generated directories. The IconHeight and IconWidth parameters require the server to include HTML HEIGHT and WIDTH tags for the icons in server generated webpages. The IconsAreLinks parameter combines the graphical icon with the HTML link anchor, which contains the URL link target.

KeepAlive

KeepAlive sets whether the server allows more than one request per connection and can be used to prevent any one client from consuming too much of the server's resources.

By default Keepalive is set to off. If Keepalive is set to on and the server becomes very busy, the server can quickly spawn the maximum number of child processes. In this situation, the server slows down significantly. If Keepalive is enabled, it is a good idea to set the KeepAliveTimeout low (refer to KeepAliveTimeout for more information about the KeepAliveTimeout directive) and monitor the /var/log/httpd/error_log log file on the server. This log reports when the server is running out of child processes.

KeepAliveTimeout

KeepAliveTimeout sets the number of seconds the server waits after a request has been served before it closes the connection. Once the server receives a request, the Timeout directive applies instead. The KeepAliveTimeout directive is set to 15 seconds by default.

LanguagePriority

LanguagePriority sets precedence for different languages in case the client Web browser has no language preference set.

Listen

The Listen command identifies the ports on which the Web server accepts incoming requests. By default, the Apache HTTP Server is set to listen to port 80 for non-secure Web communications and (in the /etc/httpd/conf.d/ssl.conf file which defines any secure servers) to port 443 for secure Web communications.

If the Apache HTTP Server is configured to listen to a port under 1024, only the root user can start it. For port 1024 and above, httpd can be started as a regular user.

The Listen directive can also be used to specify particular IP addresses over which the server accepts connections.

LoadModule

LoadModule is used to load Dynamic Shared Object (DSO) modules. More information on the Apache HTTP Server's DSO support, including instructions for using the LoadModule directive, can be found in 23.6절. "Adding Modules" . Note, the load order of the modules is no longer important with Apache HTTP Server 2.0. Refer to 23.2.2.1.3절. "Dynamic Shared Object (DSO) Support" for more information about Apache HTTP Server 2.0 DSO support.

Location

The <Location> and </Location> tags create a container in which access control based on URL can be specified.

For instance, to allow people connecting from within the server's domain to see status reports, use the following directives:

```
<Location /server-status>
  SetHandler server-status
  Order deny,allow
  Deny from all
  Allow from <.example.com>
</Location>
```

Replace <.example.com> with the second-level domain name for the Web server.

To provide server configuration reports (including installed modules and configuration directives) to requests from inside the domain, use the following directives:

```
<Location /server-info>
  SetHandler server-info
  Order deny,allow
  Deny from all
  Allow from <.example.com>
</Location>
```

Again, replace <.example.com> with the second-level domain name for the Web server.

LogFormat

The LogFormat directive configures the format of the various Web server log files. The actual LogFormat used depends on the settings given in the CustomLog directive (refer to CustomLog).

The following are the format options if the CustomLog directive is set to combined:

%h (remote host's IP address or hostname)
Lists the remote IP address of the requesting client. If HostnameLookups is set to on, the client hostname is recorded unless it is not available from DNS.

%l (rfc931)
Not used. A hyphen - appears in the log file for this field.

%u (authenticated user)
Lists the username of the user recorded if authentication was required. Usually, this is not used, so a hyphen - appears in the log file for this field.

%t (date)
Lists the date and time of the request.

%r (request string)
Lists the request string exactly as it came from the browser or client.

%s (status)
Lists the HTTP status code which was returned to the client host.

%b (bytes)
Lists the size of the document.

%\"%{Referer}i\" (referrer)
Lists the URL of the webpage which referred the client host to Web server.

%\"%{User-Agent}i\" (user-agent)
Lists the type of Web browser making the request.

LogLevel

LogLevel sets how verbose the error messages in the error logs are. LogLevel can be set (from least verbose to most verbose) to emerg, alert, crit, error, warn, notice, info, or debug. The default LogLevel is warn.

MaxKeepAliveRequests

This directive sets the maximum number of requests allowed per persistent connection. The Apache Project recommends a high setting, which improves the server's performance. MaxKeepAliveRequests is set to 100 by default, which should be appropriate for most situations.

NameVirtualHost

The NameVirtualHost directive associates an IP address and port number, if necessary, for any name-based virtual hosts. Name-based virtual hosting allows one Apache HTTP Server to serve different domains without using multiple IP addresses.

> **Note**
>
> Name-based virtual hosts only work with non-secure HTTP connections. If using virtual hosts with a secure server, use IP address-based virtual hosts instead.

To enable name-based virtual hosting, uncomment the NameVirtualHost configuration directive and add the correct IP address. Then add additional VirtualHost containers for each virtual host as is necessary for your configuration.

Options

The Options directive controls which server features are available in a particular directory. For example, under the restrictive parameters specified for the root directory, Options is only set to the FollowSymLinks directive. No features are enabled, except that the server is allowed to follow symbolic links in the root directory.

By default, in the DocumentRoot directory, Options is set to include Indexes and FollowSymLinks. Indexes permits the server to generate a directory listing for a directory if no DirectoryIndex (for example, index.html) is specified. FollowSymLinks allows the server to follow symbolic links in that directory.

> **Note**
>
> Options statements from the main server configuration section need to be replicated to each VirtualHost container individually. Refer to VirtualHost for more information.

Order

The Order directive controls the order in which allow and deny directives are evaluated. The server is configured to evaluate the Allow directives before the Deny directives for the DocumentRoot directory.

PidFile

PidFile names the file where the server records its process ID (PID). By default the PID is listed in /var/run/httpd.pid.

Proxy

<Proxy *> and </Proxy> tags create a container which encloses a group of configuration directives meant to apply only to the proxy server. Many directives which are allowed within a <Directory> container may also be used within <Proxy> container.

ProxyRequests

To configure the Apache HTTP Server to function as a proxy server, remove the hash mark (#) from the beginning of the <IfModule mod_proxy.c> line, the ProxyRequests, and each line in the <Proxy> stanza. Set the ProxyRequests directive to On, and set which domains are allowed access to the server in the Allow from directive of the <Proxy> stanza.

ReadmeName

ReadmeName names the file which, if it exists in the directory, is appended to the end of server generated directory listings. The Web server first tries to include the file as an HTML document and then tries to include it as plain text. By default, ReadmeName is set to README.html.

Redirect

When a webpage is moved, Redirect can be used to map the file location to a new URL. The format is as follows:

```
Redirect /<old-path>/<file-name> http://<current-domain>/<current-path>/<file-name>
```

In this example, replace <old-path> with the old path information for <file-name> and <current-domain> and <current-path> with the current domain and path information for <file-name>.

In this example, any requests for <file-name> at the old location is automatically redirected to the new location.

For more advanced redirection techniques, use the mod_rewrite module included with the Apache HTTP Server. For more information about configuring the mod_rewrite module, refer to the Apache Software Foundation documentation online at http://httpd.apache.org/docs/2.2/mod/mod_rewrite.html[23].

[23] http://httpd.apache.org/docs/2.2/mod/mod_rewrite.html

ScriptAlias

The ScriptAlias directive defines where CGI scripts are located. Generally, it is not good practice to leave CGI scripts within the DocumentRoot, where they can potentially be viewed as text documents. For this reason, a special directory outside of the DocumentRoot directory containing server-side executables and scripts is designated by the ScriptAlias directive. This directory is known as a cgi-bin and is set to /var/www/cgi-bin/ by default.

It is possible to establish directories for storing executables outside of the cgi-bin/ directory. For instructions on doing so, refer to AddHandler and Directory.

ServerAdmin

Sets the ServerAdmin directive to the email address of the Web server administrator. This email address shows up in error messages on server-generated Web pages, so users can report a problem by sending email to the server administrator.

By default, ServerAdmin is set to root@localhost.

A common way to set up ServerAdmin is to set it to webmaster@example.com. Once set, alias webmaster to the person responsible for the Web server in /etc/aliases and run /usr/bin/newaliases.

ServerName

ServerName specifies a hostname and port number (matching the Listen directive) for the server. The ServerName does not need to match the machine's actual hostname. For example, the Web server may be www.example.com, but the server's hostname is actually foo.example.com. The value specified in ServerName must be a valid Domain Name Service (DNS) name that can be resolved by the system — do not make something up.

The following is a sample ServerName directive:

```
ServerName www.example.com:80
```

When specifying a ServerName, be sure the IP address and server name pair are included in the /etc/hosts file.

ServerRoot

The ServerRoot directive specifies the top-level directory containing website content. By default, ServerRoot is set to "/etc/httpd" for both secure and non-secure servers.

ServerSignature

The ServerSignature directive adds a line containing the Apache HTTP Server server version and the ServerName to any server-generated documents, such as error messages sent back to clients. ServerSignature is set to on by default.

ServerSignature can be set to EMail which adds a mailto:ServerAdmin HTML tag to the signature line of auto-generated responses. ServerSignature can also be set to Off to stop Apache from sending out its version number and module information. Please also check the ServerTokens settings.

ServerTokens

The ServerTokens directive determines if the Server response header field sent back to clients should include details of the Operating System type and information about compiled-in modules. By default, ServerTokens is set to Full which sends information about the Operating System type and compiled-in modules. Setting the ServerTokens to Prod sends the product name only and is recommended as many hackers check information in the Server header when scanning for vulnerabilities. You can also set the ServerTokens to Min (minimal) or to OS (operating system).

SuexecUserGroup

The SuexecUserGroup directive, which originates from the mod_suexec module, allows the specification of user and group execution privileges for CGI programs. Non-CGI requests are still processed with the user and group specified in the User and Group directives.

> **Note**
>
> From version 2.0, the SuexecUserGroup directive replaced the Apache HTTP Server 1.3 configuration of using the User and Group directives inside the configuration of VirtualHosts sections.

Timeout

Timeout defines, in seconds, the amount of time that the server waits for receipts and transmissions during communications. Timeout is set to 300 seconds by default, which is appropriate for most situations.

TypesConfig

TypesConfig names the file which sets the default list of MIME type mappings (file name extensions to content types). The default TypesConfig file is /etc/mime.types. Instead of editing /etc/mime.types, the recommended way to add MIME type mappings is to use the AddType directive.

For more information about AddType, refer to AddType.

UseCanonicalName

When set to on, this directive configures the Apache HTTP Server to reference itself using the value specified in the ServerName and Port directives. When UseCanonicalName is set to off, the server instead uses the value used by the requesting client when referring to itself.

UseCanonicalName is set to off by default.

User

The User directive sets the username of the server process and determines what files the server is allowed to access. Any files inaccessible to this user are also inaccessible to clients connecting to the Apache HTTP Server.

By default User is set to apache.

This directive has been deprecated for the configuration of virtual hosts.

> **Note**
>
> For security reasons, the Apache HTTP Server does not run as the root user.

UserDir

UserDir is the subdirectory within each user's home directory where they should place personal HTML files which are served by the Web server. This directive is set to disable by default.

The name for the subdirectory is set to public_html in the default configuration. For example, the server might receive the following request:

```
http://example.com/~username/foo.html
```

The server would look for the file:

```
/home/username/public_html/foo.html
```

In the above example, /home/username/ is the user's home directory (note that the default path to users' home directories may vary).

Make sure that the permissions on the users' home directories are set correctly. Users' home directories must be set to 0711. The read (r) and execute (x) bits must be set on the users' public_html directories (0755 also works). Files that are served in a users' public_html directories must be set to at least 0644.

VirtualHost

<VirtualHost> and </VirtualHost> tags create a container outlining the characteristics of a virtual host. The VirtualHost container accepts most configuration directives.

A commented VirtualHost container is provided in httpd.conf, which illustrates the minimum set of configuration directives necessary for each virtual host. Refer to 23.7절. "Virtual Hosts" for more information about virtual hosts.

> **Note**
>
> The default SSL virtual host container now resides in the file /etc/httpd/conf.d/ssl.conf.

23.5.2. Configuration Directives for SSL

The directives in /etc/httpd/conf.d/ssl.conf file can be configured to enable secure Web communications using SSL and TLS.

SetEnvIf

SetEnvIf sets environment variables based on the headers of incoming connections. It is not solely an SSL directive, though it is present in the supplied /etc/httpd/conf.d/ssl.conf file. It's purpose in this context is to disable HTTP keepalive and to allow SSL to close the connection without a closing notification from the client browser. This setting is necessary for certain browsers that do not reliably shut down the SSL connection.

For more information on other directives within the SSL configuration file, refer to the following URLs:

- http://localhost/manual/mod/mod_ssl.html
- http://httpd.apache.org/docs/2.2/mod/mod_ssl.html

Note

In most cases, SSL directives are configured appropriately during the installation of Red Hat Enterprise Linux. Be careful when altering Apache HTTP Secure Server directives, misconfiguration can lead to security vulnerabilities.

23.5.3. MPM Specific Server-Pool Directives

As explained in 23.2.2.1.2절. "Server-Pool Size Regulation" , the responsibility for managing characteristics of the server-pool falls to a module group called MPMs under Apache HTTP Server 2.0. The characteristics of the server-pool differ depending upon which MPM is used. For this reason, an IfModule container is necessary to define the server-pool for the MPM in use.

By default, Apache HTTP Server 2.0 defines the server-pool for both the prefork and worker MPMs.

The following section list directives found within the MPM-specific server-pool containers.

MaxClients

MaxClients sets a limit on the total number of server processes, or simultaneously connected clients, that can run at one time. The main purpose of this directive is to keep a runaway Apache HTTP Server from crashing the operating system. For busy servers this value should be set to a high value. The server's default is set to 150 regardless of the MPM in use. However, it is not recommended that the value for MaxClients exceeds 256 when using the prefork MPM.

MaxRequestsPerChild

MaxRequestsPerChild sets the total number of requests each child server process serves before the child dies. The main reason for setting MaxRequestsPerChild is to avoid long-lived process induced memory leaks. The default MaxRequestsPerChild for the prefork MPM is 4000 and for the worker MPM is 0.

MinSpareServers and MaxSpareServers

These values are only used with the prefork MPM. They adjust how the Apache HTTP Server dynamically adapts to the perceived load by maintaining an appropriate number of spare server processes based on the number of incoming requests. The server checks the number of servers waiting for a request and kills some if there are more than MaxSpareServers or creates some if the number of servers is less than MinSpareServers.

The default MinSpareServers value is 5; the default MaxSpareServers value is 20. These default settings should be appropriate for most situations. Be careful not to increase the MinSpareServers to a large number as doing so creates a heavy processing load on the server even when traffic is light.

MinSpareThreads and MaxSpareThreads

These values are only used with the worker MPM. They adjust how the Apache HTTP Server dynamically adapts to the perceived load by maintaining an appropriate number of spare server threads based on the number of incoming requests. The server checks the number of server threads waiting for a request and kills some if there are more than MaxSpareThreads or creates some if the number of servers is less than MinSpareThreads.

The default MinSpareThreads value is 25; the default MaxSpareThreads value is 75. These default settings should be appropriate for most situations. The value for MaxSpareThreads must be greater than or equal to the sum of MinSpareThreads and ThreadsPerChild, else the Apache HTTP Server automatically corrects it.

StartServers

The StartServers directive sets how many server processes are created upon startup. Since the Web server dynamically kills and creates server processes based on traffic load, it is not necessary to change this parameter. The Web server is set to start 8 server processes at startup for the prefork MPM and 2 for the worker MPM.

ThreadsPerChild

This value is only used with the worker MPM. It sets the number of threads within each child process. The default value for this directive is 25.

23.6. Adding Modules

The Apache HTTP Server is distributed with a number of modules. More information about Apache HTTP modules can be found on http://httpd.apache.org/docs/2.2/mod/.

The Apache HTTP Server supports Dynamically Shared Objects (DSOs), or modules, which can easily be loaded at runtime as necessary.

The Apache Project provides complete DSO documentation online at http://httpd.apache.org/docs/2.2/dso.html. Or, if the http-manual package is installed, documentation about DSOs can be found online at http://localhost/manual/mod/.

For the Apache HTTP Server to use a DSO, it must be specified in a LoadModule directive within /etc/httpd/conf/httpd.conf. If the module is provided by a separate package, the line must appear within the modules configuration file in the /etc/httpd/conf.d/ directory. Refer to LoadModule for more information.

If adding or deleting modules from http.conf, Apache HTTP Server must be reloaded or restarted, as referred to in 23.3절. “Starting and Stopping httpd” .

If creating a new module, first install the httpd-devel package which contains the include files, the header files, as well as the APache eXtenSion (/usr/sbin/apxs) application, which uses the include files and the header files to compile DSOs.

After writing a module, use /usr/sbin/apxs to compile the module sources outside the Apache source tree. For more information about using the /usr/sbin/apxs command, refer to the Apache documentation online at http://httpd.apache.org/docs/2.2/dso.html as well as the apxs man page.

Once compiled, put the module in the /usr/lib/httpd/modules/ directory. For RHEL platforms using default-64-bit userspace (x86_64, ia64, ?) this path will be /usr/lib64/httpd/modules/. Then add a LoadModule line to the httpd.conf, using the following structure:

```
LoadModule <module-name> <path/to/module.so>
```

Where <module-name> is the name of the module and <path/to/module.so> is the path to the DSO.

23.7. Virtual Hosts

The Apache HTTP Server's built in virtual hosting allows the server to provide different information based on which IP address, hostname, or port is being requested. A complete guide to using virtual hosts is available online at http://httpd.apache.org/docs/2.2/vhosts/.

23.7.1. Setting Up Virtual Hosts

To create a name-based virtual host, it is best to use the virtual host container provided in httpd.conf as an example.

The virtual host example read as follows:

```
#NameVirtualHost *:80
#
#<VirtualHost *:80>
# ServerAdmin webmaster@dummy-host.example.com
# DocumentRoot /www/docs/dummy-host.example.com
# ServerName dummy-host.example.com
# ErrorLog logs/dummy-host.example.com-error_log
# CustomLog logs/dummy-host.example.com-access_log common #</VirtualHost>
```

To activate name-based virtual hosting, uncomment the NameVirtualHost line by removing the hash mark (#) and replace the asterisk (*) with the IP address assigned to the machine.

Next, configure a virtual host by uncommenting and customizing the <VirtualHost> container.

On the <VirtualHost> line, change the asterisk (*) to the server's IP address. Change the ServerName to a valid DNS name assigned to the machine, and configure the other directives as necessary.

The <VirtualHost> container is highly customizable and accepts almost every directive available within the main server configuration.

Tip

If configuring a virtual host to listen on a non-default port, that port must be added to the Listen directive in the global settings section of /etc/httpd/conf/httpd.conf file.

To activate a newly created virtual host, the Apache HTTP Server must be reloaded or restarted. Refer to 23.3절. "Starting and Stopping httpd" for further instructions.

Comprehensive information about creating and configuring both name-based and IP address-based virtual hosts is provided online at http://httpd.apache.org/docs/2.2/vhosts/.

23.8. Apache HTTP Secure Server Configuration

This section provides basic information on the Apache HTTP Server with the mod_ssl security module enabled to use the OpenSSL library and toolkit. The combination of these three components are referred to in this section as the secure Web server or just as the secure server.

The mod_ssl module is a security module for the Apache HTTP Server. The mod_ssl module uses the tools provided by the OpenSSL Project to add a very important feature to the Apache HTTP Server — the ability to encrypt communications. In contrast, regular HTTP communications between a browser and a Web server are sent in plain text, which could be intercepted and read by someone along the route between the browser and the server.

This section is not meant to be complete and exclusive documentation for any of these programs. When possible, this guide points to appropriate places where you can find more in-depth documentation on particular subjects.

This section shows you how to install these programs. You can also learn the steps necessary to generate a private key and a certificate request, how to generate your own self-signed certificate, and how to install a certificate to use with your secure server.

The mod_ssl configuration file is located at /etc/httpd/conf.d/ssl.conf. For this file to be loaded, and hence for mod_ssl to work, you must have the statement Include conf.d/*.conf in the /etc/httpd/conf/httpd.conf file. This statement is included by default in the default Apache HTTP Server configuration file.

23.8.1. An Overview of Security-Related Packages

To enable the secure server, you must have the following packages installed at a minimum:

httpd

The httpd package contains the httpd daemon and related utilities, configuration files, icons, Apache HTTP Server modules, man pages, and other files used by the Apache HTTP Server.

mod_ssl

The mod_ssl package includes the mod_ssl module, which provides strong cryptography for the Apache HTTP Server via the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols.

openssl

The openssl package contains the OpenSSL toolkit. The OpenSSL toolkit implements the SSL and TLS protocols, and also includes a general purpose cryptography library.

Additionally, other software packages provide certain security functionalities (but are not required by the secure server to function):

23.8.2. An Overview of Certificates and Security

Your secure server provides security using a combination of the Secure Sockets Layer (SSL) protocol and (in most cases) a digital certificate from a Certificate Authority (CA). SSL handles the encrypted communications as well as the mutual authentication between browsers and your secure server. The CA-approved digital certificate provides authentication for your secure server (the CA puts its reputation behind its certification of your organization's identity). When your browser is communicating using SSL encryption, the https:// prefix is used at the beginning of the Uniform Resource Locator (URL) in the navigation bar.

Encryption depends upon the use of keys (think of them as secret encoder/decoder rings in data format). In conventional or symmetric cryptography, both ends of the transaction have the same key, which they use to decode each other's transmissions. In public or asymmetric cryptography, two keys co-exist: a public key and a private key. A person or an organization keeps their private key a secret and publishes their public key. Data encoded with the public key can only be decoded with the private key; data encoded with the private key can only be decoded with the public key.

To set up your secure server, use public cryptography to create a public and private key pair. In most cases, you send your certificate request (including your public key), proof of your company's identity, and payment to a CA. The CA verifies the certificate request and your identity, and then sends back a certificate for your secure server.

A secure server uses a certificate to identify itself to Web browsers. You can generate your own certificate (called a "self-signed" certificate), or you can get a certificate from a CA. A certificate from a reputable CA guarantees that a website is associated with a particular company or organization.

Alternatively, you can create your own self-signed certificate. Note, however, that self-signed certificates should not be used in most production environments. Self-signed certificates are not automatically accepted by a user's browser — users are prompted by the browser to accept the certificate and create the secure connection. Refer to 23.8.4절. "Types of Certificates" for more information on the differences between self-signed and CA-signed certificates.

Once you have a self-signed certificate or a signed certificate from the CA of your choice, you must install it on your secure server.

23.8.3. Using Pre-Existing Keys and Certificates

If you already have an existing key and certificate (for example, if you are installing the secure server to replace another company's secure server product), you can probably use your existing key and certificate with the secure server. The following two situations provide instances where you are not able to use your existing key and certificate:

- If you are changing your IP address or domain name — Certificates are issued for a particular IP address and domain name pair. You must get a new certificate if you are changing your IP address or domain name.

- If you have a certificate from VeriSign and you are changing your server software — VeriSign is a widely used CA. If you already have a VeriSign certificate for another purpose, you may have been considering using your existing VeriSign certificate with your new secure server. However, you are not be allowed to because VeriSign issues certificates for one specific server software and IP address/domain name combination.

  If you change either of those parameters (for example, if you previously used a different secure server product), the VeriSign certificate you obtained to use with the previous configuration will not work with the new configuration. You must obtain a new certificate.

If you have an existing key and certificate that you can use, you do not have to generate a new key and obtain a new certificate. However, you may need to move and rename the files which contain your key and certificate.

Move your existing key file to:

```
/etc/pki/tls/private/server.key
```

Move your existing certificate file to:

```
/etc/pki/tls/certs/server.crt
```

If you are upgrading from the Red Hat Secure Web Server, your old key (httpsd.key) and certificate (httpsd.crt) are located in /etc/httpd/conf/. Move and rename your key and certificate so that the secure server can use them. Use the following two commands to move and rename your key and certificate files:

```
mv /etc/httpd/conf/httpsd.key /etc/pki/tls/private/server.key
mv /etc/httpd/conf/httpsd.crt /etc/pki/tls/certs/server.crt
```

Then, start your secure server with the command:

```
service httpd start
```

23.8.4. Types of Certificates

If you installed your secure server from the RPM package provided by Red Hat, a randomly generated private key and a test certificate are generated and put into the appropriate directories. Before you begin using your secure server, however, you must generate your own key and obtain a certificate which correctly identifies your server.

You need a key and a certificate to operate your secure server — which means that you can either generate a self-signed certificate or purchase a CA-signed certificate from a CA. What are the differences between the two?

A CA-signed certificate provides two important capabilities for your server:

- Browsers (usually) automatically recognize the certificate and allow a secure connection to be made, without prompting the user.

- When a CA issues a signed certificate, they are guaranteeing the identity of the organization that is providing the webpages to the browser.

If your secure server is being accessed by the public at large, your secure server needs a certificate signed by a CA so that people who visit your website know that the website is owned by the organization who claims to own it. Before signing a certificate, a CA verifies that the organization requesting the certificate was actually who they claimed to be.

Most Web browsers that support SSL have a list of CAs whose certificates they automatically accept. If a browser encounters a certificate whose authorizing CA is not in the list, the browser asks the user to either accept or decline the connection.

You can generate a self-signed certificate for your secure server, but be aware that a self-signed certificate does not provide the same functionality as a CA-signed certificate. A self-signed certificate is not automatically recognized by most Web browsers and does not provide any guarantee concerning the identity of the organization that is providing the website. A CA-signed certificate provides both of these important capabilities for a secure server. If your secure server is to be used in a production environment, a CA-signed certificate is recommended.

The process of getting a certificate from a CA is fairly easy. A quick overview is as follows:

1. Create an encryption private and public key pair.

2. Create a certificate request based on the public key. The certificate request contains information about your server and the company hosting it.

3. Send the certificate request, along with documents proving your identity, to a CA. Red Hat does not make recommendations on which certificate authority to choose. Your decision may be based on your past experiences, on the experiences of your friends or colleagues, or purely on monetary factors.

   Once you have decided upon a CA, you need to follow the instructions they provide on how to obtain a certificate from them.

4. When the CA is satisfied that you are indeed who you claim to be, they provide you with a digital certificate.

5. Install this certificate on your secure server and begin handling secure transactions.

Whether you are getting a certificate from a CA or generating your own self-signed certificate, the first step is to generate a key. Refer to 23.8.5절. “Generating a Key” for instructions.

23.8.5. Generating a Key

You must be root to generate a key.

First, use the cd command to change to the /etc/pki/tls/ directory. Remove the fake key and certificate that were generated during the installation with the following commands:

```
rm private/server.key
rm certs/server.crt
```

The crypto-utils package contains the genkey utility which you can use to generate keys as the name implies. To create your own private key, please ensure the crypto-utils package is installed. You can view more options by typing man genkey in your terminal. Assuming you wish to generate keys for www.example.com using the genkey utility, type in the following command in your terminal:

```
genkey www.example.com
```

Please note that the make based process is no longer shipped with RHEL 5. This will start the genkey graphical user interface. The figure below illustrates the first screen. To navigate, use the keyboard arrow and tab keys. This windows indicates where your key will be stored and prompts you to proceed or cancel the operation. To proceed to the next step, select Next and press the Return (Enter) key.

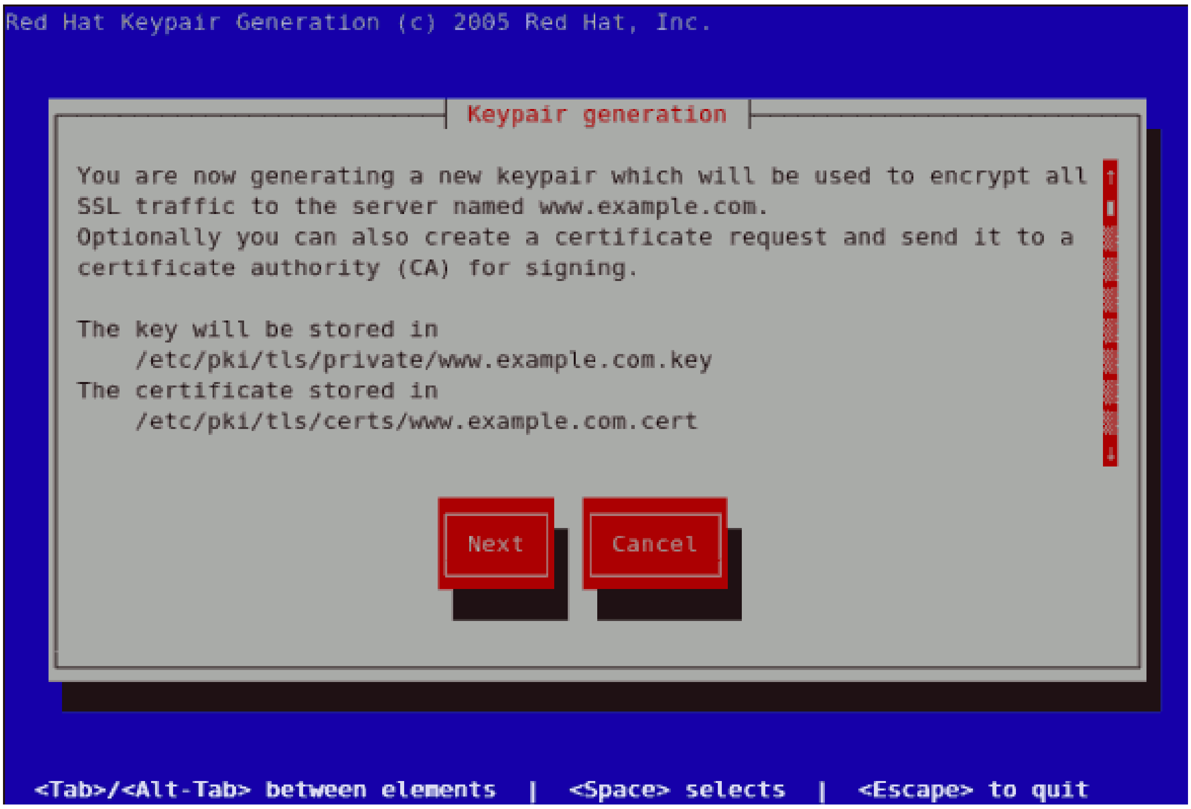

```
Red Hat Keypair Generation (c) 2005 Red Hat, Inc.

                         Keypair generation

  You are now generating a new keypair which will be used to encrypt all
  SSL traffic to the server named www.example.com.
  Optionally you can also create a certificate request and send it to a
  certificate authority (CA) for signing.

  The key will be stored in
      /etc/pki/tls/private/www.example.com.key
  The certificate stored in
      /etc/pki/tls/certs/www.example.com.cert

<Tab>/<Alt-Tab> between elements   |  <Space> selects  |  <Escape> to quit
```

그림 23.11. Keypair generation

The next screen prompts you to choose the size of your key. As indicated, the smaller the size of your key, the faster will the response from your server be and the lesser your level of security. On selecting your preferred, key size using the arrow keys, select Next to proceed to the next step. The figure below illustrates the key size selection screen.

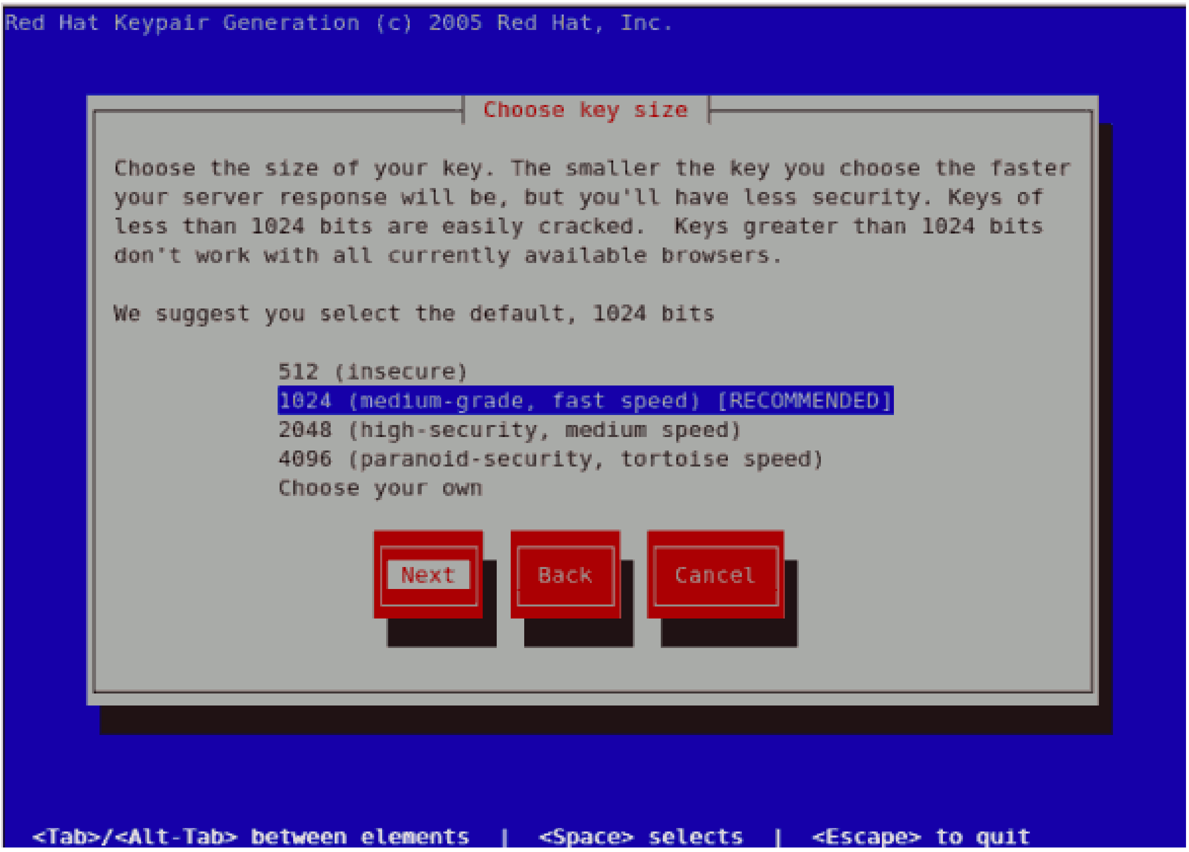

```
Red Hat Keypair Generation (c) 2005 Red Hat, Inc.

                    ┤ Choose key size ├

Choose the size of your key. The smaller the key you choose the faster
your server response will be, but you'll have less security. Keys of
less than 1024 bits are easily cracked.  Keys greater than 1024 bits
don't work with all currently available browsers.

We suggest you select the default, 1024 bits

            512 (insecure)
            1024 (medium-grade, fast speed) [RECOMMENDED]
            2048 (high-security, medium speed)
            4096 (paranoid-security, tortoise speed)
            Choose your own

<Tab>/<Alt-Tab> between elements   |  <Space> selects   |  <Escape> to quit
```

그림 23.12. Choose key size

Selecting the next step will initiate the random bits generation process which may take some time depending on the size of your selected key. The larger the size of your key, the longer it will take to generate it.

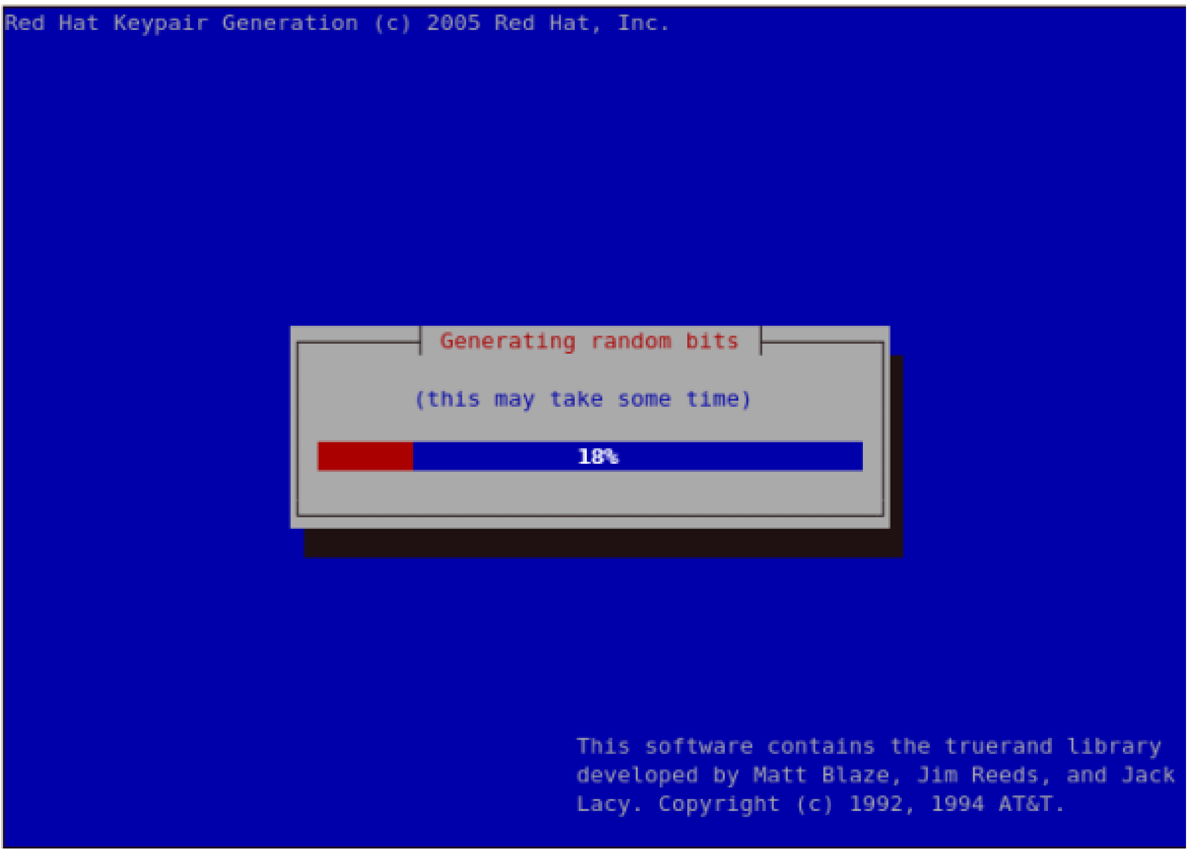

Red Hat Keypair Generation (c) 2005 Red Hat, Inc.

Generating random bits

(this may take some time)

18%

This software contains the truerand library
developed by Matt Blaze, Jim Reeds, and Jack
Lacy. Copyright (c) 1992, 1994 AT&T.

그림 23.13. Generating random bits

On generating your key, you will be prompted to send a Certificate Request (CSR) to a Certificate Authority (CA).

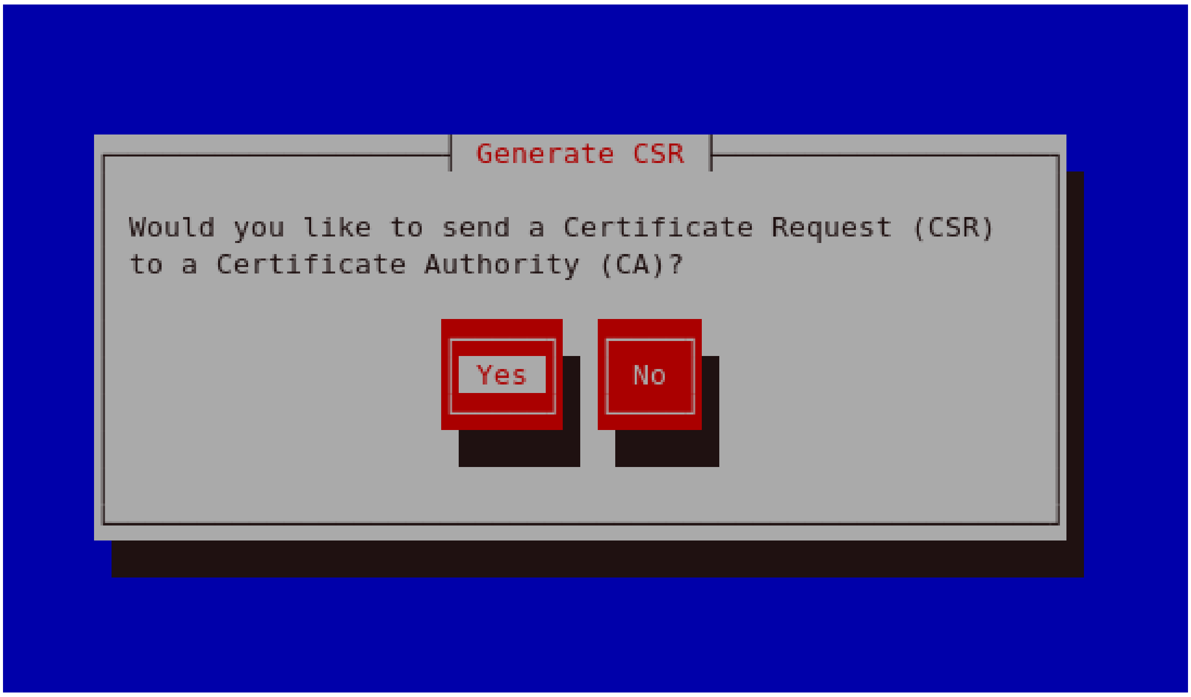

Generate CSR

Would you like to send a Certificate Request (CSR)
to a Certificate Authority (CA)?

그림 23.14. Generate CSR

Selecting Yes will prompt you to select the Certificate Authority you wish to send your request to. Selecting No will allow you to generate a self-signed certificate. The next step for this is illustrated in 그림 23.17. "Generating a self signed certificate for your server" .

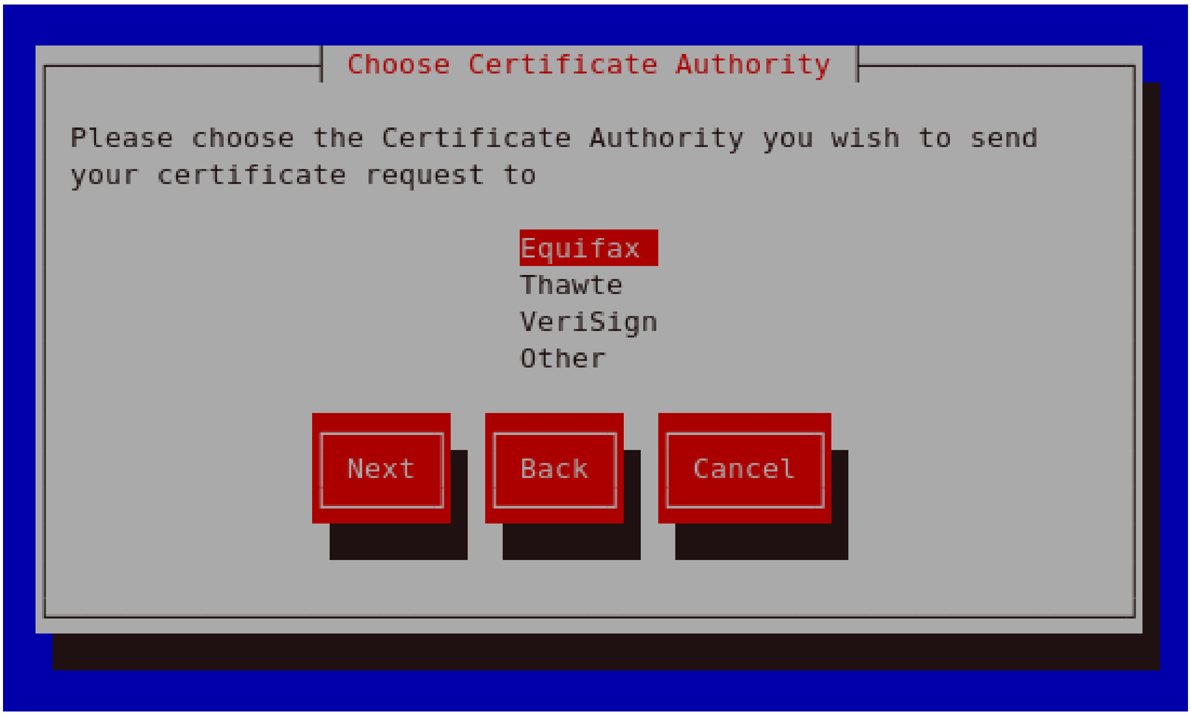

```
Choose Certificate Authority

Please choose the Certificate Authority you wish to send
your certificate request to

                    Equifax
                    Thawte
                    VeriSign
                    Other
```

그림 23.15. Choose Certificate Authority (CA)

On Selecting your preferred option, select Next to proceed to the next step. The next screen allows you to enter the details of your certificate.

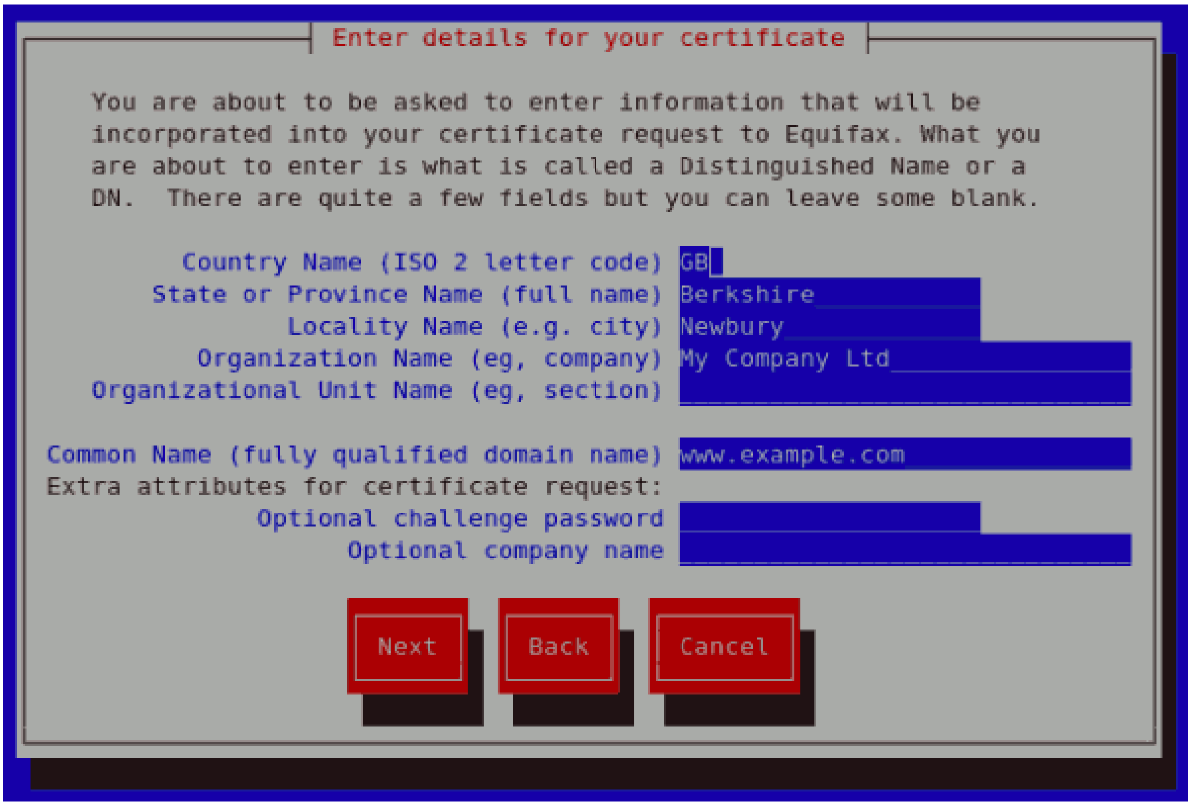

```
Enter details for your certificate

  You are about to be asked to enter information that will be
  incorporated into your certificate request to Equifax. What you
  are about to enter is what is called a Distinguished Name or a
  DN.  There are quite a few fields but you can leave some blank.

        Country Name (ISO 2 letter code) GB
      State or Province Name (full name) Berkshire
           Locality Name (e.g. city) Newbury
       Organization Name (eg, company) My Company Ltd
  Organizational Unit Name (eg, section)

Common Name (fully qualified domain name) www.example.com
Extra attributes for certificate request:
            Optional challenge password
               Optional company name
```

그림 23.16. Enter details for your certificate

If you prefer to generate a self signed cert key pair, you should not generate a CSR. To do this, select No as your preferred option in the Generate CSR screen. This will display the figure below from which you can enter your certificate details. Entering your certificate details and pressing the return key will display the 그림 23.19. "Protecting your private key" from which you can choose to encrypt your private key or not.

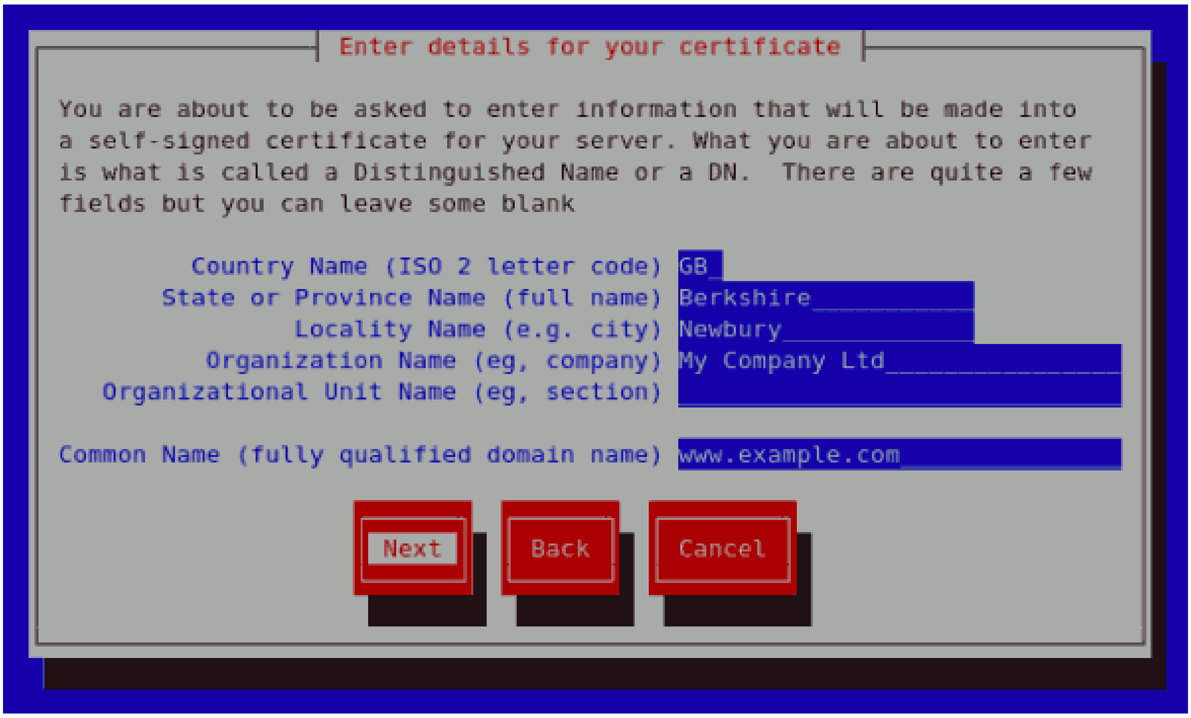

```
Enter details for your certificate

You are about to be asked to enter information that will be made into
a self-signed certificate for your server. What you are about to enter
is what is called a Distinguished Name or a DN.  There are quite a few
fields but you can leave some blank

          Country Name (ISO 2 letter code) GB
        State or Province Name (full name) Berkshire
                 Locality Name (e.g. city) Newbury
       Organization Name (eg, company) My Company Ltd
   Organizational Unit Name (eg, section)

Common Name (fully qualified domain name) www.example.com
```

그림 23.17. Generating a self signed certificate for your server

On entering the details of your certificate, select Next to proceed. The figure below illustrates an example of a the next screen displayed after completing the details for a certificate to be sent to Equifax. Please note that if you are generating a self signed key, for your server, this screen is not displayed.

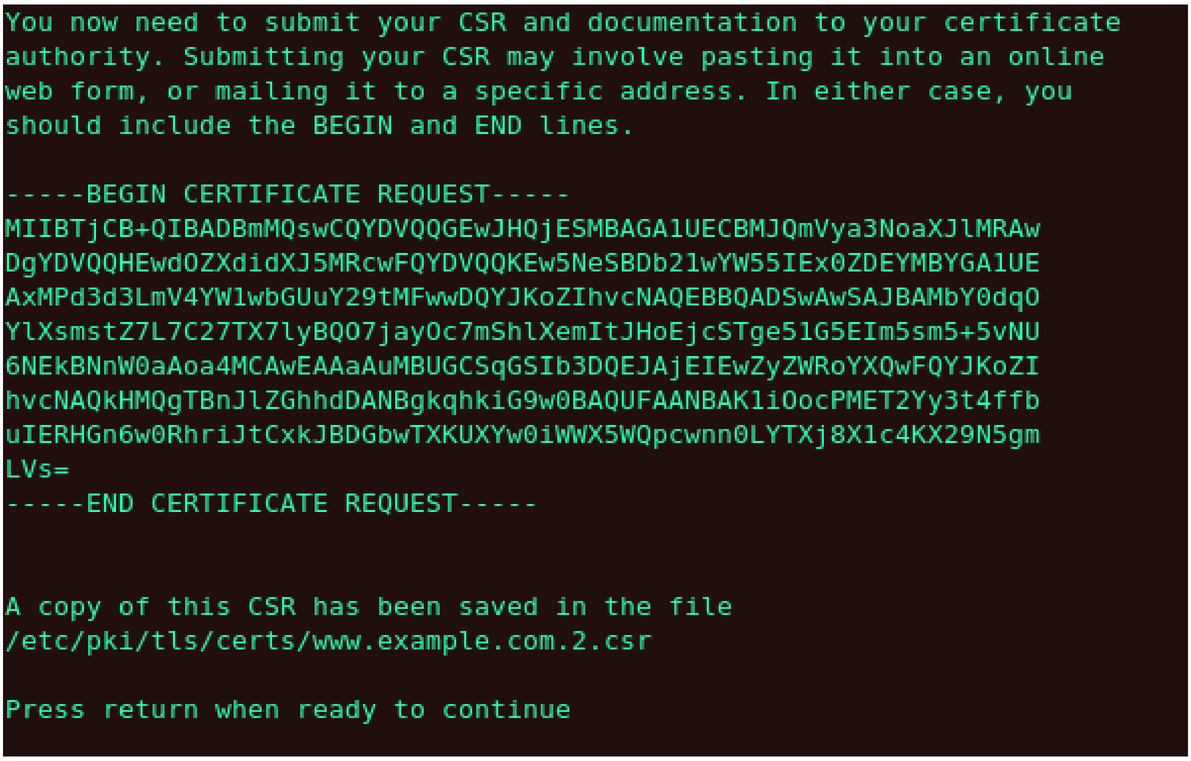

```
You now need to submit your CSR and documentation to your certificate
authority. Submitting your CSR may involve pasting it into an online
web form, or mailing it to a specific address. In either case, you
should include the BEGIN and END lines.

-----BEGIN CERTIFICATE REQUEST-----
MIIBTjCB+QIBADBmMQswCQYDVQQGEwJHQjESMBAGA1UECBMJQmVya3NoaXJlMRAw
DgYDVQQHEwdOZXdidXJ5MRcwFQYDVQQKEw5NeSBDb21wYW55IEx0ZDEYMBYGA1UE
AxMPd3d3LmV4YW1wbGUuY29tMFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAMbY0dqO
YlXsmstZ7L7C27TX7lyBQO7jayOc7mShlXemItJHoEjcSTge51G5EIm5sm5+5vNU
6NEkBNnW0aAoa4MCAwEAAaAuMBUGCSqGSIb3DQEJAjEIEwZyZWRoYXQwFQYJKoZI
hvcNAQkHMQgTBnJlZGhhdDANBgkqhkiG9w0BAQUFAANBAK1iOocPMET2Yy3t4ffb
uIERHGn6w0RhriJtCxkJBDGbwTXKUXYw0iWWX5WQpcwnn0LYTXj8X1c4KX29N5gm
LVs=
-----END CERTIFICATE REQUEST-----


A copy of this CSR has been saved in the file
/etc/pki/tls/certs/www.example.com.2.csr

Press return when ready to continue
```

그림 23.18. Begin certificate request

Pressing the return key, will display the next screen from which you can enable or disable the encryption of the private key. Use the spacebar to enable or disable this. When enabled, a [*] character will be displayed. On selecting your preferred option, select Next to proceed to the next step.

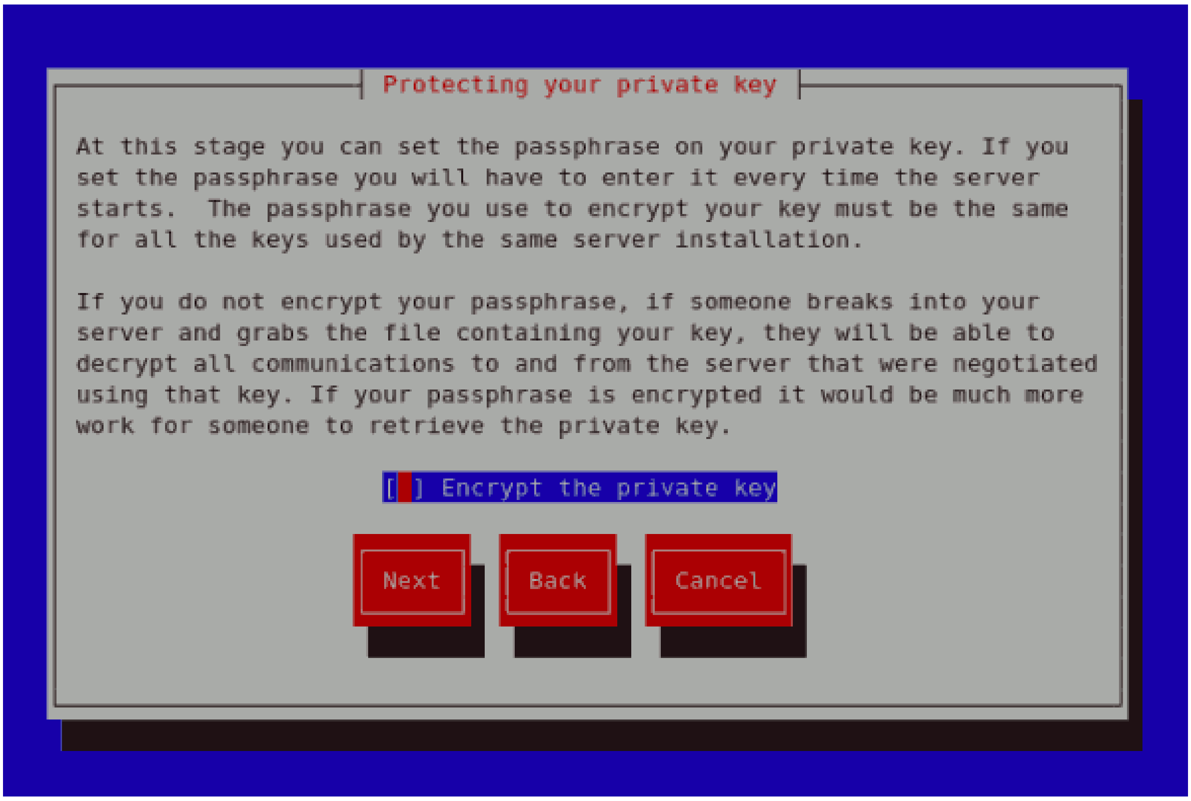

```
Protecting your private key

At this stage you can set the passphrase on your private key. If you
set the passphrase you will have to enter it every time the server
starts.  The passphrase you use to encrypt your key must be the same
for all the keys used by the same server installation.

If you do not encrypt your passphrase, if someone breaks into your
server and grabs the file containing your key, they will be able to
decrypt all communications to and from the server that were negotiated
using that key. If your passphrase is encrypted it would be much more
work for someone to retrieve the private key.

[ ] Encrypt the private key

Next   Back   Cancel
```

그림 23.19. Protecting your private key

The next screen allows you to set your key passphrase. Please do not lose this passphrase as you will not be able to run the server without it. You will need to regenerate a new private or public key pair and request a new certificate from your CA as indicated. For security, the passphrase is not displayed as you type. On typing your preferred passphrase, select Next to go back to your terminal.

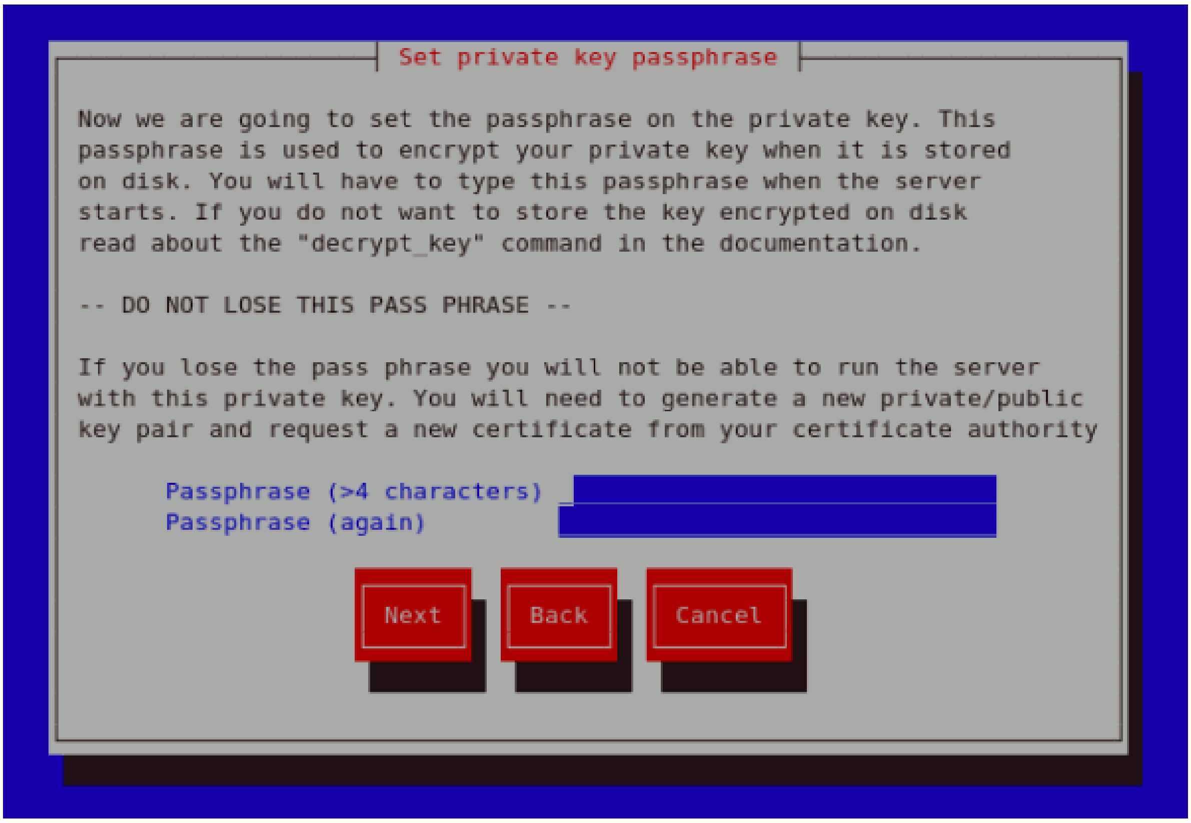

```
Set private key passphrase

Now we are going to set the passphrase on the private key. This
passphrase is used to encrypt your private key when it is stored
on disk. You will have to type this passphrase when the server
starts. If you do not want to store the key encrypted on disk
read about the "decrypt_key" command in the documentation.

-- DO NOT LOSE THIS PASS PHRASE --

If you lose the pass phrase you will not be able to run the server
with this private key. You will need to generate a new private/public
key pair and request a new certificate from your certificate authority

      Passphrase (>4 characters) _
      Passphrase (again)
```

그림 23.20. Set passphrase

If you attempt to run genkey www.example.com on a server that already has an existing key pair for the particular hostname, an error message will be displayed as illustrated below. You need to delete your existing key file as indicated to generate a new key pair.

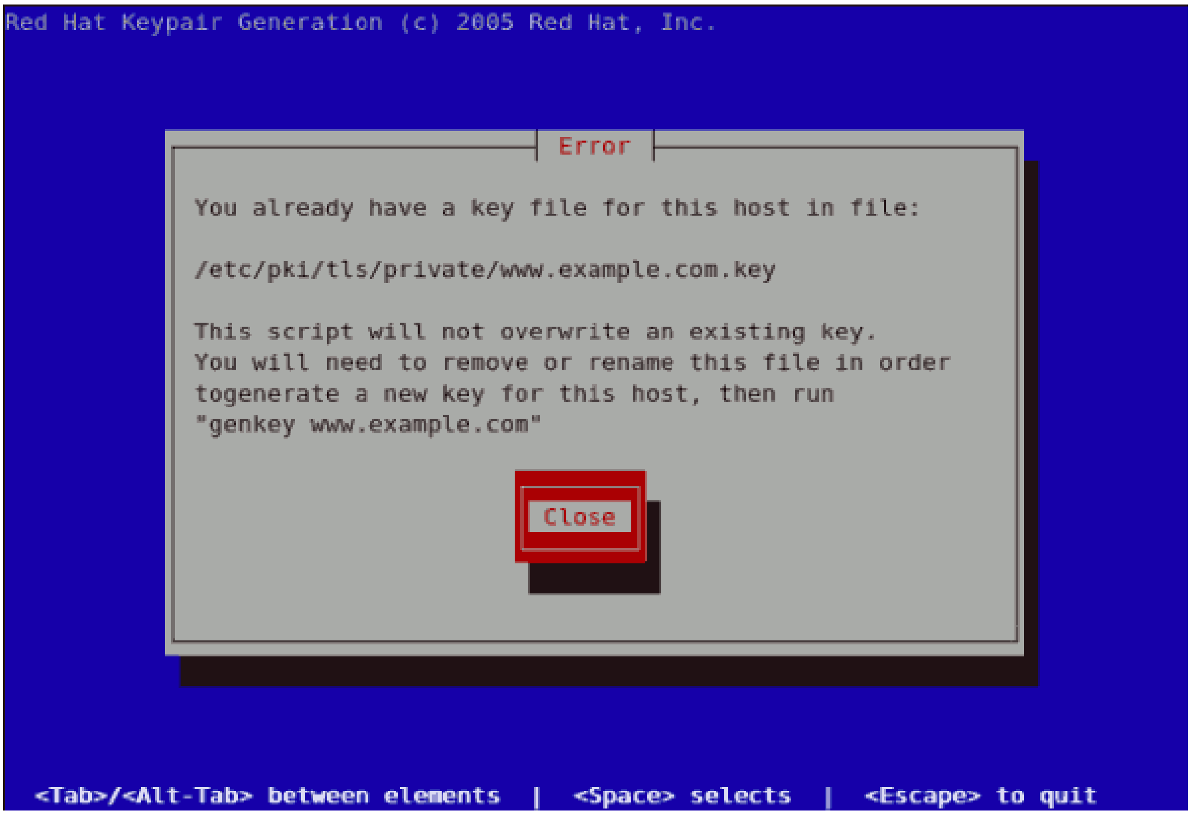

```
Red Hat Keypair Generation (c) 2005 Red Hat, Inc.

Error

You already have a key file for this host in file:

/etc/pki/tls/private/www.example.com.key

This script will not overwrite an existing key.
You will need to remove or rename this file in order
togenerate a new key for this host, then run
"genkey www.example.com"
```

Close
<Tab>/<Alt-Tab> between elements | <Space> selects | <Escape> to quit

그림 23.21. genkey error

- http://httpd.apache.org/docs/2.2/ssl/
- http://httpd.apache.org/docs/2.2/vhosts/

23.8.6. How to configure the server to use the new key

The steps to configure the Apache HTTP Server to use the new key are:

- Obtain the signed certificate from the CA after submitting the CSR.
- Copy the certificate to the path, for example /etc/pki/tls/certs/www.example.com.crt
- Edit /etc/httpd/conf.d/ssl.conf. Change the SSLCertificateFile and SSLCertificateKey lines to be.

  ```
  SSLCertificateFile /etc/pki/tls/certs/www.example.com.crt
  SSLCertificateKeyFile /etc/pki/tls/private/www.example.com.key
  ```

  Note that the "www.example.com" part should match the argument passed on the genkey command.

23.9. Additional Resources

To learn more about the Apache HTTP Server, refer to the following resources.

23.9.1. Useful Websites

- http://httpd.apache.org/ — The official website for the Apache HTTP Server with documentation on all the directives and default modules.
- http://www.modssl.org/ — The official website for mod_ssl.
- http://www.apacheweek.com/ — A comprehensive online weekly newsletter about all things Apache.

FTP

File Transfer Protocol (FTP) is one of the oldest and most commonly used protocols found on the Internet today. Its purpose is to reliably transfer files between computer hosts on a network without requiring the user to log directly into the remote host or have knowledge of how to use the remote system. It allows users to access files on remote systems using a standard set of simple commands.

This chapter outlines the basics of the FTP protocol, as well as configuration options for the primary FTP server shipped with Red Hat Enterprise Linux, vsftpd.

24.1. The File Transfer Protocol

However, because FTP is so prevalent on the Internet, it is often required to share files to the public. System administrators, therefore, should be aware of the FTP protocol's unique characteristics.

24.1.1. Multiple Ports, Multiple Modes

Unlike most protocols used on the Internet, FTP requires multiple network ports to work properly. When an FTP client application initiates a connection to an FTP server, it opens port 21 on the server — known as the command port. This port is used to issue all commands to the server. Any data requested from the server is returned to the client via a data port. The port number for data connections, and the way in which data connections are initialized, vary depending upon whether the client requests the data in active or passive mode.

The following defines these modes:

active mode

Active mode is the original method used by the FTP protocol for transferring data to the client application. When an active mode data transfer is initiated by the FTP client, the server opens a connection from port 20 on the server to the IP address and a random, unprivileged port (greater than 1024) specified by the client. This arrangement means that the client machine must be allowed to accept connections over any port above 1024. With the growth of insecure networks, such as the Internet, the use of firewalls to protect client machines is now prevalent. Because these client-side firewalls often deny incoming connections from active mode FTP servers, passive mode was devised.

passive mode

Passive mode, like active mode, is initiated by the FTP client application. When requesting data from the server, the FTP client indicates it wants to access the data in passive mode and the server provides the IP address and a random, unprivileged port (greater than 1024) on the server. The client then connects to that port on the server to download the requested information.

While passive mode resolves issues for client-side firewall interference with data connections, it can complicate administration of the server-side firewall. You can reduce the number of open ports on a server by limiting the range of unprivileged ports on the FTP server. This also simplifies the process of configuring firewall rules for the server. Refer to 24.5.8절. "Network Options" for more about limiting passive ports.

24.2. FTP Servers

Red Hat Enterprise Linux ships with two different FTP servers:

- Red Hat Content Accelerator — A kernel-based Web server that delivers high performance Web server and FTP services. Since speed as its primary design goal, it has limited functionality and runs only as an anonymous FTP server. For more information about configuring and administering Red Hat Content Accelerator, consult the documentation available online at http://www.redhat.com/docs/manuals/tux/.

- vsftpd — A fast, secure FTP daemon which is the preferred FTP server for Red Hat Enterprise Linux. The remainder of this chapter focuses on vsftpd.

24.2.1. vsftpd

The Very Secure FTP Daemon (vsftpd) is designed from the ground up to be fast, stable, and, most importantly, secure. vsftpd is the only stand-alone FTP server distributed with Red Hat Enterprise Linux, due to its ability to handle large numbers of connections efficiently and securely.

The security model used by vsftpd has three primary aspects:

- Strong separation of privileged and non-privileged processes — Separate processes handle different tasks, and each of these processes run with the minimal privileges required for the task.

- Tasks requiring elevated privileges are handled by processes with the minimal privilege necessary — By leveraging compatibilities found in the libcap library, tasks that usually require full root privileges can be executed more safely from a less privileged process.

- Most processes run in a chroot jail — Whenever possible, processes are change-rooted to the directory being shared; this directory is then considered a chroot jail. For example, if the directory /var/ftp/ is the primary shared directory, vsftpd reassigns /var/ftp/ to the new root directory, known as /. This disallows any potential malicious hacker activities for any directories not contained below the new root directory.

Use of these security practices has the following effect on how vsftpd deals with requests:

- The parent process runs with the least privileges required — The parent process dynamically calculates the level of privileges it requires to minimize the level of risk. Child processes handle direct interaction with the FTP clients and run with as close to no privileges as possible.

- All operations requiring elevated privileges are handled by a small parent process — Much like the Apache HTTP Server, vsftpd launches unprivileged child processes to handle incoming connections. This allows the privileged, parent process to be as small as possible and handle relatively few tasks.

- All requests from unprivileged child processes are distrusted by the parent process — Communication with child processes are received over a socket, and the validity of any information from child processes is checked before being acted on.

- Most interaction with FTP clients is handled by unprivileged child processes in a chroot jail — Because these child processes are unprivileged and only have access to the directory being shared, any crashed processes only allows the attacker access to the shared files.

24.3. Files Installed with vsftpd

The vsftpd RPM installs the daemon (/usr/sbin/vsftpd), its configuration and related files, as well as FTP directories onto the system. The following lists the files and directories related to vsftpd configuration:

- /etc/rc.d/init.d/vsftpd — The initialization script (initscript) used by the /sbin/service command to start, stop, or reload vsftpd. Refer to 24.4절. "Starting and Stopping vsftpd" for more information about using this script.

- /etc/pam.d/vsftpd — The Pluggable Authentication Modules (PAM) configuration file for vsftpd. This file specifies the requirements a user must meet to login to the FTP server. For more information, refer to 46.4절. "PAM (Pluggable Authentication Modules)" .

- /etc/vsftpd/vsftpd.conf — The configuration file for vsftpd. Refer to 24.5절. "vsftpd Configuration Options" for a list of important options contained within this file.

- /etc/vsftpd.ftpusers — A list of users not allowed to log into vsftpd. By default, this list includes the root, bin, and daemon users, among others.

- /etc/vsftpd.user_list — This file can be configured to either deny or allow access to the users listed, depending on whether the userlist_deny directive is set to YES (default) or NO in /etc/vsftpd/vsftpd.conf. If /etc/vsftpd.user_list is used to grant access to users, the usernames listed must not appear in /etc/vsftpd.ftpusers.

- /var/ftp/ — The directory containing files served by vsftpd. It also contains the /var/ftp/pub/ directory for anonymous users. Both directories are world-readable, but writable only by the root user.

24.4. Starting and Stopping vsftpd

The vsftpd RPM installs the /etc/rc.d/init.d/vsftpd script, which can be accessed using the /sbin/service command.

To start the server, as root type:

```
service vsftpd start
```

To stop the server, as root type:

```
service vsftpd stop
```

The restart option is a shorthand way of stopping and then starting vsftpd. This is the most efficient way to make configuration changes take effect after editing the configuration file for vsftpd.

To restart the server, as root type:

```
service vsftpd restart
```

The condrestart (conditional restart) option only starts vsftpd if it is currently running. This option is useful for scripts, because it does not start the daemon if it is not running.

To conditionally restart the server, as root type:

```
service vsftpd condrestart
```

By default, the vsftpd service does not start automatically at boot time. To configure the vsftpd service to start at boot time, use an initscript utility, such as /sbin/chkconfig, /usr/sbin/ntsysv, or the Services Configuration Tool program. Refer to 17장. 서비스로의 접근 통제 for more information regarding these tools.

24.4.1. Starting Multiple Copies of vsftpd

Sometimes one computer is used to serve multiple FTP domains. This is a technique called multihoming. One way to multihome using vsftpd is by running multiple copies of the daemon, each with its own configuration file.

To do this, first assign all relevant IP addresses to network devices or alias network devices on the system. Refer to 16장. 네트워크 설정 for more information about configuring network devices and device aliases. Additional information can be found about network configuration scripts in 15장. 네트워크 인터페이스.

Next, the DNS server for the FTP domains must be configured to reference the correct machine. For information about BIND and its configuration files, refer to 18장. Berkeley Internet Name Domain (BIND).

For vsftpd to answer requests on different IP addresses, multiple copies of the daemon must be running. The first copy must be run using the vsftpd initscripts, as outlined in 24.4절. "Starting and Stopping vsftpd" . This copy uses the standard configuration file, /etc/vsftpd/vsftpd.conf.

Each additional FTP site must have a configuration file with a unique name in the /etc/vsftpd/ directory, such as /etc/vsftpd/vsftpd-site-2.conf. Each configuration file must be readable and writable only by root. Within each configuration file for each FTP server listening on an IPv4 network, the following directive must be unique:

```
listen_address=N.N.N.N
```

Replace N.N.N.N with the unique IP address for the FTP site being served. If the site is using IPv6, use the listen_address6 directive instead.

Once each additional server has a configuration file, the vsftpd daemon must be launched from a root shell prompt using the following command:

```
vsftpd /etc/vsftpd/<configuration-file> [amp     ]
```

In the above command, replace <configuration-file> with the unique name for the server's configuration file, such as /etc/vsftpd/vsftpd-site-2.conf.

Other directives to consider altering on a per-server basis are:

- anon_root
- local_root
- vsftpd_log_file
- xferlog_file

For a detailed list of directives available within vsftpd's configuration file, refer to 24.5절. "vsftpd Configuration Options" .

To configure any additional servers to start automatically at boot time, add the above command to the end of the /etc/rc.local file.

24.5. vsftpd Configuration Options

Although vsftpd may not offer the level of customization other widely available FTP servers have, it offers enough options to fill most administrator's needs. The fact that it is not overly feature-laden limits configuration and programmatic errors.

All configuration of vsftpd is handled by its configuration file, /etc/vsftpd/vsftpd.conf. Each directive is on its own line within the file and follows the following format:

```
<directive>=<value>
```

For each directive, replace <directive> with a valid directive and <value> with a valid value.

Important

There must not be any spaces between the <directive>, equal symbol, and the <value> in a directive.

Comment lines must be preceded by a hash mark (#) and are ignored by the daemon.

For a complete list of all directives available, refer to the man page for vsftpd.conf.

Important

For an overview of ways to secure vsftpd, refer to 46.2절. "서버 보안" .

The following is a list of some of the more important directives within /etc/vsftpd/vsftpd.conf. All directives not explicitly found within vsftpd's configuration file are set to their default value.

24.5.1. Daemon Options

The following is a list of directives which control the overall behavior of the vsftpd daemon.

- listen — When enabled, vsftpd runs in stand-alone mode. Red Hat Enterprise Linux sets this value to YES. This directive cannot be used in conjunction with the listen_ipv6 directive.

  The default value is NO.

- listen_ipv6 — When enabled, vsftpd runs in stand-alone mode, but listens only to IPv6 sockets. This directive cannot be used in conjunction with the listen directive.

  The default value is NO.

- session_support — When enabled, vsftpd attempts to maintain login sessions for each user through Pluggable Authentication Modules (PAM). Refer to 46.4절. "PAM (Pluggable Authentication Modules)" for more information. If session logging is not necessary, disabling this option allows vsftpd to run with less processes and lower privileges.

  The default value is YES.

24.5.2. Log In Options and Access Controls

The following is a list of directives which control the login behavior and access control mechanisms.

- anonymous_enable — When enabled, anonymous users are allowed to log in. The usernames anonymous and ftp are accepted.

  The default value is YES.

  Refer to 24.5.3절. "Anonymous User Options" for a list of directives affecting anonymous users.

- banned_email_file — If the deny_email_enable directive is set to YES, this directive specifies the file containing a list of anonymous email passwords which are not permitted access to the server.

  The default value is /etc/vsftpd.banned_emails.

- banner_file — Specifies the file containing text displayed when a connection is established to the server. This option overrides any text specified in the ftpd_banner directive.

  There is no default value for this directive.

- cmds_allowed — Specifies a comma-delimited list of FTP commands allowed by the server. All other commands are rejected.

  There is no default value for this directive.

- deny_email_enable — When enabled, any anonymous user utilizing email passwords specified in the /etc/vsftpd.banned_emails are denied access to the server. The name of the file referenced by this directive can be specified using the banned_email_file directive.

  The default value is NO.

- ftpd_banner — When enabled, the string specified within this directive is displayed when a connection is established to the server. This option can be overridden by the banner_file directive.

  By default vsftpd displays its standard banner.

- local_enable — When enabled, local users are allowed to log into the system.

  The default value is YES.

  Refer to 24.5.4절. "Local User Options" for a list of directives affecting local users.

- pam_service_name — Specifies the PAM service name for vsftpd.

  The default value is ftp. On Red Hat Enterprise Linux 5.8, this option is set to vsftpd in the configuration file.

- tcp_wrappers — When enabled, TCP wrappers are used to grant access to the server. If the FTP server is configured on multiple IP addresses, the VSFTPD_LOAD_CONF option can be used to load different configuration files based on the IP address being requested by the client.

  The default value is NO. On Red Hat Enterprise Linux 5.8, this option is set to YES in the configuration file.

  Refer to 46.5절. "TCP Wrappers and xinetd" for more information about TCP wrappers.

- userlist_deny — When used in conjunction with the userlist_enable directive and set to NO, all local users are denied access unless the username is listed in the file specified by the userlist_file directive. Because access is denied before the client is asked for a password, setting this directive to NO prevents local users from submitting unencrypted passwords over the network.

  The default value is YES.

- userlist_enable — When enabled, the users listed in the file specified by the userlist_file directive are denied access. Because access is denied before the client is asked for a password, users are prevented from submitting unencrypted passwords over the network.

  The default value is NO. On Red Hat Enterprise Linux 5.8, this option is set to YES in the configuration file.

- userlist_file — Specifies the file referenced by vsftpd when the userlist_enable directive is enabled.

  The default value is /etc/vsftpd.user_list and is created during installation.

24.5.3. Anonymous User Options

The following lists directives which control anonymous user access to the server. To use these options, the anonymous_enable directive must be set to YES.

- anon_mkdir_write_enable — When enabled in conjunction with the write_enable directive, anonymous users are allowed to create new directories within a parent directory which has write permissions.

  The default value is NO.

- anon_root — Specifies the directory vsftpd changes to after an anonymous user logs in.

  There is no default value for this directive.

- anon_upload_enable — When enabled in conjunction with the write_enable directive, anonymous users are allowed to upload files within a parent directory which has write permissions.

  The default value is NO.

- anon_world_readable_only — When enabled, anonymous users are only allowed to download world-readable files.

  The default value is YES.

- ftp_username — Specifies the local user account (listed in /etc/passwd) used for the anonymous FTP user. The home directory specified in /etc/passwd for the user is the root directory of the anonymous FTP user.

  The default value is ftp.

- no_anon_password — When enabled, the anonymous user is not asked for a password.

  The default value is NO.

- secure_email_list_enable — When enabled, only a specified list of email passwords for anonymous logins are accepted. This is a convenient way to offer limited security to public content without the need for virtual users.

Anonymous logins are prevented unless the password provided is listed in /etc/vsftpd.email_passwords. The file format is one password per line, with no trailing white spaces.

The default value is NO.

24.5.4. Local User Options

The following lists directives which characterize the way local users access the server. To use these options, the local_enable directive must be set to YES.

- chmod_enable — When enabled, the FTP command SITE CHMOD is allowed for local users. This command allows the users to change the permissions on files.

  The default value is YES.

- chroot_list_enable — When enabled, the local users listed in the file specified in the chroot_list_file directive are placed in a chroot jail upon log in.

  If enabled in conjunction with the chroot_local_user directive, the local users listed in the file specified in the chroot_list_file directive are not placed in a chroot jail upon log in.

  The default value is NO.

- chroot_list_file — Specifies the file containing a list of local users referenced when the chroot_list_enable directive is set to YES.

  The default value is /etc/vsftpd.chroot_list.

- chroot_local_user — When enabled, local users are change-rooted to their home directories after logging in.

  The default value is NO.

  Warning

  Enabling chroot_local_user opens up a number of security issues, especially for users with upload privileges. For this reason, it is not recommended.

- guest_enable — When enabled, all non-anonymous users are logged in as the user guest, which is the local user specified in the guest_username directive.

  The default value is NO.

- guest_username — Specifies the username the guest user is mapped to.

  The default value is ftp.

- local_root — Specifies the directory vsftpd changes to after a local user logs in.

  There is no default value for this directive.

- local_umask — Specifies the umask value for file creation. Note that the default value is in octal form (a numerical system with a base of eight), which includes a "0" prefix. Otherwise the value is treated as a base-10 integer.

  The default value is 022.

- passwd_chroot_enable — When enabled in conjunction with the chroot_local_user directive, vsftpd change-roots local users based on the occurrence of the /./ in the home directory field within /etc/passwd.

  The default value is NO.

- user_config_dir — Specifies the path to a directory containing configuration files bearing the name of local system users that contain specific setting for that user. Any directive in the user's configuration file overrides those found in /etc/vsftpd/vsftpd.conf.

  There is no default value for this directive.

24.5.5. Directory Options

The following lists directives which affect directories.

- dirlist_enable — When enabled, users are allowed to view directory lists.

  The default value is YES.

- dirmessage_enable — When enabled, a message is displayed whenever a user enters a directory with a message file. This message resides within the current directory. The name of this file is specified in the message_file directive and is .message by default.

  The default value is NO. On Red Hat Enterprise Linux 5.8, this option is set to YES in the configuration file.

- force_dot_files — When enabled, files beginning with a dot (.) are listed in directory listings, with the exception of the . and .. files.

  The default value is NO.

- hide_ids — When enabled, all directory listings show ftp as the user and group for each file.

  The default value is NO.

- message_file — Specifies the name of the message file when using the dirmessage_enable directive.

  The default value is .message.

- text_userdb_names — When enabled, text usernames and group names are used in place of UID and GID entries. Enabling this option may slow performance of the server.

  The default value is NO.

- use_localtime — When enabled, directory listings reveal the local time for the computer instead of GMT.

  The default value is NO.

24.5.6. File Transfer Options

The following lists directives which affect directories.

- download_enable — When enabled, file downloads are permitted.

  The default value is YES.

- chown_uploads — When enabled, all files uploaded by anonymous users are owned by the user specified in the chown_username directive.

  The default value is NO.

- chown_username — Specifies the ownership of anonymously uploaded files if the chown_uploads directive is enabled.

  The default value is root.

- write_enable — When enabled, FTP commands which can change the file system are allowed, such as DELE, RNFR, and STOR.

  The default value is YES.

24.5.7. Logging Options

The following lists directives which affect vsftpd's logging behavior.

- dual_log_enable — When enabled in conjunction with xferlog_enable, vsftpd writes two files simultaneously: a wu-ftpd-compatible log to the file specified in the xferlog_file directive (/var/log/xferlog by default) and a standard vsftpd log file specified in the vsftpd_log_file directive (/var/log/vsftpd.log by default).

  The default value is NO.

- log_ftp_protocol — When enabled in conjunction with xferlog_enable and with xferlog_std_format set to NO, all FTP commands and responses are logged. This directive is useful for debugging.

  The default value is NO.

- syslog_enable — When enabled in conjunction with xferlog_enable, all logging normally written to the standard vsftpd log file specified in the vsftpd_log_file directive (/var/log/vsftpd.log by default) is sent to the system logger instead under the FTPD facility.

  The default value is NO.

- vsftpd_log_file — Specifies the vsftpd log file. For this file to be used, xferlog_enable must be enabled and xferlog_std_format must either be set to NO or, if xferlog_std_format is set to YES, dual_log_enable must be enabled. It is important to note that if syslog_enable is set to YES, the system log is used instead of the file specified in this directive.

  The default value is /var/log/vsftpd.log.

- xferlog_enable — When enabled, vsftpd logs connections (vsftpd format only) and file transfer information to the log file specified in the vsftpd_log_file directive (/var/log/vsftpd.log by default). If xferlog_std_format is set to YES, file transfer information is logged but connections are not, and the log file specified in xferlog_file (/var/log/xferlog by default) is used instead. It is important to note that both log files and log formats are used if dual_log_enable is set to YES.

The default value is NO. On Red Hat Enterprise Linux 5.8, this option is set to YES in the configuration file.

- xferlog_file — Specifies the wu-ftpd-compatible log file. For this file to be used, xferlog_enable must be enabled and xferlog_std_format must be set to YES. It is also used if dual_log_enable is set to YES.

  The default value is /var/log/xferlog.

- xferlog_std_format — When enabled in conjunction with xferlog_enable, only a wu-ftpd-compatible file transfer log is written to the file specified in the xferlog_file directive (/var/log/xferlog by default). It is important to note that this file only logs file transfers and does not log connections to the server.

  The default value is NO. On Red Hat Enterprise Linux 5.8, this option is set to YES in the configuration file.

Important

To maintain compatibility with log files written by the older wu-ftpd FTP server, the xferlog_std_format directive is set to YES under Red Hat Enterprise Linux. However, this setting means that connections to the server are not logged.

To both log connections in vsftpd format and maintain a wu-ftpd-compatible file transfer log, set dual_log_enable to YES.

If maintaining a wu-ftpd-compatible file transfer log is not important, either set xferlog_std_format to NO, comment the line with a hash mark (#), or delete the line entirely.

24.5.8. Network Options

The following lists directives which affect how vsftpd interacts with the network.

- accept_timeout — Specifies the amount of time for a client using passive mode to establish a connection.

  The default value is 60.

- anon_max_rate — Specifies the maximum data transfer rate for anonymous users in bytes per second.

  The default value is 0, which does not limit the transfer rate.

- connect_from_port_20 When enabled, vsftpd runs with enough privileges to open port 20 on the server during active mode data transfers. Disabling this option allows vsftpd to run with less privileges, but may be incompatible with some FTP clients.

  The default value is NO. On Red Hat Enterprise Linux 5.8, this option is set to YES in the configuration file.

- connect_timeout — Specifies the maximum amount of time a client using active mode has to respond to a data connection, in seconds.

  The default value is 60.

- data_connection_timeout — Specifies maximum amount of time data transfers are allowed to stall, in seconds. Once triggered, the connection to the remote client is closed.

  The default value is 300.

- ftp_data_port — Specifies the port used for active data connections when connect_from_port_20 is set to YES.

  The default value is 20.

- idle_session_timeout — Specifies the maximum amount of time between commands from a remote client. Once triggered, the connection to the remote client is closed.

  The default value is 300.

- listen_address — Specifies the IP address on which vsftpd listens for network connections.

  There is no default value for this directive.

  > Tip
  >
  > If running multiple copies of vsftpd serving different IP addresses, the configuration file for each copy of the vsftpd daemon must have a different value for this directive. Refer to 24.4.1절. "Starting Multiple Copies of vsftpd" for more information about multihomed FTP servers.

- listen_address6 — Specifies the IPv6 address on which vsftpd listens for network connections when listen_ipv6 is set to YES.

  There is no default value for this directive.

  > Tip
  >
  > If running multiple copies of vsftpd serving different IP addresses, the configuration file for each copy of the vsftpd daemon must have a different value for this directive. Refer to 24.4.1절. "Starting Multiple Copies of vsftpd" for more information about multihomed FTP servers.

- listen_port — Specifies the port on which vsftpd listens for network connections.

  The default value is 21.

- local_max_rate — Specifies the maximum rate data is transferred for local users logged into the server in bytes per second.

  The default value is 0, which does not limit the transfer rate.

- max_clients — Specifies the maximum number of simultaneous clients allowed to connect to the server when it is running in standalone mode. Any additional client connections would result in an error message.

  The default value is 0, which does not limit connections.

- max_per_ip — Specifies the maximum of clients allowed to connected from the same source IP address.

  The default value is 0, which does not limit connections.

- pasv_address — Specifies the IP address for the public facing IP address of the server for servers behind Network Address Translation (NAT) firewalls. This enables vsftpd to hand out the correct return address for passive mode connections.

  There is no default value for this directive.

- pasv_enable — When enabled, passive mode connects are allowed.

  The default value is YES.

- pasv_max_port — Specifies the highest possible port sent to the FTP clients for passive mode connections. This setting is used to limit the port range so that firewall rules are easier to create.

  The default value is 0, which does not limit the highest passive port range. The value must not exceed 65535.

- pasv_min_port — Specifies the lowest possible port sent to the FTP clients for passive mode connections. This setting is used to limit the port range so that firewall rules are easier to create.

  The default value is 0, which does not limit the lowest passive port range. The value must not be lower 1024.

- pasv_promiscuous — When enabled, data connections are not checked to make sure they are originating from the same IP address. This setting is only useful for certain types of tunneling.

  The default value is NO.

  Caution

  Do not enable this option unless absolutely necessary as it disables an important security feature which verifies that passive mode connections originate from the same IP address as the control connection that initiates the data transfer.

- port_enable — When enabled, active mode connects are allowed.

  The default value is YES.

24.6. Additional Resources

For more information about vsftpd, refer to the following resources.

24.6.1. Installed Documentation

- The /usr/share/doc/vsftpd-<version-number>/ directory — Replace <version-number> with the installed version of the vsftpd package. This directory contains a README with basic information about the software. The TUNING file contains basic performance tuning tips and the SECURITY/ directory contains information about the security model employed by vsftpd.

- vsftpd related man pages — There are a number of man pages for the daemon and configuration files. The following lists some of the more important man pages.

  Server Applications
    - man vsftpd — Describes available command line options for vsftpd.

  Configuration Files
    - man vsftpd.conf — Contains a detailed list of options available within the configuration file for vsftpd.

    - man 5 hosts_access — Describes the format and options available within the TCP wrappers configuration files: hosts.allow and hosts.deny.

24.6.2. Useful Websites

- http://vsftpd.beasts.org/ — The vsftpd project page is a great place to locate the latest documentation and to contact the author of the software.

- http://slacksite.com/other/ftp.html — This website provides a concise explanation of the differences between active and passive mode FTP.

- http://www.ietf.org/rfc/rfc0959.txt — The original Request for Comments (RFC) of the FTP protocol from the IETF.

Email

The birth of electronic mail (email) occurred in the early 1960s. The mailbox was a file in a user's home directory that was readable only by that user. Primitive mail applications appended new text messages to the bottom of the file, making the user wade through the constantly growing file to find any particular message. This system was only capable of sending messages to users on the same system.

The first network transfer of an electronic mail message file took place in 1971 when a computer engineer named Ray Tomlinson sent a test message between two machines via ARPANET — the precursor to the Internet. Communication via email soon became very popular, comprising 75 percent of ARPANET's traffic in less than two years.

Today, email systems based on standardized network protocols have evolved into some of the most widely used services on the Internet. Red Hat Enterprise Linux offers many advanced applications to serve and access email.

This chapter reviews modern email protocols in use today and some of the programs designed to send and receive email.

25.1. Email Protocols

Today, email is delivered using a client/server architecture. An email message is created using a mail client program. This program then sends the message to a server. The server then forwards the message to the recipient's email server, where the message is then supplied to the recipient's email client.

To enable this process, a variety of standard network protocols allow different machines, often running different operating systems and using different email programs, to send and receive email.

The following protocols discussed are the most commonly used in the transfer of email.

25.1.1. Mail Transport Protocols

Mail delivery from a client application to the server, and from an originating server to the destination server, is handled by the Simple Mail Transfer Protocol (SMTP).

25.1.1.1. SMTP

The primary purpose of SMTP is to transfer email between mail servers. However, it is critical for email clients as well. To send email, the client sends the message to an outgoing mail server, which in turn contacts the destination mail server for delivery. For this reason, it is necessary to specify an SMTP server when configuring an email client.

Under Red Hat Enterprise Linux, a user can configure an SMTP server on the local machine to handle mail delivery. However, it is also possible to configure remote SMTP servers for outgoing mail.

One important point to make about the SMTP protocol is that it does not require authentication. This allows anyone on the Internet to send email to anyone else or even to large groups of people. It is this characteristic of SMTP that makes junk email or spam possible. Imposing relay restrictions limits random users on the Internet from sending email through your SMTP server, to other servers on the internet. Servers that do not impose such restrictions are called open relay servers.

By default, Sendmail (/usr/sbin/sendmail) is the default SMTP program under Red Hat Enterprise Linux. However, a simpler mail server application called Postfix (/usr/sbin/postfix) is also available.

25.1.2. Mail Access Protocols

There are two primary protocols used by email client applications to retrieve email from mail servers: the Post Office Protocol (POP) and the Internet Message Access Protocol (IMAP).

25.1.2.1. POP

The default POP server under Red Hat Enterprise Linux is /usr/lib/cyrus-imapd/pop3d and is provided by the cyrus-imapd package. When using a POP server, email messages are downloaded by email client applications. By default, most POP email clients are automatically configured to delete the message on the email server after it has been successfully transferred, however this setting usually can be changed.

POP is fully compatible with important Internet messaging standards, such as Multipurpose Internet Mail Extensions (MIME), which allow for email attachments.

POP works best for users who have one system on which to read email. It also works well for users who do not have a persistent connection to the Internet or the network containing the mail server. Unfortunately for those with slow network connections, POP requires client programs upon authentication to download the entire content of each message. This can take a long time if any messages have large attachments.

The most current version of the standard POP protocol is POP3.

There are, however, a variety of lesser-used POP protocol variants:

- APOP — POP3 with MDS authentication. An encoded hash of the user's password is sent from the email client to the server rather then sending an unencrypted password.

- KPOP — POP3 with Kerberos authentication. Refer to 46.6절. "Kerberos" for more information.

- RPOP — POP3 with RPOP authentication. This uses a per-user ID, similar to a password, to authenticate POP requests. However, this ID is not encrypted, so RPOP is no more secure than standard POP.

For added security, it is possible to use Secure Socket Layer (SSL) encryption for client authentication and data transfer sessions. This can be enabled by using the ipop3s service or by using the /usr/sbin/stunnel program. Refer to 25.6.1절. "Securing Communication" for more information.

25.1.2.2. IMAP

The default IMAP server under Red Hat Enterprise Linux is /usr/lib/cyrus-imapd/imapd and is provided by the cyrus-imapd package. When using an IMAP mail server, email messages remain on the server where users can read or delete them. IMAP also allows client applications to create, rename, or delete mail directories on the server to organize and store email.

IMAP is particularly useful for those who access their email using multiple machines. The protocol is also convenient for users connecting to the mail server via a slow connection, because only the email header information is downloaded for messages until opened, saving bandwidth. The user also has the ability to delete messages without viewing or downloading them.

For convenience, IMAP client applications are capable of caching copies of messages locally, so the user can browse previously read messages when not directly connected to the IMAP server.

IMAP, like POP, is fully compatible with important Internet messaging standards, such as MIME, which allow for email attachments.

For added security, it is possible to use SSL encryption for client authentication and data transfer sessions. This can be enabled by using the imaps service, or by using the /usr/sbin/stunnel program. Refer to 25.6.1절. "Securing Communication" for more information.

Other free, as well as commercial, IMAP clients and servers are available, many of which extend the IMAP protocol and provide additional functionality. A comprehensive list can be found online at http://www.imap.org/products/longlist.htm.

25.1.2.3. Dovecot

The imap-login and pop3-login daemons which implement the IMAP and POP3 protocols are included in the dovecot package. The use of IMAP and POP is configured through dovecot; by default dovecot runs only IMAP. To configure dovecot to use POP:

1. Edit /etc/dovecot.conf to have the line:

   ```
   protocols = imap imaps pop3 pop3s
   ```

2. Make that change operational for the current session by running the command:

   ```
   service dovecot restart
   ```

3. Make that change operational after the next reboot by running the command:

   ```
   chkconfig dovecot on
   ```

   Please note that dovecot only reports that it started the IMAP server, but also starts the POP3 server.

Unlike SMTP, both of these protocols require connecting clients to authenticate using a username and password. By default, passwords for both protocols are passed over the network unencrypted.

To configure SSL on dovecot:

- Edit the dovecot configuration file /etc/pki/dovecot/dovecot-openssl.conf as you prefer. However in a typical installation, this file does not require modification.

- Rename, move or delete the files /etc/pki/dovecot/certs/dovecot.pem and /etc/pki/dovecot/private/dovecot.pem.

- Execute the /usr/share/doc/dovecot-1.0/examples/mkcert.sh script which creates the dovecot self signed certificates. The certificates are copied in the /etc/pki/dovecot/certs and /etc/pki/dovecot/private directories. To implement the changes, restart dovecot (/sbin/service dovecot restart).

More details on dovecot can be found online at http://www.dovecot.org.

25.2. Email Program Classifications

In general, all email applications fall into at least one of three classifications. Each classification plays a specific role in the process of moving and managing email messages. While most users are only aware of the specific email program they use to receive and send messages, each one is important for ensuring that email arrives at the correct destination.

25.2.1. Mail Transport Agent

A Mail Transport Agent (MTA) transports email messages between hosts using SMTP. A message may involve several MTAs as it moves to its intended destination.

While the delivery of messages between machines may seem rather straightforward, the entire process of deciding if a particular MTA can or should accept a message for delivery is quite complicated. In addition, due to problems from spam, use of a particular MTA is usually restricted by the MTA's configuration or the access configuration for the network on which the MTA resides.

Many modern email client programs can act as an MTA when sending email. However, this action should not be confused with the role of a true MTA. The sole reason email client programs are capable of sending email like an MTA is because the host running the application does not have its own MTA. This is particularly true for email client programs on non-UNIX-based operating systems. However, these client programs only send outbound messages to an MTA they are authorized to use and do not directly deliver the message to the intended recipient's email server.

Since Red Hat Enterprise Linux installs two MTAs, Sendmail and Postfix, email client programs are often not required to act as an MTA. Red Hat Enterprise Linux also includes a special purpose MTA called Fetchmail.

For more information on Sendmail, Postfix, and Fetchmail, refer to 25.3절. "Mail Transport Agents" .

25.2.2. Mail Delivery Agent

A Mail Delivery Agent (MDA) is invoked by the MTA to file incoming email in the proper user's mailbox. In many cases, the MDA is actually a Local Delivery Agent (LDA), such as mail or Procmail.

Any program that actually handles a message for delivery to the point where it can be read by an email client application can be considered an MDA. For this reason, some MTAs (such as Sendmail and Postfix) can fill the role of an MDA when they append new email messages to a local user's mail spool file. In general, MDAs do not transport messages between systems nor do they provide a user interface; MDAs distribute and sort messages on the local machine for an email client application to access.

25.2.3. Mail User Agent

A Mail User Agent (MUA) is synonymous with an email client application. An MUA is a program that, at the very least, allows a user to read and compose email messages. Many MUAs are capable of retrieving messages via the POP or IMAP protocols, setting up mailboxes to store messages, and sending outbound messages to an MTA.

MUAs may be graphical, such as Evolution, or have a very simple, text-based interface, such as mutt.

25.3. Mail Transport Agents

Red Hat Enterprise Linux includes two primary MTAs, Sendmail and Postfix. Sendmail is configured as the default MTA, although it is easy to switch the default MTA to Postfix.

25.3.1. Sendmail

Sendmail's core purpose, like other MTAs, is to safely transfer email among hosts, usually using the SMTP protocol. However, Sendmail is highly configurable, allowing control over almost every aspect of how email is handled, including the protocol used. Many system administrators elect to use Sendmail as their MTA due to its power and scalability.

25.3.1.1. Purpose and Limitations

It is important to be aware of what Sendmail is and what it can do, as opposed to what it is not. In these days of monolithic applications that fulfill multiple roles, Sendmail may seem like the only application needed to run an email server within an organization. Technically, this is true, as Sendmail can spool mail to each users' directory and deliver outbound mail for users. However, most users actually require much more than simple email delivery. Users usually want to interact with their email using an MUA, that uses POP or IMAP, to download their messages to their local machine. Or, they may prefer a Web interface to gain access to their mailbox. These other applications can work in conjunction with Sendmail, but they actually exist for different reasons and can operate separately from one another.

It is beyond the scope of this section to go into all that Sendmail should or could be configured to do. With literally hundreds of different options and rule sets, entire volumes have been dedicated to helping explain everything that can be done and how to fix things that go wrong. Refer to the 25.7 절. "Additional Resources" for a list of Sendmail resources.

This section reviews the files installed with Sendmail by default and reviews basic configuration changes, including how to stop unwanted email (spam) and how to extend Sendmail with the Lightweight Directory Access Protocol (LDAP).

25.3.1.2. The Default Sendmail Installation

The Sendmail executable is /usr/sbin/sendmail.

Sendmail's lengthy and detailed configuration file is /etc/mail/sendmail.cf. Avoid editing the sendmail.cf file directly. To make configuration changes to Sendmail, edit the /etc/mail/sendmail.mc file, back up the original /etc/mail/sendmail.cf, and use the following alternatives to generate a new configuration file:

- Use the included makefile in /etc/mail (make all -C /etc/mail) to create a new /etc/mail/sendmail.cf configuration file. All other generated files in /etc/mail (db files) will be regenerated if needed. The old makemap commands are still usable. The make command will automatically be used by service sendmail start | restart | reload if the make package is installed.

- Alternatively you may use the included m4 macro processor to create a new /etc/mail/sendmail.cf.

More information on configuring Sendmail can be found in 25.3.1.3절. "Common Sendmail Configuration Changes" .

Various Sendmail configuration files are installed in the /etc/mail/ directory including:

- access — Specifies which systems can use Sendmail for outbound email.
- domaintable — Specifies domain name mapping.
- local-host-names — Specifies aliases for the host.
- mailertable — Specifies instructions that override routing for particular domains.
- virtusertable — Specifies a domain-specific form of aliasing, allowing multiple virtual domains to be hosted on one machine.

Several of the configuration files in /etc/mail/, such as access, domaintable, mailertable and virtusertable, must actually store their information in database files before Sendmail can use any configuration changes. To include any changes made to these configurations in their database files, run the following command:

makemap hash /etc/mail/<name> < /etc/mail/<name>

where <name> is replaced with the name of the configuration file to convert.

For example, to have all emails addressed to the example.com domain delivered to bob@other-example.com , add the following line to the virtusertable file:

```
@example.com        bob@other-example.com
```

To finalize the change, the virtusertable.db file must be updated using the following command as root:

```
makemap hash /etc/mail/virtusertable < /etc/mail/virtusertable
```

This creates an updated virtusertable.db file containing the new configuration.

25.3.1.3. Common Sendmail Configuration Changes

When altering the Sendmail configuration file, it is best not to edit an existing file, but to generate an entirely new /etc/mail/sendmail.cf file.

Caution

Before changing the sendmail.cf file, it is a good idea to create a backup copy.

To add the desired functionality to Sendmail, edit the /etc/mail/sendmail.mc file as the root user. When finished, use the m4 macro processor to generate a new sendmail.cf by executing the following command:

```
m4 /etc/mail/sendmail.mc > /etc/mail/sendmail.cf
```

By default, the m4 macro processor is installed with Sendmail but is part of the m4 package.

After creating a new /etc/mail/sendmail.cf file, restart Sendmail for the changes to take effect. The easiest way to do this is to type the following command:

```
service sendmail restart
```

Important

The default sendmail.cf file does not allow Sendmail to accept network connections from any host other than the local computer. To configure Sendmail as a server for other clients, edit the /etc/mail/sendmail.mc file, and either change the address specified in the Addr= option of the DAEMON_OPTIONS directive from 127.0.0.1 to the IP address of an active network device or comment out the DAEMON_OPTIONS directive all together by placing dnl at the beginning of the line. When finished, regenerate /etc/mail/sendmail.cf by executing the following command:

```
m4 /etc/mail/sendmail.mc > /etc/mail/sendmail.cf
```

The default configuration which ships with Red Hat Enterprise Linux works for most SMTP-only sites. However, it does not work for UUCP (UNIX to UNIX Copy) sites. If using UUCP mail transfers, the /etc/mail/sendmail.mc file must be reconfigured and a new /etc/mail/sendmail.cf must be generated.

Consult the /usr/share/sendmail-cf/README file before editing any files in the directories under the /usr/share/sendmail-cf directory, as they can affect the future configuration of /etc/mail/sendmail.cf files.

25.3.1.4. Masquerading

One common Sendmail configuration is to have a single machine act as a mail gateway for all machines on the network. For instance, a company may want to have a machine called mail.example.com that handles all of their email and assigns a consistent return address to all outgoing mail.

In this situation, the Sendmail server must masquerade the machine names on the company network so that their return address is user@example.com instead of user@host.example.com.

To do this, add the following lines to /etc/mail/sendmail.mc:

```
FEATURE(always_add_domain)dnl
FEATURE(`masquerade_entire_domain')dnl
FEATURE(`masquerade_envelope')dnl
FEATURE(`allmasquerade')dnl
MASQUERADE_AS(`bigcorp.com.')dnl
MASQUERADE_DOMAIN(`bigcorp.com.')dnl
MASQUERADE_AS(bigcorp.com)dnl
```

After generating a new sendmail.cf using m4, this configuration makes all mail from inside the network appear as if it were sent from bigcorp.com.

25.3.1.5. Stopping Spam

Email spam can be defined as unnecessary and unwanted email received by a user who never requested the communication. It is a disruptive, costly, and widespread abuse of Internet communication standards.

Sendmail makes it relatively easy to block new spamming techniques being employed to send junk email. It even blocks many of the more usual spamming methods by default. Main anti-spam features available in sendmail are header checks, relaying denial (default from version 8.9), access database and sender information checks.

For example, forwarding of SMTP messages, also called relaying, has been disabled by default since Sendmail version 8.9. Before this change occurred, Sendmail directed the mail host (x.edu) to accept messages from one party (y.com) and sent them to a different party (z.net). Now, however, Sendmail must be configured to permit any domain to relay mail through the server. To configure relay domains, edit the /etc/mail/relay-domains file and restart Sendmail.

However, many times users are bombarded with spam from other servers throughout the Internet. In these instances, Sendmail's access control features available through the /etc/mail/access file can be used to prevent connections from unwanted hosts. The following example illustrates how this file can be used to both block and specifically allow access to the Sendmail server:

```
badspammer.com          ERROR:550 "Go away and do not spam us anymore"
tux.badspammer.com     OK
10.0                   RELAY
```

This example shows that any email sent from badspammer.com is blocked with a 550 RFC-821 compliant error code, with a message sent back to the spammer. Email sent from the tux.badspammer.com sub-domain, is accepted. The last line shows that any email sent from the 10.0.*.* network can be relayed through the mail server.

Because /etc/mail/access.db is a database, use makemap to activate any changes. Do this using the following command as root:

```
makemap hash /etc/mail/access < /etc/mail/access
```

Message header analysis allows you to reject mail based on header contents. SMTP servers store information about an emails journey in the message header. As the message travels from one MTA to another, each puts in a "Received" header above all the other Received headers. It is however important to note that this information may be altered by spammers.

The above examples only represent a small part of what Sendmail can do in terms of allowing or blocking access. Refer to the /usr/share/sendmail-cf/README for more information and examples.

Since Sendmail calls the Procmail MDA when delivering mail, it is also possible to use a spam filtering program, such as SpamAssassin, to identify and file spam for users. Refer to 25.5.2.6절. "Spam Filters" for more about using SpamAssassin.

25.3.1.6. Using Sendmail with LDAP

Using the Lightweight Directory Access Protocol (LDAP) is a very quick and powerful way to find specific information about a particular user from a much larger group. For example, an LDAP server can be used to look up a particular email address from a common corporate directory by the user's last name. In this kind of implementation, LDAP is largely separate from Sendmail, with LDAP storing the hierarchical user information and Sendmail only being given the result of LDAP queries in pre-addressed email messages.

However, Sendmail supports a much greater integration with LDAP, where it uses LDAP to replace separately maintained files, such as aliases and virtusertables, on different mail servers that work together to support a medium- to enterprise-level organization. In short, LDAP abstracts the mail routing level from Sendmail and its separate configuration files to a powerful LDAP cluster that can be leveraged by many different applications.

The current version of Sendmail contains support for LDAP. To extend the Sendmail server using LDAP, first get an LDAP server, such as OpenLDAP, running and properly configured. Then edit the /etc/mail/sendmail.mc to include the following:

```
LDAPROUTE_DOMAIN('yourdomain.com')dnl
FEATURE('ldap_routing')dnl
```

Note

This is only for a very basic configuration of Sendmail with LDAP. The configuration can differ greatly from this depending on the implementation of LDAP, especially when configuring several Sendmail machines to use a common LDAP server.

Consult /usr/share/sendmail-cf/README for detailed LDAP routing configuration instructions and examples.

Next, recreate the /etc/mail/sendmail.cf file by running m4 and restarting Sendmail. Refer to 25.3.1.3 절. "Common Sendmail Configuration Changes" for instructions.

For more information on LDAP, refer to 26장. Lightweight Directory Access Protocol (LDAP).

25.3.2. Postfix

Originally developed at IBM by security expert and programmer Wietse Venema, Postfix is a Sendmail-compatible MTA that is designed to be secure, fast, and easy to configure.

To improve security, Postfix uses a modular design, where small processes with limited privileges are launched by a master daemon. The smaller, less privileged processes perform very specific tasks related to the various stages of mail delivery and run in a change rooted environment to limit the effects of attacks.

Configuring Postfix to accept network connections from hosts other than the local computer takes only a few minor changes in its configuration file. Yet for those with more complex needs, Postfix provides a variety of configuration options, as well as third party add ons that make it a very versatile and full-featured MTA.

The configuration files for Postfix are human readable and support upward of 250 directives. Unlike Sendmail, no macro processing is required for changes to take effect and the majority of the most commonly used options are described in the heavily commented files.

Important

Before using Postfix, the default MTA must be switched from Sendmail to Postfix.

25.3.2.1. The Default Postfix Installation

The Postfix executable is /usr/sbin/postfix. This daemon launches all related processes needed to handle mail delivery.

Postfix stores its configuration files in the /etc/postfix/ directory. The following is a list of the more commonly used files:

- access — Used for access control, this file specifies which hosts are allowed to connect to Postfix.
- aliases — A configurable list required by the mail protocol.
- main.cf — The global Postfix configuration file. The majority of configuration options are specified in this file.
- master.cf — Specifies how Postfix interacts with various processes to accomplish mail delivery.
- transport — Maps email addresses to relay hosts.

> **Important**
>
> The default /etc/postfix/main.cf file does not allow Postfix to accept network connections from a host other than the local computer. For instructions on configuring Postfix as a server for other clients, refer to 25.3.2.2절. “Basic Postfix Configuration” .

When changing some options within files in the /etc/postfix/ directory, it may be necessary to restart the postfix service for the changes to take effect. The easiest way to do this is to type the following command:

```
service postfix restart
```

25.3.2.2. Basic Postfix Configuration

By default, Postfix does not accept network connections from any host other than the local host. Perform the following steps as root to enable mail delivery for other hosts on the network:

- Edit the /etc/postfix/main.cf file with a text editor, such as vi.
- Uncomment the mydomain line by removing the hash mark (#), and replace domain.tld with the domain the mail server is servicing, such as example.com.
- Uncomment the myorigin = $mydomain line.
- Uncomment the myhostname line, and replace host.domain.tld with the hostname for the machine.
- Uncomment the mydestination = $myhostname, localhost.$mydomain line.
- Uncomment the mynetworks line, and replace 168.100.189.0/28 with a valid network setting for hosts that can connect to the server.
- Uncomment the inet_interfaces = all line.
- Comment the inet_interfaces = localhost line.

- Restart the postfix service.

Once these steps are complete, the host accepts outside emails for delivery.

Postfix has a large assortment of configuration options. One of the best ways to learn how to configure Postfix is to read the comments within /etc/postfix/main.cf. Additional resources including information about LDAP and SpamAssassin integration are available online at http://www.postfix.org/.

25.3.3. Fetchmail

Fetchmail is an MTA which retrieves email from remote servers and delivers it to the local MTA. Many users appreciate the ability to separate the process of downloading their messages located on a remote server from the process of reading and organizing their email in an MUA. Designed with the needs of dial-up users in mind, Fetchmail connects and quickly downloads all of the email messages to the mail spool file using any number of protocols, including POP3 and IMAP. It can even forward email messages to an SMTP server, if necessary.

Fetchmail is configured for each user through the use of a .fetchmailrc file in the user's home directory.

Using preferences in the .fetchmailrc file, Fetchmail checks for email on a remote server and downloads it. It then delivers it to port 25 on the local machine, using the local MTA to place the email in the correct user's spool file. If Procmail is available, it is launched to filter the email and place it in a mailbox so that it can be read by an MUA.

25.3.3.1. Fetchmail Configuration Options

Although it is possible to pass all necessary options on the command line to check for email on a remote server when executing Fetchmail, using a .fetchmailrc file is much easier. Place any desired configuration options in the .fetchmailrc file for those options to be used each time the fetchmail command is issued. It is possible to override these at the time Fetchmail is run by specifying that option on the command line.

A user's .fetchmailrc file contains three classes of configuration options:

- global options — Gives Fetchmail instructions that control the operation of the program or provide settings for every connection that checks for email.
- server options — Specifies necessary information about the server being polled, such as the hostname, as well as preferences for specific email servers, such as the port to check or number of seconds to wait before timing out. These options affect every user using that server.
- user options — Contains information, such as username and password, necessary to authenticate and check for email using a specified email server.

Global options appear at the top of the .fetchmailrc file, followed by one or more server options, each of which designate a different email server that Fetchmail should check. User options follow server options for each user account checking that email server. Like server options, multiple user options may be specified for use with a particular server as well as to check multiple email accounts on the same server.

Server options are called into service in the .fetchmailrc file by the use of a special option verb, poll or skip, that precedes any of the server information. The poll action tells Fetchmail to use this server option when it is run, which checks for email using the specified user options. Any server options after a skip action, however, are not checked unless this server's hostname is specified when Fetchmail is invoked. The skip option is useful when testing configurations in .fetchmailrc because it only checks skipped servers when specifically invoked, and does not affect any currently working configurations.

A sample .fetchmailrc file looks similar to the following example:

```
set postmaster "user1"
set bouncemail
 poll pop.domain.com proto pop3
  user 'user1' there with password 'secret' is user1 here
 poll mail.domain2.com
  user 'user5' there with password 'secret2' is user1 here
 user 'user7' there with password 'secret3' is user1 here
```

In this example, the global options specify that the user is sent email as a last resort (postmaster option) and all email errors are sent to the postmaster instead of the sender (bouncemail option). The set action tells Fetchmail that this line contains a global option. Then, two email servers are specified, one set to check using POP3, the other for trying various protocols to find one that works. Two users are checked using the second server option, but all email found for any user is sent to user1's mail spool. This allows multiple mailboxes to be checked on multiple servers, while appearing in a single MUA inbox. Each user's specific information begins with the user action.

Note

Users are not required to place their password in the .fetchmailrc file. Omitting the with password '<password>' section causes Fetchmail to ask for a password when it is launched.

Fetchmail has numerous global, server, and local options. Many of these options are rarely used or only apply to very specific situations. The fetchmail man page explains each option in detail, but the most common ones are listed here.

25.3.3.2. Global Options

Each global option should be placed on a single line after a set action.

- daemon <seconds> — Specifies daemon-mode, where Fetchmail stays in the background. Replace <seconds> with the number of seconds Fetchmail is to wait before polling the server.
- postmaster — Specifies a local user to send mail to in case of delivery problems.
- syslog — Specifies the log file for errors and status messages. By default, this is /var/log/maillog.

25.3.3.3. Server Options

Server options must be placed on their own line in .fetchmailrc after a poll or skip action.

- auth <auth-type> — Replace <auth-type> with the type of authentication to be used. By default, password authentication is used, but some protocols support other types of authentication, including kerberos_v5, kerberos_v4, and ssh. If the any authentication type is used, Fetchmail first tries methods that do not require a password, then methods that mask the password, and finally attempts to send the password unencrypted to authenticate to the server.

- interval <number> — Polls the specified server every <number> of times that it checks for email on all configured servers. This option is generally used for email servers where the user rarely receives messages.
- port <port-number> — Replace <port-number> with the port number. This value overrides the default port number for the specified protocol.
- proto <protocol> — Replace <protocol> with the protocol, such as pop3 or imap, to use when checking for messages on the server.
- timeout <seconds> — Replace <seconds> with the number of seconds of server inactivity after which Fetchmail gives up on a connection attempt. If this value is not set, a default of 300 seconds is assumed.

25.3.3.4. User Options

User options may be placed on their own lines beneath a server option or on the same line as the server option. In either case, the defined options must follow the user option (defined below).

- fetchall — Orders Fetchmail to download all messages in the queue, including messages that have already been viewed. By default, Fetchmail only pulls down new messages.
- fetchlimit <number> — Replace <number> with the number of messages to be retrieved before stopping.
- flush — Deletes all previously viewed messages in the queue before retrieving new messages.
- limit <max-number-bytes> — Replace <max-number-bytes> with the maximum size in bytes that messages are allowed to be when retrieved by Fetchmail. This option is useful with slow network links, when a large message takes too long to download.
- password '<password>' — Replace <password> with the user's password.
- preconnect "<command>" — Replace <command> with a command to be executed before retrieving messages for the user.
- postconnect "<command>" — Replace <command> with a command to be executed after retrieving messages for the user.
- ssl — Activates SSL encryption.
- user "<username>" — Replace <username> with the username used by Fetchmail to retrieve messages. This option must precede all other user options.

25.3.3.5. Fetchmail Command Options

Most Fetchmail options used on the command line when executing the fetchmail command mirror the .fetchmailrc configuration options. In this way, Fetchmail may be used with or without a configuration file. These options are not used on the command line by most users because it is easier to leave them in the .fetchmailrc file.

There may be times when it is desirable to run the fetchmail command with other options for a particular purpose. It is possible to issue command options to temporarily override a .fetchmailrc setting that is causing an error, as any options specified at the command line override configuration file options.

25.3.3.6. Informational or Debugging Options

Certain options used after the fetchmail command can supply important information.

- --configdump — Displays every possible option based on information from .fetchmailrc and Fetchmail defaults. No email is retrieved for any users when using this option.
- -s — Executes Fetchmail in silent mode, preventing any messages, other than errors, from appearing after the fetchmail command.
- -v — Executes Fetchmail in verbose mode, displaying every communication between Fetchmail and remote email servers.
- -V — Displays detailed version information, lists its global options, and shows settings to be used with each user, including the email protocol and authentication method. No email is retrieved for any users when using this option.

25.3.3.7. Special Options

These options are occasionally useful for overriding defaults often found in the .fetchmailrc file.

- -a — Fetchmail downloads all messages from the remote email server, whether new or previously viewed. By default, Fetchmail only downloads new messages.
- -k — Fetchmail leaves the messages on the remote email server after downloading them. This option overrides the default behavior of deleting messages after downloading them.
- -l <max-number-bytes> — Fetchmail does not download any messages over a particular size and leaves them on the remote email server.
- --quit — Quits the Fetchmail daemon process.

More commands and .fetchmailrc options can be found in the fetchmail man page.

25.4. Mail Transport Agent (MTA) Configuration

A Mail Transport Agent (MTA) is essential for sending email. A Mail User Agent (MUA) such as Evolution, Thunderbird, and Mutt, is used to read and compose email. When a user sends an email from an MUA, the message is handed off to the MTA, which sends the message through a series of MTAs until it reaches its destination.

Even if a user does not plan to send email from the system, some automated tasks or system programs might use the /bin/mail command to send email containing log messages to the root user of the local system.

Red Hat Enterprise Linux 5 provides three MTAs: Sendmail, Postfix, and Exim. If all three are installed, sendmail is the default MTA. The Mail Transport Agent Switcher allows for the selection of either sendmail, postfix, or exim as the default MTA for the system.

The system-switch-mail RPM package must be installed to use the text-based version of the Mail Transport Agent Switcher program. If you want to use the graphical version, the system-switch-mail-gnome package must also be installed.

알림

For more information on installing RPM packages, refer to II부. 패키지 관리.

To start the Mail Transport Agent Switcher, select System (the main menu on the panel) > Administration > Mail Transport Agent Switcher, or type the command system-switch-mail at a shell prompt (for example, in an XTerm or GNOME terminal).

The program automatically detects if the X Window System is running. If it is running, the program starts in graphical mode as shown in 그림 25.1. " Mail Transport Agent Switcher " . If X is not detected, it starts in text-mode. To force Mail Transport Agent Switcher to run in text-mode, use the command system-switch-mail-nox.

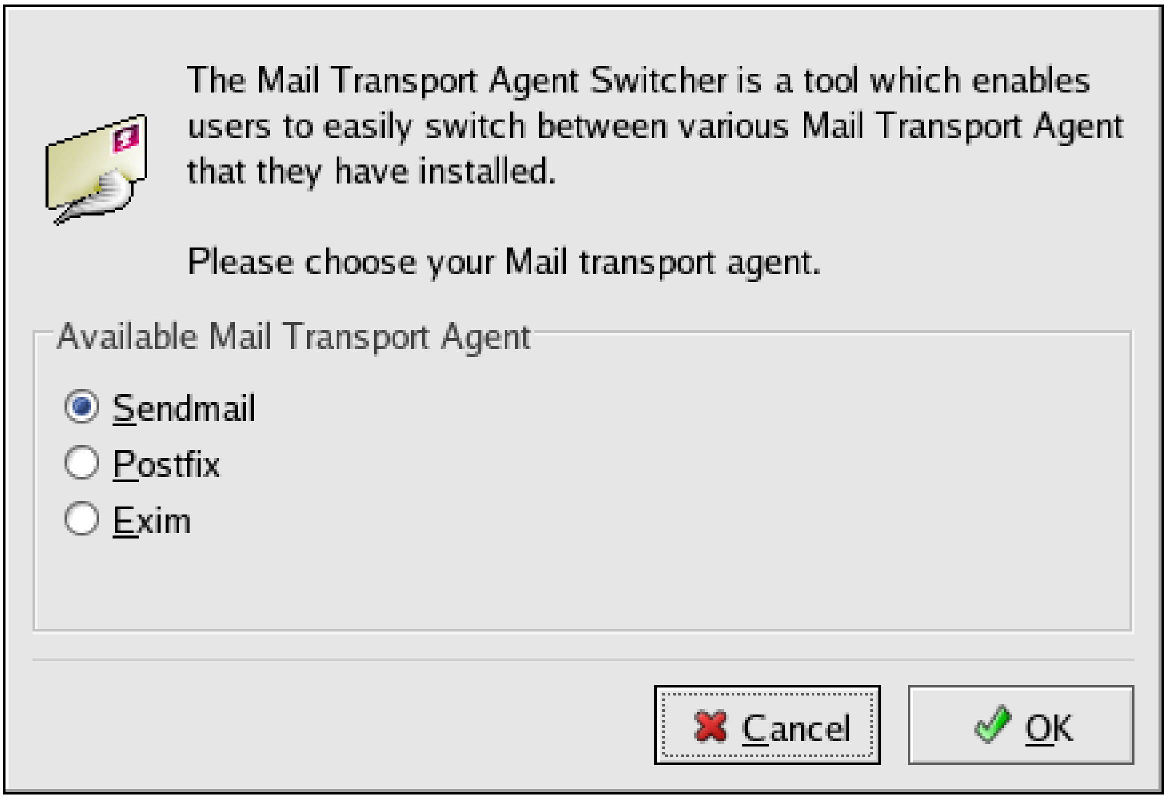

The Mail Transport Agent Switcher is a tool which enables users to easily switch between various Mail Transport Agent that they have installed.

Please choose your Mail transport agent.

Available Mail Transport Agent

● Sendmail
○ Postfix
○ Exim

그림 25.1. Mail Transport Agent Switcher

If you select OK to change the MTA, the selected mail daemon is enabled to start at boot time, and the unselected mail daemons are disabled so that they do not start at boot time. The selected mail daemon is started, and any other mail daemon is stopped; thus making the changes take place immediately.

25.5. Mail Delivery Agents

Red Hat Enterprise Linux includes two primary MDAs, Procmail and mail. Both of the applications are considered LDAs and both move email from the MTA's spool file into the user's mailbox. However, Procmail provides a robust filtering system.

This section details only Procmail. For information on the mail command, consult its man page.

Procmail delivers and filters email as it is placed in the mail spool file of the localhost. It is powerful, gentle on system resources, and widely used. Procmail can play a critical role in delivering email to be read by email client applications.

Procmail can be invoked in several different ways. Whenever an MTA places an email into the mail spool file, Procmail is launched. Procmail then filters and files the email for the MUA and quits. Alternatively, the MUA can be configured to execute Procmail any time a message is received so that messages are moved into their correct mailboxes. By default, the presence of /etc/procmailrc or of a .procmailrc file (also called an rc file) in the user's home directory invokes Procmail whenever an MTA receives a new message.

Whether Procmail acts upon an email message depends upon whether the message matches a specified set of conditions or recipes in the rc file. If a message matches a recipe, then the email is placed in a specified file, is deleted, or is otherwise processed.

When Procmail starts, it reads the email message and separates the body from the header information. Next, Procmail looks for /etc/procmailrc and rc files in the /etc/procmailrcs directory for default, system-wide, Procmail environmental variables and recipes. Procmail then searches for a .procmailrc file in the user's home directory. Many users also create additional rc files for Procmail that are referred to within the .procmailrc file in their home directory.

By default, no system-wide rc files exist in the /etc/ directory and no .procmailrc files exist in any user's home directory. Therefore, to use Procmail, each user must construct a .procmailrc file with specific environment variables and rules.

25.5.1. Procmail Configuration

The Procmail configuration file contains important environmental variables. These variables specify things such as which messages to sort and what to do with the messages that do not match any recipes.

These environmental variables usually appear at the beginning of .procmailrc in the following format:

```
<env-variable>="<value>"
```

In this example, <env-variable> is the name of the variable and <value> defines the variable.

There are many environment variables not used by most Procmail users and many of the more important environment variables are already defined by a default value. Most of the time, the following variables are used:

- DEFAULT — Sets the default mailbox where messages that do not match any recipes are placed.

  The default DEFAULT value is the same as $ORGMAIL.

- INCLUDERC — Specifies additional rc files containing more recipes for messages to be checked against. This breaks up the Procmail recipe lists into individual files that fulfill different roles, such as blocking spam and managing email lists, that can then be turned off or on by using comment characters in the user's .procmailrc file.

  For example, lines in a user's .procmailrc file may look like this:

  ```
  MAILDIR=$HOME/Msgs
  INCLUDERC=$MAILDIR/lists.rc
  INCLUDERC=$MAILDIR/spam.rc
  ```

If the user wants to turn off Procmail filtering of their email lists but leave spam control in place, they would comment out the first INCLUDERC line with a hash mark character (#).

- LOCKSLEEP — Sets the amount of time, in seconds, between attempts by Procmail to use a particular lockfile. The default is eight seconds.
- LOCKTIMEOUT — Sets the amount of time, in seconds, that must pass after a lockfile was last modified before Procmail assumes that the lockfile is old and can be deleted. The default is 1024 seconds.
- LOGFILE — The file to which any Procmail information or error messages are written.
- MAILDIR — Sets the current working directory for Procmail. If set, all other Procmail paths are relative to this directory.
- ORGMAIL — Specifies the original mailbox, or another place to put the messages if they cannot be placed in the default or recipe-required location.

  By default, a value of /var/spool/mail/$LOGNAME is used.
- SUSPEND — Sets the amount of time, in seconds, that Procmail pauses if a necessary resource, such as swap space, is not available.
- SWITCHRC — Allows a user to specify an external file containing additional Procmail recipes, much like the INCLUDERC option, except that recipe checking is actually stopped on the referring configuration file and only the recipes on the SWITCHRC-specified file are used.
- VERBOSE — Causes Procmail to log more information. This option is useful for debugging.

Other important environmental variables are pulled from the shell, such as LOGNAME, which is the login name; HOME, which is the location of the home directory; and SHELL, which is the default shell.

A comprehensive explanation of all environments variables, as well as their default values, is available in the procmailrc man page.

25.5.2. Procmail Recipes

New users often find the construction of recipes the most difficult part of learning to use Procmail. To some extent, this is understandable, as recipes do their message matching using regular expressions, which is a particular format used to specify qualifications for a matching string. However, regular expressions are not very difficult to construct and even less difficult to understand when read. Additionally, the consistency of the way Procmail recipes are written, regardless of regular expressions, makes it easy to learn by example. To see example Procmail recipes, refer to 25.5.2.5절. “Recipe Examples” .

Procmail recipes take the following form:

```
:0<flags>: <lockfile-name>
* <special-condition-character>
    <condition-1>
* <special-condition-character>
    <condition-2>
* <special-condition-character>
    <condition-N>
    <special-action-character>
    <action-to-perform>
```

The first two characters in a Procmail recipe are a colon and a zero. Various flags can be placed after the zero to control how Procmail processes the recipe. A colon after the <flags> section specifies that a lockfile is created for this message. If a lockfile is created, the name can be specified by replacing <lockfile-name> .

A recipe can contain several conditions to match against the message. If it has no conditions, every message matches the recipe. Regular expressions are placed in some conditions to facilitate message matching. If multiple conditions are used, they must all match for the action to be performed. Conditions are checked based on the flags set in the recipe's first line. Optional special characters placed after the * character can further control the condition.

The <action-to-perform> specifies the action taken when the message matches one of the conditions. There can only be one action per recipe. In many cases, the name of a mailbox is used here to direct matching messages into that file, effectively sorting the email. Special action characters may also be used before the action is specified. Refer to 25.5.2.4절. "Special Conditions and Actions" for more information.

25.5.2.1. Delivering vs. Non-Delivering Recipes

The action used if the recipe matches a particular message determines whether it is considered a delivering or non-delivering recipe. A delivering recipe contains an action that writes the message to a file, sends the message to another program, or forwards the message to another email address. A non-delivering recipe covers any other actions, such as a nesting block. A nesting block is a set of actions, contained in braces { }, that are performed on messages which match the recipe's conditions. Nesting blocks can be nested inside one another, providing greater control for identifying and performing actions on messages.

When messages match a delivering recipe, Procmail performs the specified action and stops comparing the message against any other recipes. Messages that match non-delivering recipes continue to be compared against other recipes.

25.5.2.2. Flags

Flags are essential to determine how or if a recipe's conditions are compared to a message. The following flags are commonly used:

- A — Specifies that this recipe is only used if the previous recipe without an A or a flag also matched this message.
- a — Specifies that this recipe is only used if the previous recipe with an A or a flag also matched this message and was successfully completed.
- B — Parses the body of the message and looks for matching conditions.
- b — Uses the body in any resulting action, such as writing the message to a file or forwarding it. This is the default behavior.
- c — Generates a carbon copy of the email. This is useful with delivering recipes, since the required action can be performed on the message and a copy of the message can continue being processed in the rc files.
- D — Makes the egrep comparison case-sensitive. By default, the comparison process is not case-sensitive.

- E — While similar to the A flag, the conditions in the recipe are only compared to the message if the immediately preceding the recipe without an E flag did not match. This is comparable to an else action.
- e — The recipe is compared to the message only if the action specified in the immediately preceding recipe fails.
- f — Uses the pipe as a filter.
- H — Parses the header of the message and looks for matching conditions. This occurs by default.
- h — Uses the header in a resulting action. This is the default behavior.
- w — Tells Procmail to wait for the specified filter or program to finish, and reports whether or not it was successful before considering the message filtered.
- W — Is identical to w except that "Program failure" messages are suppressed.

For a detailed list of additional flags, refer to the procmailrc man page.

25.5.2.3. Specifying a Local Lockfile

Lockfiles are very useful with Procmail to ensure that more than one process does not try to alter a message simultaneously. Specify a local lockfile by placing a colon (:) after any flags on a recipe's first line. This creates a local lockfile based on the destination file name plus whatever has been set in the LOCKEXT global environment variable.

Alternatively, specify the name of the local lockfile to be used with this recipe after the colon.

25.5.2.4. Special Conditions and Actions

Special characters used before Procmail recipe conditions and actions change the way they are interpreted.

The following characters may be used after the * character at the beginning of a recipe's condition line:

- ! — In the condition line, this character inverts the condition, causing a match to occur only if the condition does not match the message.
- < — Checks if the message is under a specified number of bytes.
- > — Checks if the message is over a specified number of bytes.

The following characters are used to perform special actions:

- ! — In the action line, this character tells Procmail to forward the message to the specified email addresses.
- $ — Refers to a variable set earlier in the rc file. This is often used to set a common mailbox that is referred to by various recipes.
- | — Starts a specified program to process the message.
- { and } — Constructs a nesting block, used to contain additional recipes to apply to matching messages.

If no special character is used at the beginning of the action line, Procmail assumes that the action line is specifying the mailbox in which to write the message.

25.5.2.5. Recipe Examples

Procmail is an extremely flexible program, but as a result of this flexibility, composing Procmail recipes from scratch can be difficult for new users.

The best way to develop the skills to build Procmail recipe conditions stems from a strong understanding of regular expressions combined with looking at many examples built by others. A thorough explanation of regular expressions is beyond the scope of this section. The structure of Procmail recipes and useful sample Procmail recipes can be found at various places on the Internet (such as http://www.iki.fi/era/procmail/links.html). The proper use and adaptation of regular expressions can be derived by viewing these recipe examples. In addition, introductory information about basic regular expression rules can be found in the grep man page.

The following simple examples demonstrate the basic structure of Procmail recipes and can provide the foundation for more intricate constructions.

A basic recipe may not even contain conditions, as is illustrated in the following example:

```
:0:
new-mail.spool
```

The first line specifies that a local lockfile is to be created but does not specify a name, so Procmail uses the destination file name and appends the value specified in the LOCKEXT environment variable. No condition is specified, so every message matches this recipe and is placed in the single spool file called new-mail.spool, located within the directory specified by the MAILDIR environment variable. An MUA can then view messages in this file.

A basic recipe, such as this, can be placed at the end of all rc files to direct messages to a default location.

The following example matched messages from a specific email address and throws them away.

```
:0
* ^From: spammer@domain.com
/dev/null
```

With this example, any messages sent by spammer@domain.com are sent to the /dev/null device, deleting them.

Caution

Be certain that rules are working as intended before sending messages to /dev/null for permanent deletion. If a recipe inadvertently catches unintended messages, and those messages disappear, it becomes difficult to troubleshoot the rule.

A better solution is to point the recipe's action to a special mailbox, which can be checked from time to time to look for false positives. Once satisfied that no messages are accidentally being matched, delete the mailbox and direct the action to send the messages to /dev/null.

The following recipe grabs email sent from a particular mailing list and places it in a specified folder.

```
:0:
* ^(From|CC|To).*tux-lug
tuxlug
```

Any messages sent from the tux-lug@domain.com mailing list are placed in the tuxlug mailbox automatically for the MUA. Note that the condition in this example matches the message if it has the mailing list's email address on the From, CC, or To lines.

Consult the many Procmail online resources available in 25.7절. "Additional Resources" for more detailed and powerful recipes.

25.5.2.6. Spam Filters

Because it is called by Sendmail, Postfix, and Fetchmail upon receiving new emails, Procmail can be used as a powerful tool for combating spam.

This is particularly true when Procmail is used in conjunction with SpamAssassin. When used together, these two applications can quickly identify spam emails, and sort or destroy them.

SpamAssassin uses header analysis, text analysis, blacklists, a spam-tracking database, and self-learning Bayesian spam analysis to quickly and accurately identify and tag spam.

The easiest way for a local user to use SpamAssassin is to place the following line near the top of the ~/.procmailrc file:

```
INCLUDERC=/etc/mail/spamassassin/spamassassin-default.rc
```

The /etc/mail/spamassassin/spamassassin-default.rc contains a simple Procmail rule that activates SpamAssassin for all incoming email. If an email is determined to be spam, it is tagged in the header as such and the title is prepended with the following pattern:

```
*****SPAM*****
```

The message body of the email is also prepended with a running tally of what elements caused it to be diagnosed as spam.

To file email tagged as spam, a rule similar to the following can be used:

```
:0 Hw
* ^X-Spam-Status: Yes
spam
```

This rule files all email tagged in the header as spam into a mailbox called spam.

Since SpamAssassin is a Perl script, it may be necessary on busy servers to use the binary SpamAssassin daemon (spamd) and client application (spamc). Configuring SpamAssassin this way, however, requires root access to the host.

To start the spamd daemon, type the following command as root:

```
service spamassassin start
```

To start the SpamAssassin daemon when the system is booted, use an initscript utility, such as the Services Configuration Tool (system-config-services), to turn on the spamassassin service. Refer to 17장. 서비스로의 접근 통제 for more information about initscript utilities.

To configure Procmail to use the SpamAssassin client application instead of the Perl script, place the following line near the top of the ~/.procmailrc file. For a system-wide configuration, place it in /etc/procmailrc:

```
INCLUDERC=/etc/mail/spamassassin/spamassassin-spamc.rc
```

25.6. Mail User Agents

There are scores of mail programs available under Red Hat Enterprise Linux. There are full-featured, graphical email client programs, such as Ximian Evolution, as well as text-based email programs such as mutt.

The remainder of this section focuses on securing communication between the client and server.

25.6.1. Securing Communication

Popular MUAs included with Red Hat Enterprise Linux, such as Ximian Evolution and mutt offer SSL-encrypted email sessions.

Like any other service that flows over a network unencrypted, important email information, such as usernames, passwords, and entire messages, may be intercepted and viewed by users on the network. Additionally, since the standard POP and IMAP protocols pass authentication information unencrypted, it is possible for an attacker to gain access to user accounts by collecting usernames and passwords as they are passed over the network.

25.6.1.1. Secure Email Clients

Most Linux MUAs designed to check email on remote servers support SSL encryption. To use SSL when retrieving email, it must be enabled on both the email client and server.

SSL is easy to enable on the client-side, often done with the click of a button in the MUA's configuration window or via an option in the MUA's configuration file. Secure IMAP and POP have known port numbers (993 and 995, respectively) that the MUA uses to authenticate and download messages.

25.6.1.2. Securing Email Client Communications

Offering SSL encryption to IMAP and POP users on the email server is a simple matter.

First, create an SSL certificate. This can be done two ways: by applying to a Certificate Authority (CA) for an SSL certificate or by creating a self-signed certificate.

Caution

Self-signed certificates should be used for testing purposes only. Any server used in a production environment should use an SSL certificate granted by a CA.

To create a self-signed SSL certificate for IMAP, change to the /etc/pki/tls/certs/ directory and type the following commands as root:

```
rm -f cyrus-imapd.pem make cyrus-imapd.pem
```

Answer all of the questions to complete the process.

To create a self-signed SSL certificate for POP, change to the /etc/pki/tls/certs/ directory, and type the following commands as root:

```
rm -f ipop3d.pem make ipop3d.pem
```

Again, answer all of the questions to complete the process.

Important

Please be sure to remove the default imapd.pem and ipop3d.pem files before issuing each make command.

Once finished, execute the /sbin/service xinetd restart command to restart the xinetd daemon which controls imapd and ipop3d.

Alternatively, the stunnel command can be used as an SSL encryption wrapper around the standard, non-secure daemons, imapd or pop3d.

The stunnel program uses external OpenSSL libraries included with Red Hat Enterprise Linux to provide strong cryptography and protect the connections. It is best to apply to a CA to obtain an SSL certificate, but it is also possible to create a self-signed certificate.

To create a self-signed SSL certificate, change to the /etc/pki/tls/certs/ directory, and type the following command:

```
make stunnel.pem
```

Again, answer all of the questions to complete the process.

Once the certificate is generated, it is possible to use the stunnel command to start the imapd mail daemon using the following command:

```
/usr/sbin/stunnel -d 993 -l /usr/sbin/imapd imapd
```

Once this command is issued, it is possible to open an IMAP email client and connect to the email server using SSL encryption.

To start the pop3d using the stunnel command, type the following command:

```
/usr/sbin/stunnel -d 995 -l /usr/sbin/pop3d pop3d
```

For more information about how to use stunnel, read the stunnel man page or refer to the documents in the /usr/share/doc/stunnel-<version-number> / directory, where <version-number> is the version number for stunnel.

25.7. Additional Resources

The following is a list of additional documentation about email applications.

25.7.1. Installed Documentation

- Information on configuring Sendmail is included with the sendmail and sendmail-cf packages.

  - /usr/share/sendmail-cf/README — Contains information on m4, file locations for Sendmail, supported mailers, how to access enhanced features, and more.

  In addition, the sendmail and aliases man pages contain helpful information covering various Sendmail options and the proper configuration of the Sendmail /etc/mail/aliases file.

- /usr/share/doc/postfix-<version-number> — Contains a large amount of information about ways to configure Postfix. Replace <version-number> with the version number of Postfix.

- /usr/share/doc/fetchmail-<version-number> — Contains a full list of Fetchmail features in the FEATURES file and an introductory FAQ document. Replace <version-number> with the version number of Fetchmail.

- /usr/share/doc/procmail-<version-number> — Contains a README file that provides an overview of Procmail, a FEATURES file that explores every program feature, and an FAQ file with answers to many common configuration questions. Replace <version-number> with the version number of Procmail.

  When learning how Procmail works and creating new recipes, the following Procmail man pages are invaluable:

  - procmail — Provides an overview of how Procmail works and the steps involved with filtering email.

  - procmailrc — Explains the rc file format used to construct recipes.

  - procmailex — Gives a number of useful, real-world examples of Procmail recipes.

  - procmailsc — Explains the weighted scoring technique used by Procmail to match a particular recipe to a message.

  - /usr/share/doc/spamassassin-<version-number>/ — Contains a large amount of information pertaining to SpamAssassin. Replace <version-number> with the version number of the spamassassin package.

25.7.2. Useful Websites

- http://www.sendmail.org/ — Offers a thorough technical breakdown of Sendmail features, documentation and configuration examples.

- http://www.sendmail.com/ — Contains news, interviews and articles concerning Sendmail, including an expanded view of the many options available.

- http://www.postfix.org/ — The Postfix project home page contains a wealth of information about Postfix. The mailing list is a particularly good place to look for information.

- http://fetchmail.berlios.de/ — The home page for Fetchmail, featuring an online manual, and a thorough FAQ.

- http://www.procmail.org/ — The home page for Procmail with links to assorted mailing lists dedicated to Procmail as well as various FAQ documents.
- http://partmaps.org/era/procmail/mini-faq.html — An excellent Procmail FAQ, offers troubleshooting tips, details about file locking, and the use of wildcard characters.
- http://www.uwasa.fi/~ts/info/proctips.html — Contains dozens of tips that make using Procmail much easier. Includes instructions on how to test .procmailrc files and use Procmail scoring to decide if a particular action should be taken.
- http://www.spamassassin.org/ — The official site of the SpamAssassin project.

25.7.3. Related Books

- Sendmail Milters: A Guide for Fighting Spam by Bryan Costales and Marcia Flynt; Addison-Wesley — A good Sendmail guide that can help you customise your mail filters.
- Sendmail by Bryan Costales with Eric Allman et al; O'Reilly & Associates — A good Sendmail reference written with the assistance of the original creator of Delivermail and Sendmail.
- Removing the Spam: Email Processing and Filtering by Geoff Mulligan; Addison-Wesley Publishing Company — A volume that looks at various methods used by email administrators using established tools, such as Sendmail and Procmail, to manage spam problems.
- Internet Email Protocols: A Developer's Guide by Kevin Johnson; Addison-Wesley Publishing Company — Provides a very thorough review of major email protocols and the security they provide.
- Managing IMAP by Dianna Mullet and Kevin Mullet; O'Reilly & Associates — Details the steps required to configure an IMAP server.

Lightweight Directory Access Protocol (LDAP)

The Lightweight Directory Access Protocol (LDAP) is a set of open protocols used to access centrally stored information over a network. It is based on the X.500 standard for directory sharing, but is less complex and resource-intensive. For this reason, LDAP is sometimes referred to as "X.500 Lite." The X.500 standard is a directory that contains hierarchical and categorized information, which could include information such as names, addresses, and phone numbers.

Like X.500, LDAP organizes information in a hierarchal manner using directories. These directories can store a variety of information and can even be used in a manner similar to the Network Information Service (NIS), enabling anyone to access their account from any machine on the LDAP enabled network.

In many cases, LDAP is used as a virtual phone directory, allowing users to easily access contact information for other users. But LDAP is more flexible than a traditional phone directory, as it is capable of referring a querent to other LDAP servers throughout the world, providing an ad-hoc global repository of information. Currently, however, LDAP is more commonly used within individual organizations, like universities, government departments, and private companies.

LDAP is a client/server system. The server can use a variety of databases to store a directory, each optimized for quick and copious read operations. When an LDAP client application connects to an LDAP server, it can either query a directory or attempt to modify it. In the event of a query, the server either answers the query locally, or it can refer the querent to an LDAP server which does have the answer. If the client application is attempting to modify information within an LDAP directory, the server verifies that the user has permission to make the change and then adds or updates the information.

This chapter refers to the configuration and use of OpenLDAP 2.0, an open source implementation of the LDAPv2 and LDAPv3 protocols.

26.1. Why Use LDAP?

The main benefit of using LDAP is that information for an entire organization can be consolidated into a central repository. For example, rather than managing user lists for each group within an organization, LDAP can be used as a central directory accessible from anywhere on the network. And because LDAP supports Secure Sockets Layer (SSL) and Transport Layer Security (TLS), sensitive data can be protected from prying eyes.

LDAP also supports a number of back-end databases in which to store directories. This allows administrators the flexibility to deploy the database best suited for the type of information the server is to disseminate. Because LDAP also has a well-defined client Application Programming Interface (API), the number of LDAP-enabled applications are numerous and increasing in quantity and quality.

26.1.1. OpenLDAP Features

OpenLDAP includes a number of important features.

- LDAPv3 Support — OpenLDAP supports Simple Authentication and Security Layer (SASL), Transport Layer Security (TLS), and Secure Sockets Layer (SSL), among other improvements. Many of the changes in the protocol since LDAPv2 are designed to make LDAP more secure.

- IPv6 Support — OpenLDAP supports the next generation Internet Protocol version 6.
- LDAP Over IPC — OpenLDAP can communicate within a system using interprocess communication (IPC). This enhances security by eliminating the need to communicate over a network.
- Updated C API — Improves the way programmers can connect to and use LDAP directory servers.
- LDIFv1 Support — Provides full compliance with the LDAP Data Interchange Format (LDIF) version 1.
- Enhanced Stand-Alone LDAP Server — Includes an updated access control system, thread pooling, better tools, and much more.

26.2. LDAP Terminology

Any discussion of LDAP requires a basic understanding of a set of LDAP-specific terms:

- entry — A single unit within an LDAP directory. Each entry is identified by its unique Distinguished Name (DN).
- attributes — Information directly associated with an entry. For example, an organization could be represented as an LDAP entry. Attributes associated with the organization might include a fax number, an address, and so on. People can also be represented as entries in an LDAP directory, with common attributes such as the person's telephone number and email address.

  Some attributes are required, while other attributes are optional. An objectclass definition sets which attributes are required for each entry. Objectclass definitions are found in various schema files, located in the /etc/openldap/schema/ directory. For more information, refer to 26.5절. "The /etc/openldap/schema/ Directory" .

  The assertion of an attribute and its corresponding value is also referred to as a Relative Distinguished Name (RDN). An RDN is only unique per entry, whereas a DN is globally unique.
- LDIF — The LDAP Data Interchange Format (LDIF) is an ASCII text representation of LDAP entries. Files used for importing data to LDAP servers must be in LDIF format. An LDIF entry looks similar to the following example:

  ```
  [<id>] dn: <distinguished name>
  <attrtype>: <attrvalue>
  <attrtype>: <attrvalue>
  <attrtype>: <attrvalue>
  ```

  Each entry can contain as many <attrtype>: <attrvalue> pairs as needed. A blank line indicates the end of an entry.

  > **Caution**
  >
  > All <attrtype> and <attrvalue> pairs must be defined in a corresponding schema file to use this information.

Any value enclosed within a < and a > is a variable and can be set whenever a new LDAP entry is created. This rule does not apply, however, to <id>. The <id> is a number determined by the application used to edit the entry.

26.3. OpenLDAP Daemons and Utilities

The suite of OpenLDAP libraries and tools are included within the following packages:

- openldap — Contains the libraries necessary to run the OpenLDAP server and client applications.
- openldap-clients — Contains command line tools for viewing and modifying directories on an LDAP server.
- openldap-servers — Contains the servers and other utilities necessary to configure and run an LDAP server.

There are two servers contained in the openldap-servers package: the Standalone LDAP Daemon (/usr/sbin/slapd) and the Standalone LDAP Update Replication Daemon (/usr/sbin/slurpd).

The slapd daemon is the standalone LDAP server while the slurpd daemon is used to synchronize changes from one LDAP server to other LDAP servers on the network. The slurpd daemon is only used when dealing with multiple LDAP servers.

To perform administrative tasks, the openldap-servers package installs the following utilities into the /usr/sbin/ directory:

- slapadd — Adds entries from an LDIF file to an LDAP directory. For example, the command /usr/sbin/slapadd -l ldif-input reads in the LDIF file, ldif-input, containing the new entries.

Important

Only the root user may use /usr/sbin/slapadd. However, the directory server runs as the ldap user. Therefore the directory server is unable to modify any files created by slapadd. To correct this issue, after using slapadd, type the following command:

```
chown -R ldap /var/lib/ldap
```

- slapcat — Pulls entries from an LDAP directory in the default format, Sleepycat Software's Berkeley DB system, and saves them in an LDIF file. For example, the command /usr/sbin/slapcat -l ldif-output outputs an LDIF file called ldif-output containing the entries from the LDAP directory.
- slapindex — Re-indexes the slapd directory based on the current content. This tool should be run whenever indexing options within /etc/openldap/slapd.conf are changed.
- slappasswd — Generates an encrypted user password value for use with ldapmodify or the rootpw value in the slapd configuration file, /etc/openldap/slapd.conf. Execute the /usr/sbin/slappasswd command to create the password.

Warning

You must stop slapd by issuing the /sbin/service ldap stop command before using slapadd, slapcat or slapindex. Otherwise, the integrity of the LDAP directory is at risk.

For more information on using these utilities, refer to their respective man pages.

The openldap-clients package installs tools into /usr/bin/ which are used to add, modify, and delete entries in an LDAP directory. These tools include the following:

- ldapadd — Adds entries to an LDAP directory by accepting input via a file or standard input; ldapadd is actually a hard link to ldapmodify -a.
- ldapdelete — Deletes entries from an LDAP directory by accepting user input at a shell prompt or via a file.
- ldapmodify — Modifies entries in an LDAP directory, accepting input via a file or standard input.
- ldappasswd — Sets the password for an LDAP user.
- ldapsearch — Searches for entries in an LDAP directory using a shell prompt.
- ldapcompare — Opens a connection to an LDAP server, binds, and performs a comparison using specified parameters.
- ldapwhoami — Opens a connection to an LDAP server, binds, and performs a whoami operation.
- ldapmodrdn — Opens a connection to an LDAP server, binds, and modifies the RDNs of entries.

With the exception of ldapsearch, each of these utilities is more easily used by referencing a file containing the changes to be made rather than typing a command for each entry to be changed within an LDAP directory. The format of such a file is outlined in the man page for each utility.

26.3.1. NSS, PAM, and LDAP

In addition to the OpenLDAP packages, Red Hat Enterprise Linux includes a package called nss_ldap, which enhances LDAP's ability to integrate into both Linux and other UNIX environments.

The nss_ldap package provides the following modules (where <version> refers to the version of libnss_ldap in use):

- /lib/libnss_ldap-<version>.so
- /lib/security/pam_ldap.so

The nss_ldap package provides the following modules for Itanium or AMD64 architectures:

- /lib64/libnss_ldap-<version>.so
- /lib64/security/pam_ldap.so

The libnss_ldap-<version>.so module allows applications to look up users, groups, hosts, and other information using an LDAP directory via the Nameservice Switch (NSS) interface of glibc. NSS allows applications to authenticate using LDAP in conjunction with the NIS name service and flat authentication files.

The pam_ldap module allows PAM-aware applications to authenticate users using information stored in an LDAP directory. PAM-aware applications include console login, POP and IMAP mail servers, and Samba. By deploying an LDAP server on a network, all of these applications can authenticate using the same user ID and password combination, greatly simplifying administration.

For more about configuring PAM, refer to 46.4절. "PAM (Pluggable Authentication Modules)" and the PAM man pages.

26.3.2. PHP4, LDAP, and the Apache HTTP Server

Red Hat Enterprise Linux includes a package containing an LDAP module for the PHP server-side scripting language.

The php-ldap package adds LDAP support to the PHP4 HTML-embedded scripting language via the /usr/lib/php4/ldap.so module. This module allows PHP4 scripts to access information stored in an LDAP directory.

Red Hat Enterprise Linux ships with the mod_authz_ldap module for the Apache HTTP Server. This module uses the short form of the distinguished name for a subject and the issuer of the client SSL certificate to determine the distinguished name of the user within an LDAP directory. It is also capable of authorizing users based on attributes of that user's LDAP directory entry, determining access to assets based on the user and group privileges of the asset, and denying access for users with expired passwords. The mod_ssl module is required when using the mod_authz_ldap module.

Important

The mod_authz_ldap module does not authenticate a user to an LDAP directory using an encrypted password hash. This functionality is provided by the experimental mod_auth_ldap module, which is not included with Red Hat Enterprise Linux. Refer to the Apache Software Foundation website online at http://www.apache.org/ for details on the status of this module.

26.3.3. LDAP Client Applications

There are graphical LDAP clients available which support creating and modifying directories, but they are not included with Red Hat Enterprise Linux. One such application is LDAP Browser/Editor — A Java-based tool available online at http://www.iit.edu/~gawojar/ldap/.

Other LDAP clients access directories as read-only, using them to reference, but not alter, organization-wide information. Some examples of such applications are Sendmail, Mozilla, Gnome Meeting, and Evolution.

26.4. OpenLDAP Configuration Files

OpenLDAP configuration files are installed into the /etc/openldap/ directory. The following is a brief list highlighting the most important directories and files:

- /etc/openldap/ldap.conf — This is the configuration file for all client applications which use the OpenLDAP libraries such as ldapsearch, ldapadd, Sendmail, Evolution, and Gnome Meeting.
- /etc/openldap/slapd.conf — This is the configuration file for the slapd daemon. Refer to 26.6.1절. "Editing /etc/openldap/slapd.conf" for more information.
- /etc/openldap/schema/ directory — This subdirectory contains the schema used by the slapd daemon. Refer to 26.5절. "The /etc/openldap/schema/ Directory" for more information.

Note

If the nss_ldap package is installed, it creates a file named /etc/ldap.conf. This file is used by the PAM and NSS modules supplied by the nss_ldap package. Refer to 26.7절. "Configuring a System to Authenticate Using OpenLDAP" for more information.

26.5. The /etc/openldap/schema/ Directory

The /etc/openldap/schema/ directory holds LDAP definitions, previously located in the slapd.at.conf and slapd.oc.conf files. The /etc/openldap/schema/redhat/ directory holds customized schemas distributed by Red Hat for Red Hat Enterprise Linux.

All attribute syntax definitions and objectclass definitions are now located in the different schema files. The various schema files are referenced in /etc/openldap/slapd.conf using include lines, as shown in this example:

```
include    /etc/openldap/schema/core.schema
include    /etc/openldap/schema/cosine.schema
include    /etc/openldap/schema/inetorgperson.schema
include    /etc/openldap/schema/nis.schema
include    /etc/openldap/schema/rfc822-MailMember.schema
include    /etc/openldap/schema/redhat/autofs.schema
```

Caution

Do not modify schema items defined in the schema files installed by OpenLDAP.

It is possible to extend the schema used by OpenLDAP to support additional attribute types and object classes using the default schema files as a guide. To do this, create a local.schema file in the /etc/openldap/schema/ directory. Reference this new schema within slapd.conf by adding the following line below the default include schema lines:

```
include    /etc/openldap/schema/local.schema
```

Next, define new attribute types and object classes within the local.schema file. Many organizations use existing attribute types from the schema files installed by default and add new object classes to the local.schema file.

Extending the schema to match certain specialized requirements is quite involved and beyond the scope of this chapter. Refer to http://www.openldap.org/doc/admin/schema.html for information.

26.6. OpenLDAP Setup Overview

This section provides a quick overview for installing and configuring an OpenLDAP directory. For more details, refer to the following URLs:

- http://www.openldap.org/doc/admin/quickstart.html — The Quick-Start Guide on the OpenLDAP website.
- http://www.tldp.org/HOWTO/LDAP-HOWTO/index.html — The LDAP Linux HOWTO from the Linux Documentation Project.

The basic steps for creating an LDAP server are as follows:

1. Install the openldap, openldap-servers, and openldap-clients RPMs.

2. Edit the /etc/openldap/slapd.conf file to specify the LDAP domain and server. Refer to 26.6.1절. "Editing /etc/openldap/slapd.conf" for more information.

3. Start slapd with the command:

   ```
   service ldap start
   ```

   After configuring LDAP, use chkconfig, /usr/sbin/ntsysv, or the Services Configuration Tool to configure LDAP to start at boot time. For more information about configuring services, refer to 17장. 서비스로의 접근 통제.

4. Add entries to an LDAP directory with ldapadd.

5. Use ldapsearch to determine if slapd is accessing the information correctly.

6. At this point, the LDAP directory should be functioning properly and can be configured with LDAP-enabled applications.

26.6.1. Editing /etc/openldap/slapd.conf

To use the slapd LDAP server, modify its configuration file, /etc/openldap/slapd.conf, to specify the correct domain and server.

The suffix line names the domain for which the LDAP server provides information and should be changed from:

```
suffix          "dc=your-domain,dc=com"
```

Edit it accordingly so that it reflects a fully qualified domain name. For example:

```
suffix          "dc=example,dc=com"
```

The rootdn entry is the Distinguished Name (DN) for a user who is unrestricted by access controls or administrative limit parameters set for operations on the LDAP directory. The rootdn user can be thought of as the root user for the LDAP directory. In the configuration file, change the rootdn line from its default value as in the following example:

```
rootdn          "cn=root,dc=example,dc=com"
```

When populating an LDAP directory over a network, change the rootpw line — replacing the default value with an encrypted password string. To create an encrypted password string, type the following command:

```
slappasswd
```

When prompted, type and then re-type a password. The program prints the resulting encrypted password to the shell prompt.

Next, copy the newly created encrypted password into the /etc/openldap/slapd.conf on one of the rootpw lines and remove the hash mark (#).

When finished, the line should look similar to the following example:

```
rootpw {SSHA}vv2y+i6V6esazrIv70xSSnNAJE18bb2u
```

Warning

LDAP passwords, including the rootpw directive specified in /etc/openldap/slapd.conf, are sent over the network unencrypted, unless TLS encryption is enabled.

To enable TLS encryption, review the comments in /etc/openldap/slapd.conf and refer to the man page for slapd.conf.

For added security, the rootpw directive should be commented out after populating the LDAP directory by preceding it with a hash mark (#).

When using the /usr/sbin/slapadd command line tool locally to populate the LDAP directory, use of the rootpw directive is not necessary.

Important

Only the root user can use /usr/sbin/slapadd. However, the directory server runs as the ldap user. Therefore, the directory server is unable to modify any files created by slapadd. To correct this issue, after using slapadd, type the following command:

```
chown -R ldap /var/lib/ldap
```

26.7. Configuring a System to Authenticate Using OpenLDAP

This section provides a brief overview of how to configure OpenLDAP user authentication. Unless you are an OpenLDAP expert, more documentation than is provided here is necessary. Refer to the references provided in 26.9절. "Additional Resources" for more information.

Install the Necessary LDAP Packages.

First, make sure that the appropriate packages are installed on both the LDAP server and the LDAP client machines. The LDAP server needs the openldap-servers package.

The openldap, openldap-clients, and nss_ldap packages need to be installed on all LDAP client machines.

Edit the Configuration Files.

- On the server, edit the /etc/openldap/slapd.conf file on the LDAP server to make sure it matches the specifics of the organization. Refer to 26.6.1절. "Editing /etc/openldap/slapd.conf" for instructions about editing slapd.conf.

- On the client machines, both /etc/ldap.conf and /etc/openldap/ldap.conf need to contain the proper server and search base information for the organization.

  To do this, run the graphical Authentication Configuration Tool (system-config-authentication) and select Enable LDAP Support under the User Information tab.

  It is also possible to edit these files by hand.

- On the client machines, the /etc/nsswitch.conf must be edited to use LDAP.

  To do this, run the Authentication Configuration Tool (system-config-authentication) and select Enable LDAP Support under the User Information tab.

  If editing /etc/nsswitch.conf by hand, add ldap to the appropriate lines.

  For example:

  ```
  passwd: files ldap
  shadow: files ldap
  group: files ldap
  ```

26.7.1. PAM and LDAP

To have standard PAM-enabled applications use LDAP for authentication, run the Authentication Configuration Tool (system-config-authentication) and select Enable LDAP Support under the Authentication tab. For more about configuring PAM, refer to 46.4절. "PAM (Pluggable Authentication Modules)" and the PAM man pages.

26.7.2. Migrating Old Authentication Information to LDAP Format

The /usr/share/openldap/migration/ directory contains a set of shell and Perl scripts for migrating authentication information into an LDAP format.

Note

Perl must be installed on the system to use these scripts.

First, modify the migrate_common.ph file so that it reflects the correct domain. The default DNS domain should be changed from its default value to something like:

```
$DEFAULT_MAIL_DOMAIN = "example";
```

The default base should also be changed to something like:

```
$DEFAULT_BASE = "dc=example,dc=com";
```

The job of migrating a user database into a format that is LDAP readable falls to a group of migration scripts installed in the same directory. Using 표 26.1. “LDAP Migration Scripts” , decide which script to run to migrate the user database.

Run the appropriate script based on the existing name service.

The README and the migration-tools.txt files in the /usr/share/openldap/migration/ directory provide more details on how to migrate the information.

표 26.1. LDAP Migration Scripts

Existing name service	Is LDAP running?	Script to Use
/etc flat files	yes	migrate_all_online.sh
/etc flat files	no	migrate_all_offline.sh
NetInfo	yes	migrate_all_netinfo_online.sh
NetInfo	no	migrate_all_netinfo_offline.sh
NIS (YP)	yes	migrate_all_nis_online.sh
NIS (YP)	no	migrate_all_nis_offline.sh

26.8. Migrating Directories from Earlier Releases

With Red Hat Enterprise Linux, OpenLDAP uses Sleepycat Software's Berkeley DB system as its on-disk storage format for directories. Earlier versions of OpenLDAP used GNU Database Manager (gdbm). For this reason, before upgrading an LDAP implementation to Red Hat Enterprise Linux 5.8, original LDAP data should first be exported before the upgrade, and then reimported afterwards. This can be achieved by performing the following steps:

1. Before upgrading the operating system, run the command /usr/sbin/slapcat -l ldif-output. This outputs an LDIF file called ldif-output containing the entries from the LDAP directory.

2. Upgrade the operating system, being careful not to reformat the partition containing the LDIF file.

3. Re-import the LDAP directory to the upgraded Berkeley DB format by executing the command /usr/sbin/slapadd -l ldif-output.

26.9. Additional Resources

The following resources offer additional information on LDAP. It is highly recommended that you review these, especially the OpenLDAP website and the LDAP HOWTO, before configuring LDAP on your system(s).

26.9.1. Installed Documentation

- /usr/share/docs/openldap-<versionnumber>/ directory — Contains a general README document and miscellaneous information.
- LDAP related man pages — There are a number of man pages for the various applications and configuration files involved with LDAP. The following is a list of some of the more important man pages.

  Client Applications

  - man ldapadd — Describes how to add entries to an LDAP directory.
  - man ldapdelete — Describes how to delete entries within an LDAP directory.
  - man ldapmodify — Describes how to modify entries within an LDAP directory.
  - man ldapsearch — Describes how to search for entries within an LDAP directory.
  - man ldappasswd — Describes how to set or change the password of an LDAP user.
  - man ldapcompare — Describes how to use the ldapcompare tool.
  - man ldapwhoami — Describes how to use the ldapwhoami tool.
  - man ldapmodrdn — Describes how to modify the RDNs of entries.

  Server Applications

  - man slapd — Describes command line options for the LDAP server.
  - man slurpd — Describes command line options for the LDAP replication server.

  Administrative Applications

  - man slapadd — Describes command line options used to add entries to a slapd database.
  - man slapcat — Describes command line options used to generate an LDIF file from a slapd database.
  - man slapindex — Describes command line options used to regenerate an index based upon the contents of a slapd database.
  - man slappasswd — Describes command line options used to generate user passwords for LDAP directories.

  Configuration Files

  - man ldap.conf — Describes the format and options available within the configuration file for LDAP clients.
  - man slapd.conf — Describes the format and options available within the configuration file referenced by both the LDAP server applications (slapd and slurpd) and the LDAP administrative tools (slapadd, slapcat, and slapindex).

26.9.2. Useful Websites

- http://www.openldap.org/[1] — Home of the OpenLDAP Project. This website contains a wealth of information about configuring OpenLDAP as well as a future roadmap and version changes.

[1] http://www.openldap.org

- http://www.padl.com/[2] — Developers of nss_ldap and pam_ldap, among other useful LDAP tools.
- http://www.kingsmountain.com/ldapRoadmap.shtml — Jeff Hodges' LDAP Road Map contains links to several useful FAQs and emerging news concerning the LDAP protocol.
- http://www.ldapman.org/articles/ — Articles that offer a good introduction to LDAP, including methods to design a directory tree and customizing directory structures.

26.9.3. Related Books

- OpenLDAP by Example by John Terpstra and Benjamin Coles; Prentice Hall.
- Implementing LDAP by Mark Wilcox; Wrox Press, Inc.
- Understanding and Deploying LDAP Directory Services by Tim Howes et al.; Macmillan Technical Publishing.

[2] http://www.padl.com

인증 설정

When a user logs in to a Red Hat Enterprise Linux system, the username and password combination must be verified, or authenticated, as a valid and active user. Sometimes the information to verify the user is located on the local system, and other times the system defers the authentication to a user database on a remote system.

The Authentication Configuration Tool provides a graphical interface for configuring user information retrieval from NIS, LDAP, and Hesiod servers. This tool also allows you to configure LDAP, Kerberos, and SMB as authentication protocols.

알림

If you configured a medium or high security level during installation (or with the Security Level Configuration Tool), then the firewall will prevent NIS (Network Information Service) authentication.

This chapter does not explain each of the different authentication types in detail. Instead, it explains how to use the Authentication Configuration Tool to configure them.

To start the graphical version of the Authentication Configuration Tool from the desktop, select the System (on the panel) > Administration > Authentication or type the command system-config-authentication at a shell prompt (for example, in an XTerm or a GNOME terminal).

중요

인증 프로그램을 종료하시면, 변경 사항이 즉시 적용될 것입니다.

27.1. User Information

The User Information tab allows you to configure how users should be authenticated, and has several options. To enable an option, click the empty checkbox beside it. To disable an option, click the checkbox beside it to clear the checkbox. Click OK to exit the program and apply the changes.

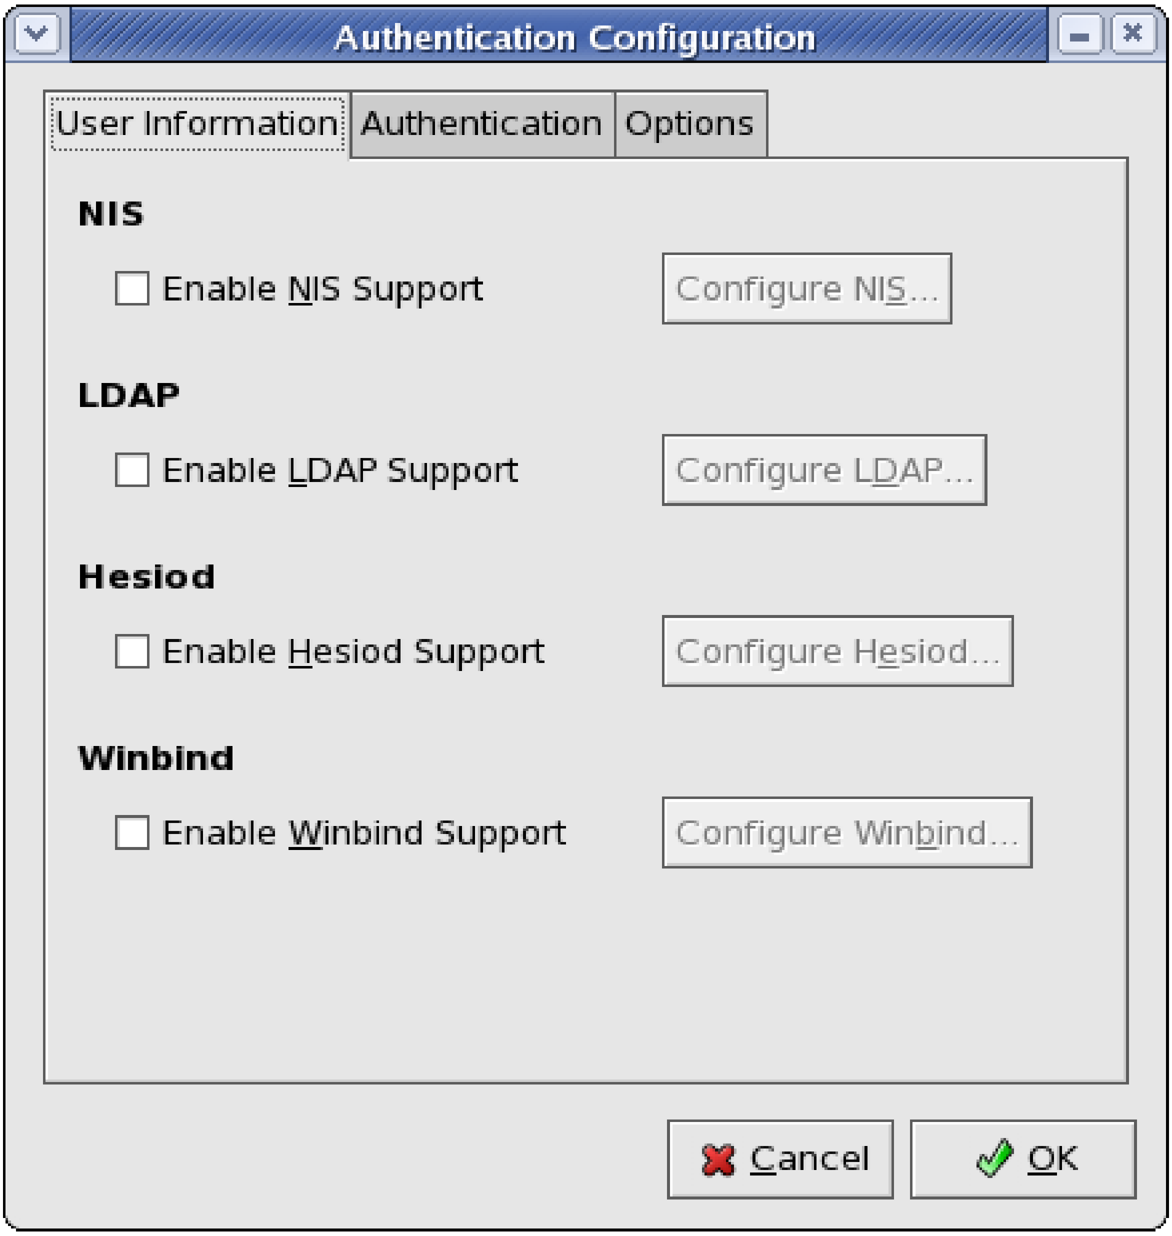

NIS

□ Enable NIS Support

LDAP

□ Enable LDAP Support

Hesiod

□ Enable Hesiod Support

Winbind

□ Enable Winbind Support

그림 27.1. User Information

다음 목록에서는 각 옵션 설정 사항에 대하여 설명하고 있습니다:

NIS

The Enable NIS Support option configures the system to connect to an NIS server (as an NIS client) for user and password authentication. Click the Configure NIS... button to specify the NIS domain and NIS server. If the NIS server is not specified, the daemon attempts to find it via broadcast.

The ypbind package must be installed for this option to work. If NIS support is enabled, the portmap and ypbind services are started and are also enabled to start at boot time.

For more information about NIS, refer to 46.2.3절. “NIS 보안 강화” .

LDAP

The Enable LDAP Support option instructs the system to retrieve user information via LDAP. Click the Configure LDAP... button to specify the following:

- LDAP Search Base DN — Specifies that user information should be retrieved using the listed Distinguished Name (DN).

- LDAP Server — Specifies the IP address of the LDAP server.

- Use TLS to encrypt connections — When enabled, Transport Layer Security will be used to encrypt passwords sent to the LDAP server. The Download CA Certificate option allows you to specify a URL from which to download a valid CA (Certificate Authority) Certificate. A valid CA Certificate must be in PEM (Privacy Enhanced Mail) format.

  For more information about CA Certificates, refer to 23.8.2절. "An Overview of Certificates and Security" .

The openldap-clients package must be installed for this option to work.

For more information about LDAP, refer to 26장. Lightweight Directory Access Protocol (LDAP).

Hesiod

The Enable Hesiod Support option configures the system to retrieve information (including user information) from a remote Hesiod database. Click the Configure Hesiod... button to specify the following:

- Hesiod LHS — Specifies the domain prefix used for Hesiod queries.

- Hesiod RHS — Specifies the default Hesiod domain.

The hesiod package must be installed for this option to work.

For more information about Hesiod, refer to its man page using the command man hesiod. You can also refer to the hesiod.conf man page (man hesiod.conf) for more information on LHS and RHS.

Winbind

The Enable Winbind Support option configures the system to connect to a Windows Active Directory or a Windows domain controller. User information from the specified directory or domain controller can then be accessed, and server authentication options can be configured. Click the Configure Winbind... button to specify the following:

- Winbind Domain — Specifies the Windows Active Directory or domain controller to connect to.

- Security Model — Allows you to select a security model, which configures how clients should respond to Samba. The drop-down list allows you select any of the following:
  - user — This is the default mode. With this level of security, a client must first log in with a valid username and password. Encrypted passwords can also be used in this security mode.

  - server — In this mode, Samba will attempt to validate the username/password by authenticating it through another SMB server (for example, a Windows NT Server). If the attempt fails, the user mode will take effect instead.

  - domain — In this mode, Samba will attempt to validate the username/password by authenticating it through a Windows NT Primary or Backup Domain Controller, similar to how a Windows NT Server would.
  - ads — This mode instructs Samba to act as a domain member in an Active Directory Server (ADS) realm. To operate in this mode, the krb5-server package must be installed, and Kerberos must be configured properly.

- Winbind ADS Realm — When the ads Security Model is selected, this allows you to specify the ADS Realm the Samba server should act as a domain member of.

- Winbind Domain Controllers — Use this option to specify which domain controller winbind should use. For more information about domain controllers, please refer to 21.6.3절. "Domain Controller" .

- Template Shell — When filling out the user information for a Windows NT user, the winbindd daemon uses the value chosen here to to specify the login shell for that user.

For more information about the winbind service, refer to winbindd under 21.2절. "Samba Daemons and Related Services" .

27.2. Authentication

The Authentication tab allows for the configuration of network authentication methods. To enable an option, click the empty checkbox beside it. To disable an option, click the checkbox beside it to clear the checkbox.

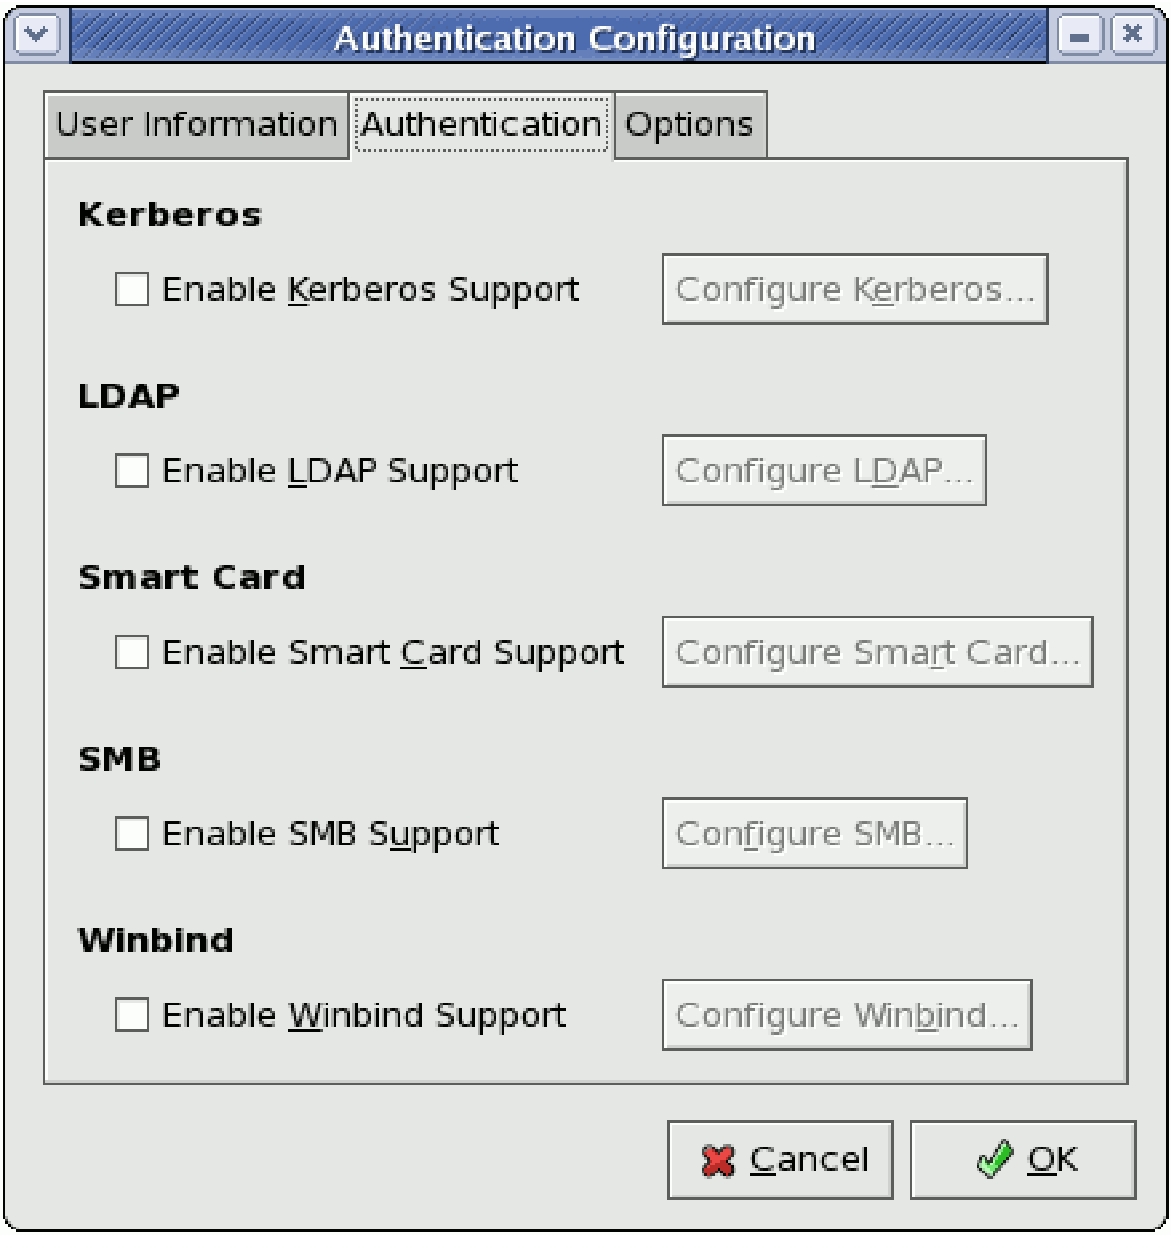

Kerberos

☐ Enable Kerberos Support

LDAP

☐ Enable LDAP Support

Smart Card

☐ Enable Smart Card Support

SMB

☐ Enable SMB Support

Winbind

☐ Enable Winbind Support

그림 27.2. Authentication

다음은 각 옵션 설정 사항에 대하여 설명합니다:

Kerberos

The Enable Kerberos Support option enables Kerberos authentication. Click the Configure Kerberos... button to open the Kerberos Settings dialogue and configure the following:

- Realm — Configures the realm for the Kerberos server. The realm is the network that uses Kerberos, composed of one or more KDCs and a potentially large number of clients.
- KDC — Defines the Key Distribution Center (KDC), which is the server that issues Kerberos tickets.
- Admin Servers — Specifies the administration server(s) running kadmind.

The Kerberos Settings dialogue also allows you to use DNS to resolve hosts to realms and locate KDCs for realms.

The krb5-libs and krb5-workstation packages must be installed for this option to work. For more information about Kerberos, refer to 46.6절. “Kerberos” .

LDAP

The Enable LDAP Support option instructs standard PAM-enabled applications to use LDAP for authentication. The Configure LDAP... button allows you to configure LDAP support with options identical to those present in Configure LDAP... under the User Information tab. For more information about these options, refer to 27.1절. “User Information” .

The openldap-clients package must be installed for this option to work.

Smart Card

The Enable Smart Card Support option enables Smart Card authentication. This allows users to log in using a certificate and key associated stored on a smart card. Click the Configure Smart Card... button for more options.

The pam_pkcs11 and coolkey packages must be installed for this option to work. For more information about smart cards, refer to 46.3.1.3절. “Supported Smart Cards” under 46.3절. “Single Sign-on (SSO)” .

SMB

The Enable SMB Support option configures PAM to use a Server Message Block (SMB) server to authenticate users. SMB refers to a client-server protocol used for cross-system communication; it is also the protocol used by Samba to appear as a Windows server to Windows clients. Click the Configure SMB... button to specify the following:

- Workgroup — Specifies the SMB workgroup to use.
- Domain Controllers — Specifies the SMB domain controllers to use.

Winbind

The Enable Winbind Support option configures the system to connect to a Windows Active Directory or a Windows domain controller. User information from the specified directory or domain controller can then be accessed, and server authentication options can be configured.

The Configure Winbind... options are identical to those in the Configure Winbind... button on the User Information tab. Please refer to Winbind (under 27.1절. “User Information”) for more information.

27.3. Options

This tab allows other configuration options, as listed below.

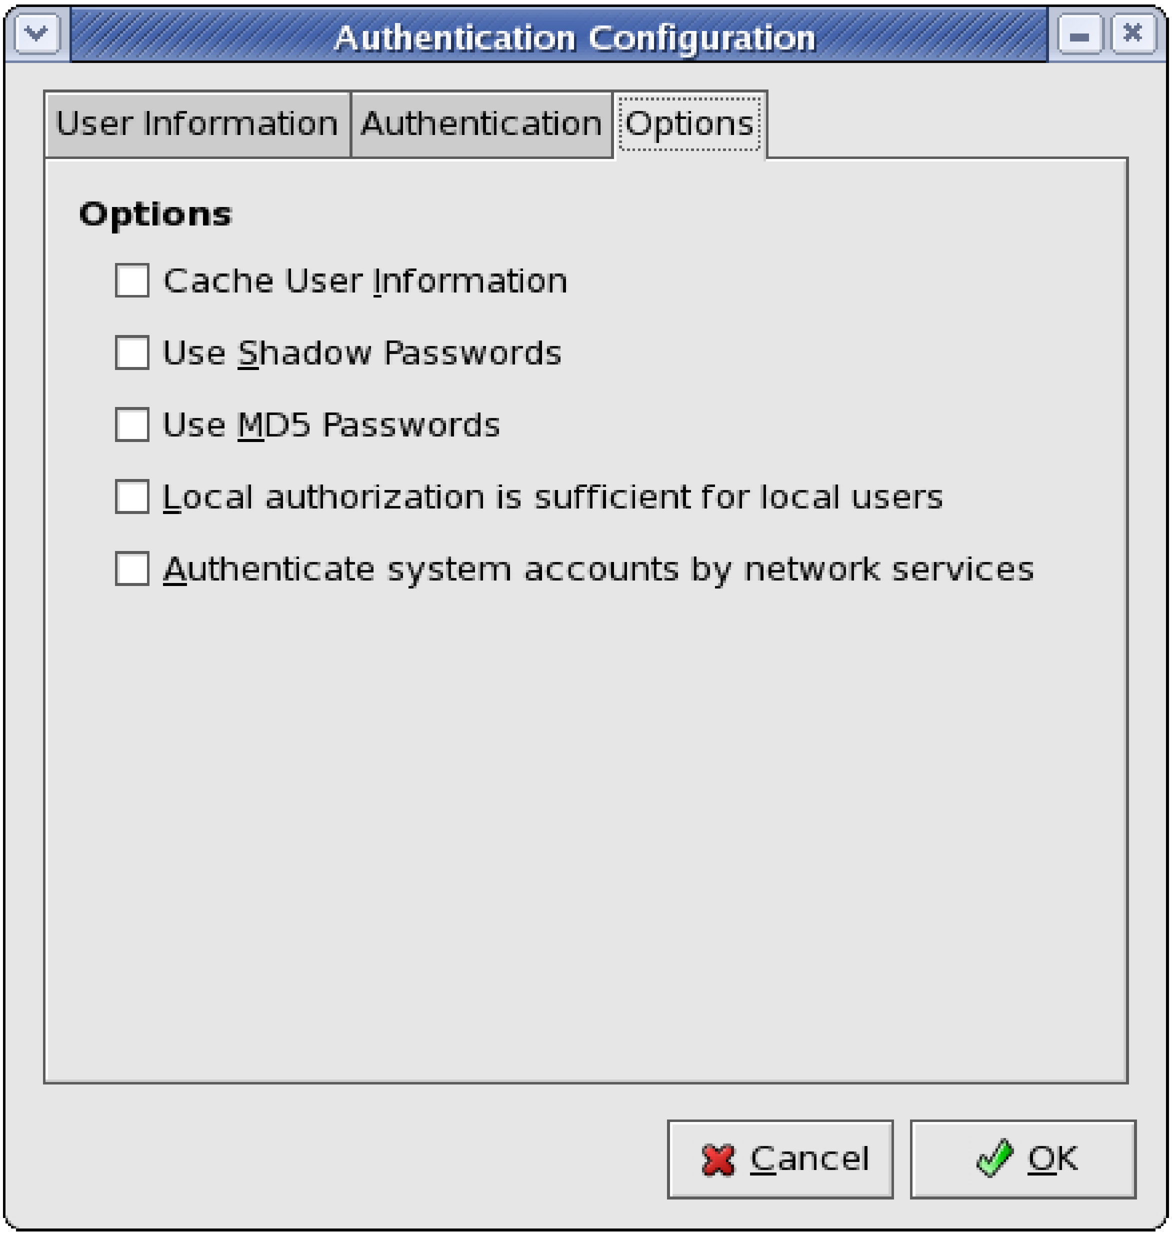

Options

□ Cache User Information

□ Use Shadow Passwords

□ Use MD5 Passwords

□ Local authorization is sufficient for local users

□ Authenticate system accounts by network services

그림 27.3. Options

Cache User Information

Select this option to enable the name service cache daemon (nscd) and configure it to start at boot time.

The nscd package must be installed for this option to work. For more information about nscd, refer to its man page using the command man nscd.

Use Shadow Passwords

Select this option to store passwords in shadow password format in the /etc/shadow file instead of /etc/passwd. Shadow passwords are enabled by default during installation and are highly recommended to increase the security of the system.

The shadow-utils package must be installed for this option to work. For more information about shadow passwords, refer to 35.6절. "Shadow Passwords" .

Use MD5 Passwords

Select this option to enable MD5 passwords, which allows passwords to be up to 256 characters instead of eight characters or less. It is selected by default during installation and is highly recommended for increased security.

Local authorization is sufficient for local users

When this option is enabled, the system will not check authorization from network services (such as LDAP or Kerberos) for user accounts maintained in its /etc/passwd file.

Authenticate system accounts by network services

Enabling this option configures the system to allow network services (such as LDAP or Kerberos) to authenticate system accounts (including root) in the machine.

27.4. 명령행 버전

The Authentication Configuration Tool can also be run as a command line tool with no interface. The command line version can be used in a configuration script or a kickstart script. The authentication options are summarized in 표 27.1. "명령행 옵션" .

Tip

These options can also be found in the authconfig man page or by typing authconfig --help at a shell prompt.

표 27.1. 명령행 옵션

옵션	설명
--enableshadow	섀도우 암호 활성화
--disableshadow	섀도우 암호 비활성화
--enablemd5	MD5 암호 활성화
--disablemd5	MD5 암호 비활성화
--enablenis	NIS 활성화
--disablenis	NIS 비활성화
--nisdomain=<domain>	NIS 도메인 지정
--nisserver=<server>	NIS 서버 지정
--enableldap	사용자 정보를 인증하는데 LDAP을 사용

옵션	설명
--disableldap	사용자 정보를 인증하는데 LDAP을 사용하지 않음
--enableldaptls	LDAP에 TLS를 사용함
--disableldaptls	LDAP에 TLS를 사용하지 않음
--enableldapauth	인증에 LDAP을 사용함
--disableldapauth	인증에 LDAP을 사용하지 않음
--ldapserver=<server>	LDAP 서버 지정
--ldapbasedn=<dn>	LDAP 기반 DN 지정
--enablekrb5	Kerberos 활성화
--disablekrb5	Kerberos 비활성화
--krb5kdc=<kdc>	Kerberos KDC 지정
--krb5adminserver=<server>	Kerberos 관리 서버 지정
--krb5realm=<realm>	Kerberos 관리 영역 지정
--enablekrb5kdcdns	Enable use of DNS to find Kerberos KDCs
--disablekrb5kdcdns	Disable use of DNS to find Kerberos KDCs
--enablekrb5realmdns	Enable use of DNS to find Kerberos realms
--disablekrb5realmdns	Disable use of DNS to find Kerberos realms
--enablesmbauth	SMB 활성화
--disablesmbauth	SMB 비활성화
--smbworkgroup=<workgroup>	SMB 작업그룹 지정
--smbservers=<server>	SMB 서버 지정
--enablewinbind	Enable winbind for user information by default
--disablewinbind	Disable winbind for user information by default
--enablewinbindauth	Enable winbindauth for authentication by default
--disablewinbindauth	Disable winbindauth for authentication by default
--smbsecurity=<user|server|domain|ads>	Security mode to use for Samba and winbind
--smbrealm=<STRING>	Default realm for Samba and winbind when security=ads
--smbidmapuid=<lowest-highest>	UID range winbind assigns to domain or ADS users
--smbidmapgid=<lowest-highest>	GID range winbind assigns to domain or ADS users

옵션	설명
--winbindseparator=<\>	Character used to separate the domain and user part of winbind usernames if winbindusedefaultdomain is not enabled
--winbindtemplatehomedir=</home/%D/%U>	Directory that winbind users have as their home
--winbindtemplateprimarygroup=<nobody>	Group that winbind users have as their primary group
--winbindtemplateshell=</bin/false>	Shell that winbind users have as their default login shell
--enablewinbindusedefaultdomain	Configures winbind to assume that users with no domain in their usernames are domain users
--disablewinbindusedefaultdomain	Configures winbind to assume that users with no domain in their usernames are not domain users
--winbindjoin=<Administrator>	Joins the winbind domain or ADS realm now as this administrator
--enablewins	Enable WINS for hostname resolution
--disablewins	Disable WINS for hostname resolution
--enablehesiod	Hesiod 활성화
--disablehesiod	Hesiod 비활성화
--hesiodlhs=<lhs>	Hesiod LHS 지정
--hesiodrhs=<rhs>	Hesiod RHS 지정
--enablecache	Enable nscd
--disablecache	Disable nscd
--nostart	Do not start or stop the portmap, ypbind, or nscd services even if they are configured
--kickstart	사용자 인터페이스를 표시하지 않음
--probe	네트워크 디폴트를 검색하여 표시

Using and Caching Credentials with SSSD

The System Security Services Daemon (SSSD) provides access to different identity and authentication providers. SSSD is an intermediary between local clients and any configured data store. The local clients connect to SSSD and then SSSD contacts the external providers. This brings a number of benefits for administrators:

- Reducing the load on identification/authentication servers. Rather than having every client service attempt to contact the identification server directly, all of the local clients can contact SSSD which can connect to the identification server or check its cache.

- Permitting offline authentication. SSSD can optionally keep a cache of user identities and credentials that it retrieves from remote services. This allows users to authenticate to resources successfully, even if the remote identification server is offline or the local machine is offline.

- Using a single user account. Remote users frequently have two (or even more) user accounts, such as one for their local system and one for the organizational system. This is necessary to connect to a virtual private network (VPN). Because SSSD supports caching and offline authentication, remote users can connect to network resources simply by authenticating to their local machine and then SSSD maintains their network credentials.

The System Security Services Daemon does not require any additional configuration or tuning to work with the Authentication Configuration Tool. However, SSSD can work with other applications, and the daemon may require configuration changes to improve the performance of those applications.

28.1. About the sssd.conf File

SSSD services and domains are configured in a .conf file. The default file is /etc/sssd/sssd.conf, although alternative files can be passed to SSSD by using the -c option with the sssd command:

```
# sssd -c /etc/sssd/customfile.conf
```

Both services and domains are configured individually, in separate sections on the configuration identified by [type/name] divisions, such as [domain/LDAP]. The configuration file uses simple key = value lines to set the configuration. Comment lines are set by either a hash sign (#) or a semicolon (;)

For example:

```
[section]
# Comment line
key1 = val1
key10 = val1,val2
```

28.2. Starting and Stopping SSSD

참고

Configure at least one domain before starting SSSD for the first time. See 28.4절. “Creating Domains” .

Either the service command or the /etc/init.d/sssd script can start SSSD. For example:

```
# service sssd start
```

By default, SSSD is configured not to start automatically. To change this behavior, use the chkconfig command:

```
[root@server ~]# chkconfig sssd on
```

28.3. Configuring Services

SSSD worked with specialized services that run in tandem with the SSSD process itself. SSSD and its associated services are configured in the sssd.conf file. on sections. The [sssd] section also lists the services that are active and should be started when sssd starts within the services directive.

SSSD currently provides several services:

- An NSS provider service that answers NSS requests from the sssd_nss module. This is configured in the [nss] section of the configuration.

- A PAM provider service that manages a PAM conversation through the sssd_pam PAM module. This is configured in the [pam] section of the configuration.

- monitor, a special service that monitors and starts or restarts all other SSSD services. Its options are specified in the [sssd] section of the /etc/sssd/sssd.conf configuration file.

참고

If a DNS lookup fails to return an IPv4 address for a hostname, SSSD attempts to look up an IPv6 address before returning a failure. This only ensures that the asynchronous resolver identifies the correct address.

The hostname resolution behavior is configured in the lookup family order option in the sssd.conf configuration file.

28.3.1. Configuring the NSS Service

SSSD provides an NSS module, sssd_nss, which instructs the system to use SSSD to retrieve user information. The NSS configuration must include a reference to the SSSD module, and then the SSSD configuration sets how SSSD interacts with NSS.

To configure the NSS service:

1. Open the sssd.conf file.

   ```
   # vim /etc/sssd/sssd.conf
   ```

2. Make sure that NSS is listed as one of the services that works with SSSD.

   ```
   [sssd]
   config_file_version = 2
   reconnection_retries = 3
   sbus_timeout = 30
   ```

```
services = nss, pam
```

3. In the [nss] section, change any of the NSS parameters. These are listed in 표 28.1. “SSSD [nss] Configuration Parameters” .

```
[nss]
filter_groups = root
filter_users = root
reconnection_retries = 3
entry_cache_timeout = 300
entry_cache_nowait_percentage = 75
```

4. Restart SSSD.

```
service sssd restart
```

표 28.1. SSSD [nss] Configuration Parameters

Parameter	Value Format	Description
enum_cache_timeout	integer	Specifies how long, in seconds, sssd_nss should cache requests for information about all users (enumerations).
entry_cache_nowait_percentage	integer	Specifies how long sssd_nss should return cached entries before refreshing the cache. Setting this to zero (0) disables the entry cache refresh. This configures the entry cache to update entries in the background automatically if they are requested if the time before the next update is a certain percentage of the next interval. For example, if the interval is 300 seconds and the cache percentage is 75, then the entry cache will begin refreshing when a request comes in at 225 seconds — 75% of the interval. The allowed values for this option are 0 to 99, which sets the percentage based on the entry_cache_timeout value. The default value is 50%.
entry_negative_timeout	integer	Specifies how long, in seconds, sssd_nss should cache negative cache hits. A negative cache hit is a query for an invalid database entries, including non-existent entries.

Parameter	Value Format	Description
filter_users, filter_groups	string	Tells SSSD to exclude certain users from being fetched from the NSS database. This is particularly useful for system accounts such as root.
filter_users_in_groups	Boolean	Sets whether users listed in the filter_users list appear in group memberships when performing group lookups. If set to FALSE, group lookups return all users that are members of that group. If not specified, this value defaults to TRUE, which filters the group member lists.

28.3.2. Configuring the PAM Service

> **주의**
>
> A mistake in the PAM configuration file can lock users out of the system completely. Always back up the configuration files before performing any changes, and keep a session open so that any changes can be reverted.

SSSD provides a PAM module, sssd_pam, which instructs the system to use SSSD to retrieve user information. The PAM configuration must include a reference to the SSSD module, and then the SSSD configuration sets how SSSD interacts with PAM.

To configure the PAM service:

1. The Authentication Configuration tool automatically writes to the /etc/pam.d/system-auth-ac file, which is symlinked to /etc/pam.d/system-auth. Any changes made to /etc/pam.d/system-auth are overwritten the next time that authconfig is run.

   So, remove the /etc/pam.d/system-auth symlink.

   ```
   [root@server ~]# rm /etc/pam.d/system-auth
   rm: remove symbolic link `/etc/pam.d/system-auth'? y
   ```

2. Create a new /etc/pam.d/system-auth-local file. One easy way to do this is simply to copy the /etc/pam.d/system-auth-ac file.

   ```
   [root@server ~]# cp /etc/pam.d/system-auth-ac /etc/pam.d/system-auth-local
   ```

3. Create a new symlink between the /etc/pam.d/system-auth-local file and /etc/pam.d/system-auth.

   ```
   [root@server ~]# ln -s /etc/pam.d/system-auth-local /etc/pam.d/system-auth
   ```

4. Edit the /etc/pam.d/system-auth-local file, and add all of the SSSD modules to the PAM configuration:

```
#%PAM-1.0
...
auth        sufficient    pam_sss.so use_first_pass
auth        required       pam_deny.so

...
account [default=bad success=ok user_unknown=ignore] pam_sss.so
account     required       pam_permit.so

...
password    sufficient    pam_sss.so use_authtok
password    required       pam_deny.so

...
session     sufficient    pam_sss.so
session     required       pam_unix.so
```

These modules can be set to include statements, as necessary.

5. Open the sssd.conf file.

```
# vim /etc/sssd/sssd.conf
```

6. Make sure that PAM is listed as one of the services that works with SSSD.

```
[sssd]
config_file_version = 2
reconnection_retries = 3
sbus_timeout = 30
services = nss, pam
```

7. In the [pam] section, change any of the PAM parameters. These are listed in 표 28.2. “SSSD [pam] Configuration Parameters” .

```
[pam]
reconnection_retries = 3
offline_credentials_expiration = 2
offline_failed_login_attempts = 3
offline_failed_login_delay = 5
```

8. Restart SSSD.

```
service sssd restart
```

표 28.2. SSSD [pam] Configuration Parameters

Parameter	Value Format	Description
offline_credentials_expiration	integer	Sets how long, in days, to allow cached logins if the authentication provider is offline. This value is measured from the last successful online login. If not specified, this

Parameter	Value Format	Description
		defaults to zero (0), which is unlimited.
offline_failed_login_attempts	integer	Sets how many failed login attempts are allowed if the authentication provider is offline. If not specified, this defaults to zero (0), which is unlimited.
offline_failed_login_delay	integer	Sets how long to prevent login attempts if a user hits the failed login attempt limit. If set to zero (0), the user cannot authenticate while the provider is offline once he hits the failed attempt limit. Only a successful online authentication can re-enable offline authentication. If not specified, this defaults to five (5).

28.4. Creating Domains

SSSD recognizes domains, which are associated with the different identity servers. Domains are a combination of an identity provider and an authentication method. SSSD works with LDAP identity providers (including OpenLDAP, Red Hat Directory Server, and Microsoft Active Directory) and can use native LDAP authentication or Kerberos authentication.

As long as they belong to different domains, SSSD can recognize different users with the same username. For example, SSSD can successfully authenticate both jsmith in the ldap.example.com domain and jsmith in the ldap.otherexample.com domain. SSSD allows requests using fully-qualified domain names, so requesting information for jsmith@ldap.example.com returns the the proper user account. Specifying only the username returns the user for whichever domain comes first in the lookup order.

Tip

SSSD has a filter_users option, which excludes the specified users from being returned in a search.

Configuring a domain defines both where user information is stored and how those users are allowed to authenticate to the system. The possible combinations are listed in 표 28.3. “Identity Store and Authentication Type Combinations” .

- 28.4.1절. “General Rules and Options for Configuring a Domain”
- 28.4.2절. “Configuring an LDAP Domain”
- 28.4.3절. “Configuring Kerberos Authentication with a Domain”

- 28.4.4절. "Configuring a Proxy Domain"

표 28.3. Identity Store and Authentication Type Combinations

Identification Provider	Authentication Provider
LDAP	LDAP
LDAP	Kerberos
proxy	LDAP
proxy	Kerberos
proxy	proxy

28.4.1. General Rules and Options for Configuring a Domain

A domain configuration defines the identity provider, the authentication provider, and any specific configuration to access the information in those providers. There are two types of identity providers — LDAP and proxy —three types of authentication providers — LDAP, Kerberos, and proxy. The identity and authentication providers can be configured in any combination in a domain entry.

Along with the domain entry itself, the domain name must be added to the list of domains that SSSD will query. For example:

```
domains = LOCAL,Name

[domain/Name]
id_provider = type
auth_provider = type
provider_specific = value
global = value
```

global attributes are available to any type of domain, such as cache and time out settings. Each identity and authentication provider has its own set of required and optional configuration parameters.

표 28.4. General [domain] Configuration Parameters

Parameter	Value Format	Description
id_provider	string	Specifies the data provider identity backend to use for this domain. The supported identity backends are: • ldap • ipa, compatible with FreeIPA version 2.x and Identity Management in Red Hat Enterprise Linux • proxy for a legacy NSS provider, such as nss_nis. Using a proxy ID provider also requires specifying the legacy NSS library to load to start successfully, set in the proxy_lib_name option. • local, the SSSD internal local provider

Parameter	Value Format	Description
auth_provider	string	Sets the authentication provider used for the domain. The default value for this option is the value of id_provider. The supported authentication providers are ldap, ipa, krb5 (Kerberos), proxy, and none.
min_id,max_id	integer	Optional. Specifies the UID and GID range for the domain. If a domain contains entries that are outside that range, they are ignored. The default value for min_id is 1; the default value for max_id is 0, which is unlimited. 중요 The default min_id value is the same for all types of identity provider. If LDAP directories are using UID numbers that start at one, it could cause conflicts with users in the local /etc/passwd file. To avoid these conflicts, set min_id to 1000 or higher as possible.
enumerate	Boolean	Optional. Specifies whether to list the users and groups of a domain. Enumeration means that the entire set of available users and groups on the remote source is cached on the local machine. When enumeration is disabled, users and groups are only cached as they are requested.

Parameter	Value Format	Description
		주의 When enumeration is enabled, reinitializing a client results in a complete refresh of the entire set of available users and groups from the remote source. Similarly, when SSSD is connected to a new server, the entire set of available users and groups from the remote source is pulled and cached on the local machine. In a domain with a large number of clients connected to a remote source, this refresh process can harm the network performance because of frequent queries from the clients. If the set of available users and groups is large enough, it degrades client performance as well. The default value for this parameter is false, which disables enumeration.
cache_credentials	Boolean	Optional. Specifies whether to store user credentials in the local SSSD domain database cache. The default value for this parameter is false. Set this value to true for domains other than the LOCAL domain to enable offline authentication.
entry_cache_timeout	integer	Optional. Specifies how long, in seconds, SSSD should cache positive cache hits. A positive cache hit is a successful query.
use_fully_qualified_names	Boolean	Optional. Specifies whether requests to this domain require fully-qualified domain names. If set to true, all requests to

Parameter	Value Format	Description
		this domain must use fully-qualified domain names. It also means that the output from the request displays the fully-qualified name. Restricting requests to fully-qualified user names allows SSSD to differentiate between domains with users with conflicting usernames. If use_fully_qualified_names is set to false, it is possible to use the fully-qualified name in the requests, but only the simplified version is displayed in the output. SSSD can only parse names based on the domain name, not the realm name. The same name can be used for both domains and realms, however.

28.4.2. Configuring an LDAP Domain

An LDAP domain simply means that SSSD uses an LDAP directory as the identity provider (and, optionally, also as an authentication provider). SSSD supports several major directory services:

- Red Hat Directory Server
- OpenLDAP
- Microsoft Active Directory 2003 and 2003R2, with Services for UNIX
- Microsoft Active Directory 2003 and 2003R2, with Subsystem for UNIX-based Applications

참고

DNS service discovery allows the LDAP backend to find the appropriate DNS servers to connect to automatically using a special DNS query.

- 28.4.2.1절. “Parameters for Configuring an LDAP Domain”
- 28.4.2.2절. “LDAP Domain Examples”
- 28.4.2.3절. “Using IP Addresses in Certificate Subject Names”

28.4.2.1. Parameters for Configuring an LDAP Domain

An LDAP directory can function as both an identity provider and an authentication provider. The configuration requires enough information to identify and connect to the user directory in the LDAP server. Other options are available to provide more fine-grained control, like specifying a user account to use to connect to the LDAP server or using different LDAP servers for password operations. The most common options are listed in 표 28.5. "LDAP Domain Configuration Parameters" . All of the options listed in 28.4.1절. "General Rules and Options for Configuring a Domain" are also available for LDAP domains.

Tip

Many other options are listed in the man page for LDAP domain configuration, sssd-ldap(5).

표 28.5. LDAP Domain Configuration Parameters

Parameter	Description
ldap_uri	Gives a comma-separated list of the URIs of the LDAP servers to which SSSD will connect. The list is given in order of preference, so the first server in the list is tried first. Listing additional servers provides failover protection. This can be detected from the DNS SRV records if it is not given.
ldap_search_base	Gives the base DN to use for performing LDAP user operations.
ldap_tls_reqcert	Specifies how to check for SSL server certificates in a TLS session. There are four options: • never disables requests for certificates. • allow requests a certificate, but proceeds normally even if no certificate is given or a bad certificate is given. • try requests a certificate and proceeds normally if no certificate is given, If a bad certificate is given, the session terminates. • demand and hard are the same option. This requires a valid certificate or the session is terminated. The default is hard.
ldap_tls_cacert	Gives the full path and file name to the file that contains the CA certificates for all of the CAs that SSSD recognizes. SSSD will accept any certificate issued by these CAs. This uses the OpenLDAP system defaults if it is not given explicitly.

Parameter	Description
ldap_referrals	Sets whether SSSD will use LDAP referrals, meaning forwarding queries from one LDAP database to another. SSSD supports database-level and subtree referrals. For referrals within the same LDAP server, SSSD will adjust the DN of the entry being queried. For referrals that go to different LDAP servers, SSSD does an exact match on the DN. Setting this value to true enables referrals; by default, referrals are enabled.

ldap_schema	Sets what version of schema to use when searching for user entries. This can be either rfc2307 or rfc2307bis. The default is rfc2307.
	In RFC 2307, group objects use a multi-valued attribute, memberuid, which lists the names of the users that belong to that group. In RFC 2307bis, group objects use the member attribute, which contains the full distinguished name (DN) of a user or group entry. RFC 2307bis allows nested groups usning the member attribute. Because these different schema use different definitions for group membership, using the wrong LDAP schema with SSSD can affect both viewing and managing network resources, even if the appropriate permissions are in place.
	For example, with RFC 2307bis, all groups are returned when using nested groups or primary/secondary groups.
	`$ id`
	`uid=500(myserver) gid=500(myserver)`
	` groups=500(myserver),510(myothergroup)`
	If SSSD is using RFC 2307 schema, only the primary group is returned.
	This setting only affects how SSSD determines the group members. It does not change the actual user data.

ldap_search_timeout	Sets the time, in seconds, that LDAP searches are allowed to run before they are canceled and cached results are returned. This defaults to five when the enumerate value is false and defaults to 30 when enumerate is true. When an LDAP search times out, SSSD automatically switches to offline mode.
ldap_network_timeout	Sets the time, in seconds, SSSD attempts to poll an LDAP server after a connection attempt fails. The default is six seconds.

Parameter	Description
ldap_opt_timeout	Sets the time, in seconds, to wait before aborting synchronous LDAP operations if no response is received from the server. This option also controls the timeout when communicating with the KDC in case of a SASL bind. The default is five seconds.

28.4.2.2. LDAP Domain Examples

The LDAP configuration is very flexible, depending on your specific environment and how general or specific you need the SSSD behavior to be. These are some common examples of an LDAP domain, but the SSSD configuration is not limited to these examples.

참고

Along with creating the domain entry, add the new domain to the list of domains for SSSD to query in the sssd.conf file. For example:

```
domains = LOCAL,LDAP1,AD,PROXYNIS
```

예 28.1. A Basic LDAP Domain Configuration

An LDAP domain requires three things:

- An LDAP server
- The search base
- A way to establish a secure connection

The last item depends on the LDAP environment. SSSD requires a secure connection since it handles sensitive information. This connection can be a dedicated TLS/SSL connection or it can use Start TLS.

Using a dedicated TLS/SSL connection simply uses an LDAPS connection to connect to the server and is therefore set as part of the ldap_uri option:

```
# An LDAP domain
[domain/LDAP]
enumerate = false
cache_credentials = TRUE

id_provider = ldap
auth_provider = ldap

ldap_uri = ldaps://ldap.example.com:636
ldap_search_base = dc=example,dc=com
```

Using Start TLS requires a way to input the certificate information to establish a secure connection dynamically over an insecure port. This is done using the ldap_id_use_start_tls option to use Start TLS and then ldap_tls_cacert to identify the CA certificate which issued the SSL server certificates.

```
# An LDAP domain
```

```
[domain/LDAP]
enumerate = false
cache_credentials = TRUE

id_provider = ldap
auth_provider = ldap

ldap_uri = ldap://ldap.example.com
ldap_search_base = dc=example,dc=com
ldap_id_use_start_tls = True
ldap_tls_reqcert = demand
ldap_tls_cacert = /etc/pki/tls/certs/ca-bundle.crt
```

To configure any Active Directory server as an LDAP domain requires two things:

- Installing Windows Services for UNIX (2003 and 2003 R2) or the Subsystem for UNIX-based Applications (2008).

> **참고**
>
> Services for Unix is not supported on 64-bit operating systems.

- Running the cacertdir_rehash function to create the appropriate symlinks.

예 28.2. An Active Directory 2003 Domain

As with an OpenLDAP or Directory Server domain, Active Directory requires the search base and the LDAP URI of the Active Directory server, but SSSD requires more information about directory entries and the user account to use to connect because of the differences between an Active Directory-style database and an OpenLDAP/Directory Server-style database.

These options are described in the man page for LDAP domain configuration, sssd-ldap(5).

```
# Example LDAP domain where the LDAP server is an Active Directory 2003 server.

[domain/AD]
description = LDAP domain with AD server
enumerate = false
;
id_provider = ldap
auth_provider = ldap
ldap_uri = ldap://your.ad.server.com
ldap_schema = rfc2307bis
ldap_search_base = dc=example,dc=com
ldap_default_bind_dn = cn=Administrator,cn=Users,dc=example,dc=com
ldap_default_authtok_type = password
ldap_default_authtok = secret
ldap_user_object_class = person
ldap_user_name = msSFU30Name
ldap_user_uid_number = msSFU30UidNumber
ldap_user_gid_number = msSFU30GidNumber
ldap_user_home_directory = msSFU30HomeDirectory
ldap_user_shell = msSFU30LoginShell
ldap_user_principal = userPrincipalName
ldap_group_object_class = group
ldap_group_name = msSFU30Name
```

```
ldap_group_gid_number = msSFU30GidNumber
```

예 28.3. A Basic Active Directory 2003 R2 or 2008 Domain

Configuring a Microsoft Active Directory 2003 R2 or 2008 domain is similar, but not identical, to configuring an Active Directory 2003 domain. Using later Active Directory servers requires less group configuration information.

These options are described in the man page for LDAP domain configuration, sssd-ldap(5).

```
# Example LDAP domain where the LDAP server is an Active Directory 2003 R2 or an Active Directory 2008 server.

[domain/AD]
description = LDAP domain with AD server
; debug_level = 9
enumerate = false

id_provider = ldap
auth_provider = ldap
chpass_provider = ldap

ldap_uri = ldap://your.ad.server.com
ldap_tls_cacertdir = /etc/openldap/cacerts
ldap_tls_cacert = /etc/openldap/cacerts/test.cer
ldap_search_base = dc=example,dc=com
ldap_default_bind_dn = cn=Administrator,cn=Users,dc=example,dc=com
ldap_default_authtok_type = password
ldap_default_authtok = secret
ldap_pwd_policy = none
ldap_user_object_class = user
ldap_group_object_class = group
```

28.4.2.3. Using IP Addresses in Certificate Subject Names

Using an IP address in the ldap_uri option instead of the server name may cause the TLS/SSL connection to fail. TLS/SSL certificates contain the server name, not the IP address. However, the subject alternative name field in the certificate can be used to include the IP address of the server, which allows a successful secure connection using an IP address.

1. Convert an existing certificate into a certificate request. The signing key (-signkey) is the key of the issuer of whatever CA originally issued the certificate. If this is done by an external CA, it requires a separate PEM file; if the certificate is self-signed, then this is the certificate itself. For example:

   ```
   openssl x509 -x509toreq -in old_cert.pem -out req.pem -signkey key.pem
   ```

   With a self-signed certificate:

   ```
   openssl x509 -x509toreq -in old_cert.pem -out req.pem -signkey old_cert.pem
   ```

2. Edit the /etc/pki/tls/openssl.cnf configuration file to include the server's IP address under the [v3_ca] section:

   ```
   subjectAltName = IP:10.0.0.10
   ```

3. Use the generated certificate request to generate a new self-signed certificate with the specified IP address:

```
openssl x509 -req -in req.pem -out new_cert.pem -extfile ./openssl.cnf -extensions v3_ca -signkey old_cert.pem
```

The -extensions option sets which extensions to use with the certificate. For this, it should be v3_ca to load the appropriate section.

4. Copy the private key block from the old_cert.pem file into the new_cert.pem file to keep all relevant information in one file.

When creating a certificate through the certutil utility provided by the nss-utils package, note that certutil supports DNS subject alternative names for certificate creation only.

28.4.3. Configuring Kerberos Authentication with a Domain

Both LDAP and proxy identity providers can use a separate Kerberos domain to supply authentication. Configuring a Kerberos authentication provider requires the key distribution center (KDC) and the Kerberos domain. All of the principal names must be available in the specified identity provider; if they are not, SSSD constructs the principals using the format username@REALM.

> **참고**
>
> Kerberos can only provide authentication; it cannot provide an identity database.

SSSD assumes that the Kerberos KDC is also a Kerberos kadmin server. However, production environments commonly have multiple, read-only replicas of the KDC and only a single kadmin server. Use the krb5_kpasswd option to specify where the password changing service is running or if it is running on a non-default port. If the krb5_kpasswd option is not defined, SSSD tries to use the Kerberos KDC to change the password.

The basic Kerberos configuration options are listed in 표 28.6. "Kerberos Authentication Configuration Parameters" . The sssd-krb5(5) man page has more information about Kerberos configuration options.

예 28.4. Basic Kerberos Authentication

```
# A domain with identities provided by LDAP and authentication by Kerberos
[domain/KRBDOMAIN]
enumerate = false
id_provider = ldap
chpass_provider = krb5
ldap_uri = ldap://ldap.example.com
ldap_search_base = dc=example,dc=com
ldap-tls_reqcert = demand
ldap_tls_cacert = /etc/pki/tls/certs/ca-bundle.crt

auth_provider = krb5
krb5_server = 192.168.1.1, kerberos.example.com
krb5_realm = EXAMPLE.COM
krb5_kpasswd = kerveros.admin.example.com
krb5_auth_timeout = 15
```

표 28.6. Kerberos Authentication Configuration Parameters

Parameter	Description
chpass_provider	Specifies which service to use for password change operations. This is assumed to be the same as the authentication provider. To use Kerberos, set this to krb5.
krb5_server	Gives a comma-separated list of IP addresses or hostnames of Kerberos servers to which SSSD will connect. The list is given in order of preference, so the first server in the list is tried first. Listing additional servers provides failover protection. When using service discovery for KDC or kpasswd servers, SSSD first searches for DNS entries that specify UDP as the connection protocol, and then falls back to TCP.
krb5_realm	Identies the Kerberos realm served by the KDC.
krb5_lifetime	Requests a Kerberos ticket with the specified lifetime in seconds (s), minutes (m), hours (h) or days (d).
krb5_renewable_lifetime	Requests a renewable Kerberos ticket with a total lifetime that is specified in seconds (s), minutes (m), hours (h) or days (d).
krb5_renew_interval	Sets the time, in seconds, for SSSD to check if tickets should be renewed. Tickets are renewed automatically once they exceed half their lifetime. If this option is missing or set to zero, then automatic ticket renewal is disabled.
krb5_store_password_if_offline	Sets whether to store user passwords if the Kerberos authentication provider is offline, and then to use that cache to request tickets when the provider is back online. The default is false, which does not store passwords.
krb5_kpasswd	Lists alternate Kerberos kadmin servers to use if the change password service is not running on the KDC.
krb5_ccname_template	Gives the directory to use to store the user's credential cache. This can be templatized, and the following tokens are supported: • %u, the user's login name • %U, the user's login UID • %p, the user's principal name • %r, the realm name • %h, the user's home directory • %d, the value of the krb5ccache_dir parameter

Parameter	Description
	• %P, the process ID of the SSSD client. • %%, a literal percent sign (%) • XXXXXX, a string at the end of the template which instructs SSSD to create a unique filename safely For example: `krb5_ccname_template = FILE:%d/krb5cc_%U_XXXXXX`
krb5_ccachedir	Specifies the directory to store credential caches. This can be templatized, using the same tokens as krb5_ccname_template, except for %d and %P. If %u, %U, %p, or %h are used, then SSSD creates a private directory for each user; otherwise, it creates a public directory.
krb5_auth_timeout	Gives the time, in seconds, before an online authentication or change password request is aborted. If possible, the authentication request is continued offline. The default is 15 seconds.

28.4.4. Configuring a Proxy Domain

A proxy with SSSD is just a relay, an intermediary configuration. SSSD connects to its proxy service, and then that proxy loads the specified libraries. This allows SSSD to use some resources that it otherwise would not be able to use. For example, SSSD only supports LDAP and Kerberos as authentication providers, but using a proxy allows SSSD to use alternative authentication methods like a fingerprint scanner or smart card.

표 28.7. Proxy Domain Configuration Parameters

Parameter	Description
proxy_pam_target	Specifies the target to which PAM must proxy as an authentication provider.. The PAM target is a file containing PAM stack information in the default PAM directory, /etc/pam.d/. This is used to proxy an authentication provider. 중요 Ensure that the proxy PAM stack does not recursively include pam_sss.so.
proxy_lib_name	Specifies which existing NSS library to proxy identity requests through. This is used to proxy an identity provider.

예 28.5. Proxy Identity and Kerberos Authentication

The proxy library is loaded using the proxy_lib_name parameter. This library can be anything as long as it is compatible with the given authentication service. For a Kerberos authentication provider, it must be a Kerberos-compatible library, like NIS.

```
[domain/PROXY_KRB5]
auth_provider = krb5
krb5_server = 192.168.1.1
krb5_realm = EXAMPLE.COM

id_provider = proxy
proxy_lib_name = nis
enumerate = true
cache_credentials = true
```

예 28.6. LDAP Identity and Proxy Authentication

The proxy library is loaded using the proxy_pam_target parameter. This library must be a PAM module that is compatible with the given identity provider. For example, this uses a PAM fingerprint module with LDAP:

```
[domain/LDAP_PROXY]
id_provider = ldap
ldap_uri = ldap://example.com
ldap_search_base = dc=example,dc=com

auth_provider = proxy
proxy_pam_target = sssdpamproxy
enumerate = true
cache_credentials = true
```

After the SSSD domain is configured, make sure that the specified PAM files are configured. In this, the target is sssdpamproxy, so create a /etc/pam.d/sssdpamproxy file and load the PAM/LDAP modules:

```
auth        required      pam_frprint.so
account     required      pam_frprint.so
password    required      pam_frprint.so
session     required      pam_frprint.so
```

예 28.7. Proxy Identity and Authentication

SSSD can have a domain with both identity and authentication proxies. The only configuration given then are the proxy settings, proxy_pam_target for the authentication PAM module and proxy_lib_name for the service, like NIS or LDAP.

This example illustrates a possible configuration, but this is not a realistic configuration. If LDAP is used for identity and authentication, then both the identity and authentication providers should be set to the LDAP configuration, not a proxy.

```
[domain/PROXY_PROXY]
auth_provider = proxy
id_provider = proxy
proxy_lib_name = ldap
```

```
proxy_pam_target = sssdproxyldap
enumerate = true
cache_credentials = true
```

Once the SSSD domain is added, then update the system settings to configure the proxy service:

1. Create a /etc/pam.d/sssdproxyldap file which requires the pam_ldap.so module:

   ```
   auth        required    pam_ldap.so
   account     required    pam_ldap.so
   password    required    pam_ldap.so
   session     required    pam_ldap.so
   ```

2. Edit the /etc/nslcd.conf file, the configuration file for the LDAP name service daemon, to contain the information for the LDAP directory:

   ```
   uid nslcd
   gid ldap
   uri ldaps://ldap.example.com:636
   base dc=example,dc=com
   ssl on
   tls_cacertdir /etc/openldap/cacerts
   ```

28.5. Configuring Access Control for SSSD Domains

SSSD provides a rudimentary access control for domain configuration, allowing either simple user allow/deny lists or using the LDAP backend itself.

28.5.1. Using the Simple Access Provider

The Simple Access Provider allows or denies access based on a list of usernames or groups.

The Simple Access Provider is a way to restrict access to certain, specific machines. For example, if a company uses laptops, the Simple Access Provider can be used to restrict access to only a specific user or a specific group, even if a different user authenticated successfully against the same authentication provider.

The most common options are simple_allow_users and simple_allow_groups, which grant access explicitly to specific users (either the given users or group members) and deny access to everyone else. It is also possible to create deny lists (which deny access only to explicit people and implicitly allow everyone else access).

The Simple Access Provider adheres to the following three rules to determine which users should or should not be granted access:

- If both the allow and deny lists are empty, access is granted.
- If any list is provided, allow rules are evaluated first, and then deny rules. Practically, this means that deny rules supersede allow rules.
- If an allowed list is provided, then all users are denied access unless they are in the list.
- If only deny lists are provided, then all users are allowed access unless they are in the list.

For example, this grants access to two users and anyone who belongs to the IT group; implicitly, all other users are denied.

```
[domain/example.com]
access_provider = simple
simple_allow_users = jsmith,bjensen
simple_allow_groups = itgroup
```

참고

The LOCAL domain in SSSD does not support simple as an access provider.

Other options are listed in the sssd-simple man page, but these are rarely used.

28.5.2. Using the LDAP Access Filter

The LDAP server itself can provide the access control rules. The associated filter option (ldap_access_filter) specifies which users are granted access to the specified host. The user filter must be used or all users are denied access.

For example:

```
[domain/example.com]
access_provider = ldap
ldap_access_filter = memberOf=cn=allowedusers,ou=Groups,dc=example,dc=com
```

참고

Offline caching for LDAP access providers is limited to determining whether the user's last online login attempt was successful. Users that were granted access during their last login will continue to be granted access while offline.

SSSD can also check results by the account expiration policy and the authorizedService attribute.

28.6. Configuring Domain Failover

SSSD attempts to connect to machines and to services separately.

When SSSD tries to connect to one of its domain backends, it first tries to resolve the hostname of a given machine. If this resolution attempt fails, the machine is considered offline, and SSSD no longer attempts to connect to this machine for any other service.

If the resolution attempt succeeds, the backend tries to connect to a service on this machine. If the service connection attempt fails, then only this particular service is considered offline and the backend automatically switches over to the next service. The machine is still considered online and might still be tried for another service.

SSSD only tries the first IP address given in the DNS A record. To find multiple servers with a single request, SSSD relies on SRV records.

Connections are retried to offline machines or services every 30 seconds, until SSSD can successfully connect to the backend.

28.6.1. Configuring Failover

Configuring failover allows SSSD to switch automatically to a different server if the primary server fails. These servers are entered as a case-insensitive, comma-separated list in the [domain/Name] sections of the /etc/sssd/sssd.conf file. The servers are listed in order of preference. This list can contain any number of servers.

For example, for a native LDAP domain:

```
ldap_uri = ldap://ldap0.example.com, ldap://ldap1.example.com, ldap://ldap2.example.com
```

The first entry, ldap://ldap0.example.com, is the primary server. If this server fails, SSSD first attempts to connect to ldap1.example.com and then ldap2.example.com.

If the server parameter is not specified, then SSSD uses service discovery to try to find another server on the network.

중요

The failover servers must be entered as a comma-separated list of values for a single key. If there are multiple keys, SSSD only recognizes the last entry.

28.6.2. Using SRV Records with Failover

SSSD supports SRV records in its failover configuration. The SSSD configuration can specify a server that is later resolved into a list of specific servers using SRV requests.

For every service with which to use service discovery, add a special DNS record to the DNS server:

```
_service._protocol._domain TTL priority weight port hostname
```

The priority and weight attributes of SRV records provide fine-grained control over which servers to contact first if the primary server fails.

A typical configuration contains multiple such records, each with a different priority for failover and different weights for load balancing.

For more information on SRV records, see RFC 2782[1].

28.7. Deleting Domain Cache Files

SSSD can define multiple domains of the same type and different types of domain. SSSD maintains a separate database file for each domain, meaning each domain has its own cache. These cache files are stored in the /var/lib/sss/db/ directory.

[1] http://tools.ietf.org/html/rfc2782

If there is ever a problem with a domain, it is easy to purge the cache by stopping SSSD and deleting the cache file for that domain.

All cache files are named for the domain. For example, for a domain named exampleldap, the cache file is named cache_exampleldap.ldb.

Be careful when you delete a cache file. This operation has significant effects:

- Deleting the cache file deletes all user data, both identification and cached credentials. Consequently, do delete a cache file unless the system is online and can authenticate with a username against the domain's servers. Without a credentials cache, offline authentication will fail.
- If the configuration is changed to reference a different identity provider, SSSD will recognize users from both providers until the cached entries from the original provider time out.

  It is possible to avoid this by purging the cache, but the better option is to use a different domain name for the new provider. When SSSD is restarted, it creates a new cache file with the new name and the old file is ignored.

28.8. Using NSCD with SSSD

SSSD is not designed to be used with the NSCD daemon. Even though SSSD does not directly conflict with NSCD, using both services can result in unexpected behavior, especially with how long entries are cached.

The most common evidence of a problem is conflicts with NFS. When using Network Manager to manage network connections, it may take several minutes for the network interface to come up. During this time, various services attempt to start. If these services start before the network is up and the DNS servers are available, these services fail to identify the forward or reverse DNS entries they need. These services will read an incorrect or possibly empty resolv.conf file. This file is typically only read once, and so any changes made to this file are not automatically applied. This can cause NFS locking to fail on the machine where the NSCD service is running, unless that service is manually restarted.

To avoid this problem, enable caching for hosts and services in the /etc/nscd.conf file and rely on the SSSD cache for the passwd, group, and netgroup entries.

Change the /etc/nscd.conf file:

```
enable-cache hosts yes
enable-cache passwd no
enable-cache group no
enable-cache netgroup no
```

With NSCD answering hosts requests, these entries will be cached by NSCD and returned by NSCD during the boot process. All other entries are handled by SSSD.

28.9. Troubleshooting SSSD

28.9.1. Using SSSD Log Files

SSSD uses a number of log files to report information about its operation, located in the /var/log/sssd/ directory. SSSD produces a log file for each domain, as well as an sssd_pam.log and an sssd_nss.log file.

Additionally, the /var/log/secure file logs authentication failures and the reason for the failure.

Increasing the log level can provide more information about problems with SSSD. To change the log level, set the debug_level parameter for each section in the sssd.conf file for which to product extra logs. For example:

```
[sssd]
config_file_version = 2
services = nss, pam
domains = LDAP
debug_level = 9
```

28.9.2. Problems with SSSD Configuration

SSSD fails to start

SSSD requires that the configuration file be properly set up, with all the required entries, before the daemon will start.

- SSSD requires at least one properly configured domain before the service will start. Without a domain, attempting to start SSSD returns an error that no domains are configured:

  ```
  # sssd -d4

  [sssd] [ldb] (3): server_sort:Unable to register control with rootdse!
  [sssd] [confdb_get_domains] (0): No domains configured, fatal error!
  [sssd] [get_monitor_config] (0): No domains configured.
  ```

  Edit the /etc/sssd/sssd.conf file and create at least one domain.

- SSSD also requires at least one available service provider before it will start. If the problem is with the service provider configuration, the error message indicates that there are no services configured:

  ```
  [sssd] [get_monitor_config] (0): No services configured!
  ```

  Edit the /etc/sssd/sssd.conf file and configure at least one service provider.

중요

SSSD requires that service providers be configured as a comma-separated list in a single services entry in the /etc/sssd/sssd.conf file. If services are listed in multiple entries, only the last entry is recognized by SSSD.

NSS fails to return user information

This usually means that SSSD cannot connect to the NSS service.

- Ensure that NSS is running:

  ```
  # service sssd status
  ```

- If NSS is running, make sure that the provider is properly configured in the [nss] section of the /etc/sssd/sssd.conf file. Especially check the filter_users and filter_groups attributes.
- Make sure that NSS is included in the list of services that SSSD uses.
- Check the configuration in the /etc/nsswitch.conf file.

NSS returns incorrect user information

If searches are returning the incorrect user information, check that there are not conflicting usernames in separate domains. When there are multiple domains, set the use_fully_qualified_domains attribute to TRUE in the /etc/sssd/sssd.conf file. This differentiates between different users in different domains with the same name.

Setting the password for the local SSSD user prompts twice for the password

When attempting to change a local SSSD user's password, it may prompt for the password twice:

```
[root@clientF11 tmp]# passwd user1000
Changing password for user user1000.
New password:
Retype new password:
New Password:
Reenter new Password:
passwd: all authentication tokens updated successfully.
```

This is the result of an incorrect PAM configuration. Ensure that the use_authtok option is correctly configured in your /etc/pam.d/system-auth file.

부 IV. 시스템 설정

Part of a system administrator's job is configuring the system for various tasks, types of users, and hardware configurations. This section explains how to configure a Red Hat Enterprise Linux system.

콘솔 사용

루트가 아닌 일반 사용자가 로컬 컴퓨터에 로그인할 때, 다음과 같은 두 가지 유형의 특수 권한이 주어집니다:

1. 실행할 수 없었던 특정 프로그램을 실행할 수 있는 권한

2. 사용할 수 없었던 특정 파일(주로 디스켓, CD-ROM과 같은 장치를 접근할 때 사용하는 특정 장치 파일)을 사용할 수 있는 권한

한 컴퓨터에 다중 콘솔이 지원되므로 많은 사용자가 로컬 컴퓨터에 동시에 로그인할 수 있습니다. 즉, 사용자는 파일을 더 빨리 사용해야 하는 경주에 이겨야 합니다. 콘솔에 가장 먼저 로그인한 사용자가 원하는 파일을 소유할 수 있습니다. 첫 번째 사용자가 로그아웃하면 다음 로그인한 사용자가 파일을 소유하게 됩니다.

In contrast, every user who logs in at the console is allowed to run programs that accomplish tasks normally restricted to the root user. If X is running, these actions can be included as menu items in a graphical user interface. As shipped, these console-accessible programs include halt, poweroff, and reboot.

29.1. Disabling Shutdown Via Ctrl+Alt+Del

By default, /etc/inittab specifies that your system is set to shutdown and reboot in response to a Ctrl+Alt+Del key combination used at the console. To completely disable this ability, comment out the following line in /etc/inittab by putting a hash mark (#) in front of it:

```
ca::ctrlaltdel:/sbin/shutdown -t3 -r now
```

Alternatively, you may want to allow certain non-root users the right to shutdown or reboot the system from the console using Ctrl+Alt+Del . You can restrict this privilege to certain users, by taking the following steps:

1. Add the -a option to the /etc/inittab line shown above, so that it reads:

   ```
   ca::ctrlaltdel:/sbin/shutdown -a -t3 -r now
   ```

   The -a flag tells shutdown to look for the /etc/shutdown.allow file.

2. Create a file named shutdown.allow in /etc. The shutdown.allow file should list the usernames of any users who are allowed to shutdown the system using Ctrl+Alt+Del . The format of the shutdown.allow file is a list of usernames, one per line, like the following:

   ```
   stephen
   jack
   sophie
   ```

According to this example shutdown.allow file, the users stephen, jack, and sophie are allowed to shutdown the system from the console using Ctrl+Alt+Del . When that key combination is used, the shutdown -a command in /etc/inittab checks to see if any of the users in /etc/shutdown.allow (or root) are logged in on a virtual console. If one of them is, the shutdown of the system continues; if not, an error message is written to the system console instead.

For more information on shutdown.allow, refer to the shutdown man page.

29.2. 콘솔 프로그램 사용 거부

사용자가 콘솔 프로그램을 사용할 수 없도록 설정하려면 루트로 로그인한 후 다음 명령어를 실행해야 합니다:

```
rm -f /etc/security/console.apps/*
```

In environments where the console is otherwise secured (BIOS and boot loader passwords are set, Ctrl+Alt+Delete is disabled, the power and reset switches are disabled, and so forth), you may not want to allow any user at the console to run poweroff, halt, and reboot, which are accessible from the console by default.

이러한 권한을 제거하려면 루트로 다음 명령어를 실행하십시오:

```
rm -f /etc/security/console.apps/poweroff
rm -f /etc/security/console.apps/halt
rm -f /etc/security/console.apps/reboot
```

29.3. 콘솔 정의

The pam_console.so module uses the /etc/security/console.perms file to determine the permissions for users at the system console. The syntax of the file is very flexible; you can edit the file so that these instructions no longer apply. However, the default file has a line that looks like this:

```
<console>=tty[0-9][0-9]* vc/[0-9][0-9]* :[0-9]\.[0-9] :[0-9]
```

When users log in, they are attached to some sort of named terminal, which can be either an X server with a name like :0 or mymachine.example.com:1.0, or a device like /dev/ttyS0 or /dev/pts/2. The default is to define that local virtual consoles and local X servers are considered local, but if you want to consider the serial terminal next to you on port /dev/ttyS1 to also be local, you can change that line to read:

```
<console>=tty[0-9][0-9]* vc/[0-9][0-9]* :[0-9]\.[0-9] :[0-9] /dev/ttyS1
```

29.4. 콘솔에서 파일 사용 가능하도록 설정

개별적인 장치 클래스 및 권한 정의에 관한 디폴트 설정은 /etc/security/console.perms.d/50-default.perms 파일에 정의되어 있습니다. 파일과 장치 권한을 편집하려면 /etc/security/console.perms.d/ 디렉토리에 특정 파일 또는 장치에 대한 필요한 설정이 지정된 새로운 디폴트 파일을 생성하는 것이 좋습니다. 새 디폴트 파일의 이름은 반드시 50-default.perms 파일 내용을 덮어쓸 수 있도록 50보다 큰 숫자(예, 51-default.perms)로 시작해야 합니다.

/etc/security/console.perms.d/ 디렉토리에 51-default.perms 파일을 생성하십시오:

```
touch /etc/security/console.perms.d/51-default.perms
```

기존의 디폴트 perms 파일인 50-default.perms 파일을 엽니다. 첫 번째 섹션은 다음과 같은 줄로 장치 클래스를 정의합니다.

```
<floppy>=/dev/fd[0-1]* \
            /dev/floppy/* /mnt/floppy*
<sound>=/dev/dsp* /dev/audio* /dev/midi* \
    /dev/mixer* /dev/sequencer \
    /dev/sound/* /dev/beep \
    /dev/snd/*
<cdrom>=/dev/cdrom* /dev/cdroms/* /dev/cdwriter* /mnt/cdrom*
```

Items enclosed in brackets name the device; in the above example, <cdrom> refers to the CD-ROM drive. To add a new device, do not define it in the default 50-default.perms file; instead, define it in 51-default.perms. For example, to define a scanner, add the following line to 51-default.perms:

```
<scanner>=/dev/scanner /dev/usb/scanner*
```

Of course, you must use the appropriate name for the device. Ensure that /dev/scanner is really your scanner and not some other device, such as your hard drive.

장치나 파일을 정확히 정의했으면 다음 단계는 장치나 파일의 권한 정의를 지정하는 것입니다. /etc/security/console.perms.d/50-default.perms의 두 번째 섹션은 다음과 같은 줄로 이러한 권한 정의를 다루고 있습니다:

```
<console> 0660 <floppy> 0660 root.floppy
<console> 0600 <sound>  0640 root
<console> 0600 <cdrom>  0600 root.disk
```

스캐너에 대한 권한을 정의하려면 51-default.perms에 다음과 같은 줄을 추가합니다:

```
<console> 0600 <scanner> 0600 root
```

Then, when you log in at the console, you are given ownership of the /dev/scanner device with the permissions of 0600 (readable and writable by you only). When you log out, the device is owned by root, and still has the permissions 0600 (now readable and writable by root only).

경고

디폴트 50-default.perms 파일을 절대 수정해서는 안 됩니다. 이미 50-default.perms에 정의된 장치 권한을 수정하려면 51-default.perms에 원하는 장치 권한 정의를 추가해야 합니다. 이것은 50-default.perms에 정의된 권한과 관계없이 덮어쓰게 됩니다.

29.5. 다른 어플리케이션에 대한 콘솔 사용 활성화

콘솔 사용자가 다른 어플리케이션을 사용하도록 설정하려면 몇 가지 추가 작업이 필요합니다.

First of all, console access only works for applications which reside in /sbin/ or /usr/sbin/, so the application that you wish to run must be there. After verifying that, perform the following steps:

1. Create a link from the name of your application, such as our sample foo program, to the /usr/bin/consolehelper application:

   ```
   cd /usr/bin
   ln -s consolehelper foo
   ```

2. Create the file /etc/security/console.apps/foo:

```
touch /etc/security/console.apps/foo
```

3. Create a PAM configuration file for the foo service in /etc/pam.d/. An easy way to do this is to copy the PAM configuration file of the halt service, and then modify the copy if you want to change the behavior:

```
cp /etc/pam.d/halt /etc/pam.d/foo
```

Now, when /usr/bin/foo is executed, consolehelper is called, which authenticates the user with the help of /usr/sbin/userhelper. To authenticate the user, consolehelper asks for the user's password if /etc/pam.d/foo is a copy of /etc/pam.d/halt (otherwise, it does precisely what is specified in /etc/pam.d/foo) and then runs /usr/sbin/foo with root permissions.

In the PAM configuration file, an application can be configured to use the pam_timestamp module to remember (or cache) a successful authentication attempt. When an application is started and proper authentication is provided (the root password), a timestamp file is created. By default, a successful authentication is cached for five minutes. During this time, any other application that is configured to use pam_timestamp and run from the same session is automatically authenticated for the user — the user does not have to enter the root password again.

This module is included in the pam package. To enable this feature, add the following lines to your PAM configuration file in etc/pam.d/:

```
auth            include         config-util
account         include         config-util
session         include         config-util
```

These lines can be copied from any of the /etc/pam.d/system-config-* configuration files. Note that these lines must be added below any other auth sufficient session optional lines in your PAM configuration file.

If an application configured to use pam_timestamp is successfully authenticated from the Applications (the main menu on the panel), the 🔒 icon is displayed in the notification area of the panel if you are running the GNOME or KDE desktop environment. After the authentication expires (the default is five minutes), the icon disappears.

아이콘에 클릭한 후 인증을 기억하지 않기 옵션을 선택하면 캐시된 인증을 기억하지 않도록 설정하실 수 있습니다.

29.6. The floppy Group

If, for whatever reason, console access is not appropriate for you and your non-root users require access to your system's diskette drive, this can be done using the floppy group. Add the user(s) to the floppy group using the tool of your choice. For example, the gpasswd command can be used to add user fred to the floppy group:

```
gpasswd -a fred floppy
```

Now, user fred is able to access the system's diskette drive from the console.

The sysconfig Directory

The /etc/sysconfig/ directory contains a variety of system configuration files for Red Hat Enterprise Linux.

This chapter outlines some of the files found in the /etc/sysconfig/ directory, their function, and their contents. The information in this chapter is not intended to be complete, as many of these files have a variety of options that are only used in very specific or rare circumstances.

30.1. Files in the /etc/sysconfig/ Directory

The following sections offer descriptions of files normally found in the /etc/sysconfig/ directory. Files not listed here, as well as extra file options, are found in the /usr/share/doc/initscripts-<version-number>/sysconfig.txt file (replace <version-number> with the version of the initscripts package). Alternatively, looking through the initscripts in the /etc/rc.d/ directory can prove helpful.

알림

If some of the files listed here are not present in the /etc/sysconfig/ directory, then the corresponding program may not be installed.

30.1.1. /etc/sysconfig/amd

The /etc/sysconfig/amd file contains various parameters used by amd; these parameters allow for the automatic mounting and unmounting of file systems.

30.1.2. /etc/sysconfig/apmd

The /etc/sysconfig/apmd file is used by apmd to configure what power settings to start/stop/change on suspend or resume. This file configures how apmd functions at boot time, depending on whether the hardware supports Advanced Power Management (APM) or whether the user has configured the system to use it. The apm daemon is a monitoring program that works with power management code within the Linux kernel. It is capable of alerting users to low battery power on laptops and other power-related settings.

30.1.3. /etc/sysconfig/arpwatch

The /etc/sysconfig/arpwatch file is used to pass arguments to the arpwatch daemon at boot time. The arpwatch daemon maintains a table of Ethernet MAC addresses and their IP address pairings. By default, this file sets the owner of the arpwatch process to the user pcap and sends any messages to the root mail queue. For more information regarding available parameters for this file, refer to the arpwatch man page.

30.1.4. /etc/sysconfig/authconfig

The /etc/sysconfig/authconfig file sets the authorization to be used on the host. It contains one or more of the following lines:

- USEMD5=<value>, where <value> is one of the following:
  - yes — MD5 is used for authentication.
  - no — MD5 is not used for authentication.
- USEKERBEROS=<value>, where <value> is one of the following:
  - yes — Kerberos is used for authentication.
  - no — Kerberos is not used for authentication.
- USELDAPAUTH=<value>, where <value> is one of the following:
  - yes — LDAP is used for authentication.
  - no — LDAP is not used for authentication.

30.1.5. /etc/sysconfig/autofs

The /etc/sysconfig/autofs file defines custom options for the automatic mounting of devices. This file controls the operation of the automount daemons, which automatically mount file systems when you use them and unmount them after a period of inactivity. File systems can include network file systems, CD-ROMs, diskettes, and other media.

The /etc/sysconfig/autofs file may contain the following:

- LOCALOPTIONS="<value>", where <value> is a string for defining machine-specific automount rules. The default value is an empty string ("").
- DAEMONOPTIONS="<value>", where <value> is the timeout length in seconds before unmounting the device. The default value is 60 seconds ("--timeout=60").
- UNDERSCORETODOT=<value>, where <value> is a binary value that controls whether to convert underscores in file names into dots. For example, auto_home to auto.home and auto_mnt to auto.mnt. The default value is 1 (true).
- DISABLE_DIRECT=<value>, where <value> is a binary value that controls whether to disable direct mount support, as the Linux implementation does not conform to the Sun Microsystems' automounter behavior. The default value is 1 (true), and allows for compatibility with the Sun automounter options specification syntax.

30.1.6. /etc/sysconfig/clock

The /etc/sysconfig/clock file controls the interpretation of values read from the system hardware clock.

설정 값:

- UTC=<value>, where <value> is one of the following boolean values:
  - true or yes — The hardware clock is set to Universal Time.
  - false or no — The hardware clock is set to local time.

- ARC=<value>, where <value> is the following:
  - false or no — This value indicates that the normal UNIX epoch is in use. Other values are used by systems not supported by Red Hat Enterprise Linux.
- SRM=<value>, where <value> is the following:
  - false or no — This value indicates that the normal UNIX epoch is in use. Other values are used by systems not supported by Red Hat Enterprise Linux.
- ZONE=<filename> — The time zone file under /usr/share/zoneinfo that /etc/localtime is a copy of. The file contains information such as:

```
ZONE="America/New York"
```

  알림: Time and Date Properties Tool(system-config-date)로 ZONE 파라미터를 읽을 수 있으며, 수동으로 ZONE 파라미터를 수정하여 시스템 시간 영역을 변경할 수 없습니다.

이전 버전의 Red Hat Enterprise Linux에서는 다음 값을 사용했습니다(현재 사용되지 않음):

- CLOCKMODE=<value>, where <value> is one of the following:
  - GMT — The clock is set to Universal Time (Greenwich Mean Time).
  - ARC — The ARC console's 42-year time offset is in effect (for Alpha-based systems only).

30.1.7. /etc/sysconfig/desktop

The /etc/sysconfig/desktop file specifies the desktop for new users and the display manager to run when entering runlevel 5.

설정 값:

- DESKTOP="<value>", where "<value>" is one of the following:
  - GNOME — Selects the GNOME desktop environment.
  - KDE — Selects the KDE desktop environment.
- DISPLAYMANAGER="<value>", where "<value>" is one of the following:
  - GNOME — Selects the GNOME Display Manager.
  - KDE — Selects the KDE Display Manager.
  - XDM — Selects the X Display Manager.

For more information, refer to 33장. X Window System.

30.1.8. /etc/sysconfig/dhcpd

The /etc/sysconfig/dhcpd file is used to pass arguments to the dhcpd daemon at boot time. The dhcpd daemon implements the Dynamic Host Configuration Protocol (DHCP) and the Internet Bootstrap Protocol (BOOTP). DHCP and BOOTP assign hostnames to machines on the network. For more information about what parameters are available in this file, refer to the dhcpd man page.

30.1.9. /etc/sysconfig/exim

The /etc/sysconfig/exim file allows messages to be sent to one or more clients, routing the messages over whatever networks are necessary. The file sets the default values for exim to run. Its default values are set to run as a background daemon and to check its queue each hour in case something has backed up.

설정 값:

- DAEMON=<value>, where <value> is one of the following:
  - yes — exim should be configured to listen to port 25 for incoming mail. yes implies the use of the Exim's -bd options.
  - no — exim should not be configured to listen to port 25 for incoming mail.
- QUEUE=1h which is given to exim as -q$QUEUE. The -q option is not given to exim if /etc/sysconfig/exim exists and QUEUE is empty or undefined.

30.1.10. /etc/sysconfig/firstboot

The first time the system boots, the /sbin/init program calls the etc/rc.d/init.d/firstboot script, which in turn launches the Setup Agent. This application allows the user to install the latest updates as well as additional applications and documentation.

The /etc/sysconfig/firstboot file tells the Setup Agent application not to run on subsequent reboots. To run it the next time the system boots, remove /etc/sysconfig/firstboot and execute chkconfig --level 5 firstboot on.

30.1.11. /etc/sysconfig/gpm

The /etc/sysconfig/gpm file is used to pass arguments to the gpm daemon at boot time. The gpm daemon is the mouse server which allows mouse acceleration and middle-click pasting. For more information about what parameters are available for this file, refer to the gpm man page. By default, the DEVICE directive is set to /dev/input/mice.

30.1.12. /etc/sysconfig/hwconf

The /etc/sysconfig/hwconf file lists all the hardware that kudzu detected on the system, as well as the drivers used, vendor ID, and device ID information. The kudzu program detects and configures new and/or changed hardware on a system. The /etc/sysconfig/hwconf file is not meant to be manually edited. If edited, devices could suddenly show up as being added or removed.

30.1.13. /etc/sysconfig/i18n

The /etc/sysconfig/i18n file sets the default language, any supported languages, and the default system font. For example:

```
LANG="en_US.UTF-8"
SUPPORTED="en_US.UTF-8:en_US:en"
SYSFONT="latarcyrheb-sun16"
```

30.1.14. /etc/sysconfig/init

The /etc/sysconfig/init file controls how the system appears and functions during the boot process.

설정 값:

- BOOTUP=<value>, where <value> is one of the following:
  - color — The standard color boot display, where the success or failure of devices and services starting up is shown in different colors.
  - verbose — An old style display which provides more information than purely a message of success or failure.
  - anything else는 ANSI 형식이 아닌 새로운 화면을 의미합니다.
- RES_COL=<value>, where <value> is the number of the column of the screen to start status labels. The default is set to 60.
- MOVE_TO_COL=<value>, where <value> moves the cursor to the value in the RES_COL line via the echo -en command.
- SETCOLOR_SUCCESS=<value>, where <value> sets the success color via the echo -en command. The default color is set to green.
- SETCOLOR_FAILURE=<value>, where <value> sets the failure color via the echo -en command. The default color is set to red.
- SETCOLOR_WARNING=<value>, where <value> sets the warning color via the echo -en command. The default color is set to yellow.
- SETCOLOR_NORMAL=<value>, where <value> resets the color to "normal" via the echo -en.
- LOGLEVEL=<value>, where <value> sets the initial console logging level for the kernel. The default is 3; 8 means everything (including debugging), while 1 means only kernel panics. The syslogd daemon overrides this setting once started.
- PROMPT=<value>, where <value> is one of the following boolean values:
  - yes — Enables the key check for interactive mode.
  - no — Disables the key check for interactive mode.

30.1.15. /etc/sysconfig/ip6tables-config

The /etc/sysconfig/ip6tables-config file stores information used by the kernel to set up IPv6 packet filtering at boot time or whenever the ip6tables service is started.

Do not modify this file by hand unless familiar with how to construct ip6tables rules. Rules also can be created manually using the /sbin/ip6tables command. Once created, add the rules to the /etc/sysconfig/ip6tables file by typing the following command:

```
service ip6tables save
```

이 파일에 저장된 모든 방화벽 규칙은 시스템 재부팅 또는 서비스 재시작과 관계없이 지속됩니다.

For more information on ip6tables, refer to 46.9절. “IPTables” .

30.1.16. /etc/sysconfig/iptables-config

The /etc/sysconfig/iptables-config file stores information used by the kernel to set up packet filtering services at boot time or whenever the service is started.

Do not modify this file by hand unless you are familiar with constructing iptables rules. The easiest way to add rules is to use the Security Level Configuration Tool (system-config-securitylevel) application to create a firewall. These applications automatically edit this file at the end of the process.

Rules can also be created manually using the /sbin/iptables command. Once created, add the rule(s) to the /etc/sysconfig/iptables file by typing the following command:

```
service iptables save
```

이 파일에 저장된 모든 방화벽 규칙은 시스템 재부팅 또는 서비스 재시작과 관계없이 지속됩니다.

For more information on iptables, refer to 46.9절. "IPTables" .

30.1.17. /etc/sysconfig/irda

The /etc/sysconfig/irda file controls how infrared devices on the system are configured at startup.

설정 값:

- IRDA=<value>, where <value> is one of the following boolean values:
  - yes — irattach runs and periodically checks to see if anything is trying to connect to the infrared port, such as another notebook computer trying to make a network connection. For infrared devices to work on the system, this line must be set to yes.
  - no — irattach does not run, preventing infrared device communication.
- DEVICE=<value>, where <value> is the device (usually a serial port) that handles infrared connections. A sample serial device entry could be /dev/ttyS2.
- DONGLE=<value>, where <value> specifies the type of dongle being used for infrared communication. This setting exists for people who use serial dongles rather than real infrared ports. A dongle is a device that is attached to a traditional serial port to communicate via infrared. This line is commented out by default because notebooks with real infrared ports are far more common than computers with add-on dongles. A sample dongle entry could be actisys+.
- DISCOVERY=<value>, where <value> is one of the following boolean values:
  - yes — Starts irattach in discovery mode, meaning it actively checks for other infrared devices. This must be turned on for the machine to actively look for an infrared connection (meaning the peer that does not initiate the connection).
  - no — Does not start irattach in discovery mode.

30.1.18. /etc/sysconfig/keyboard

The /etc/sysconfig/keyboard file controls the behavior of the keyboard. The following values may be used:

- KEYBOARDTYPE="sun|pc" where sun means a Sun keyboard is attached on /dev/kbd, or pc means a PS/2 keyboard connected to a PS/2 port.

- KEYTABLE="<file>", where <file> is the name of a keytable file.

  For example: KEYTABLE="us". The files that can be used as keytables start in /lib/kbd/keymaps/i386 and branch into different keyboard layouts from there, all labeled <file>.kmap.gz. The first file found beneath /lib/kbd/keymaps/i386 that matches the KEYTABLE setting is used.

30.1.19. /etc/sysconfig/kudzu

The /etc/sysconfig/kuzdu file triggers a safe probe of the system hardware by kudzu at boot time. A safe probe is one that disables serial port probing.

- SAFE=<value>, where <value> is one of the following:

  - yes — kuzdu does a safe probe.

  - no — kuzdu does a normal probe.

30.1.20. /etc/sysconfig/named

The /etc/sysconfig/named file is used to pass arguments to the named daemon at boot time. The named daemon is a Domain Name System (DNS) server which implements the Berkeley Internet Name Domain (BIND) version 9 distribution. This server maintains a table of which hostnames are associated with IP addresses on the network.

현재, 다음 값만이 사용됩니다:

- ROOTDIR="</some/where>", where </some/where> refers to the full directory path of a configured chroot environment under which named runs. This chroot environment must first be configured. Type info chroot for more information.

- OPTIONS="<value>", where <value> is any option listed in the man page for named except -t. In place of -t, use the ROOTDIR line above.

For more information about available parameters for this file, refer to the named man page. For detailed information on how to configure a BIND DNS server, refer to 18장. Berkeley Internet Name Domain (BIND). By default, the file contains no parameters.

30.1.21. /etc/sysconfig/network

The /etc/sysconfig/network file is used to specify information about the desired network configuration. The following values may be used:

- NETWORKING=<value>, where <value> is one of the following boolean values:

  - yes — Networking should be configured.

  - no — Networking should not be configured.

- HOSTNAME=<value>, where <value> should be the Fully Qualified Domain Name (FQDN), such as hostname.expample.com, but can be whatever hostname is necessary.

- GATEWAY=<value>, where <value> is the IP address of the network's gateway.

- GATEWAYDEV=<value>, where <value> is the gateway device, such as eth0. Configure this option if you have multiple interfaces on the same subnet, and require one of those interfaces to be the preferred route to the default gateway.

- NISDOMAIN=<value>, where <value> is the NIS domain name.

- NOZEROCONF=<value>, where setting <value> to true disables the zeroconf route.

  By default, the zeroconf route (169.254.0.0) is enabled when the system boots. For more information about zeroconf, refer to http://www.zeroconf.org/.

> **Warning**
>
> Do not use custom initscripts to configure network settings. When performing a post-boot network service restart, custom initscripts configuring network settings that are run outside of the network init script lead to unpredictable results.

30.1.22. /etc/sysconfig/nfs

NFS requires portmap, which dynamically assigns ports for RPC services. This causes problems for configuring firewall rules. To overcome this problem, use the /etc/sysconfig/nfs file to control which ports the required RPC services run on.

The /etc/sysconfig/nfs may not exist by default on all systems. If it does not exist, create it and add the following variables (alternatively, if the file exists, un-comment and change the default entries as required):

MOUNTD_PORT=x
: control which TCP and UDP port mountd (rpc.mountd) uses. Replace x with an unused port number.

STATD_PORT=x
: control which TCP and UDP port status (rpc.statd) uses. Replace x with an unused port number.

LOCKD_TCPPORT=x
: control which TCP port nlockmgr (rpc.lockd) uses. Replace x with an unused port number.

LOCKD_UDPPORT=x
: control which UDP port nlockmgr (rpc.lockd) uses. Replace x with an unused port number.

If NFS fails to start, check /var/log/messages. Normally, NFS will fail to start if you specify a port number that is already in use. After editing /etc/sysconfig/nfs restart the NFS service by running the service nfs restart command. Run the rpcinfo -p command to confirm the changes.

To configure a firewall to allow NFS:

1. Allow TCP and UDP port 2049 for NFS.

2. Allow TCP and UDP port 111 (portmap/sunrpc).

3. Allow the TCP and UDP port specified with MOUNTD_PORT="x"

4. Allow the TCP and UDP port specified with STATD_PORT="x"

5. Allow the TCP port specified with LOCKD_TCPPORT="x"

6. Allow the UDP port specified with LOCKD_UDPPORT="x"

30.1.23. /etc/sysconfig/ntpd

The /etc/sysconfig/ntpd file is used to pass arguments to the ntpd daemon at boot time. The ntpd daemon sets and maintains the system clock to synchronize with an Internet standard time server. It implements version 4 of the Network Time Protocol (NTP). For more information about what parameters are available for this file, use a Web browser to view the following file: /usr/share/doc/ntp-<version>/ntpd.htm (where <version> is the version number of ntpd). By default, this file sets the owner of the ntpd process to the user ntp.

30.1.24. /etc/sysconfig/radvd

The /etc/sysconfig/radvd file is used to pass arguments to the radvd daemon at boot time. The radvd daemon listens for router requests and sends router advertisements for the IP version 6 protocol. This service allows hosts on a network to dynamically change their default routers based on these router advertisements. For more information about available parameters for this file, refer to the radvd man page. By default, this file sets the owner of the radvd process to the user radvd.

30.1.25. /etc/sysconfig/samba

The /etc/sysconfig/samba file is used to pass arguments to the smbd and the nmbd daemons at boot time. The smbd daemon offers file sharing connectivity for Windows clients on the network. The nmbd daemon offers NetBIOS over IP naming services. For more information about what parameters are available for this file, refer to the smbd man page. By default, this file sets smbd and nmbd to run in daemon mode.

30.1.26. /etc/sysconfig/selinux

The /etc/sysconfig/selinux file contains the basic configuration options for SELinux. This file is a symbolic link to /etc/selinux/config.

30.1.27. /etc/sysconfig/sendmail

The /etc/sysconfig/sendmail file allows messages to be sent to one or more clients, routing the messages over whatever networks are necessary. The file sets the default values for the Sendmail application to run. Its default values are set to run as a background daemon and to check its queue each hour in case something has backed up.

설정 값:

- DAEMON=<value>, where <value> is one of the following:
  - yes — Sendmail should be configured to listen to port 25 for incoming mail. yes implies the use of Sendmail's -bd options.
  - no — Sendmail should not be configured to listen to port 25 for incoming mail.

- QUEUE=1h which is given to Sendmail as -q$QUEUE. The -q option is not given to Sendmail if /etc/sysconfig/sendmail exists and QUEUE is empty or undefined.

30.1.28. /etc/sysconfig/spamassassin

The /etc/sysconfig/spamassassin file is used to pass arguments to the spamd daemon (a daemonized version of Spamassassin) at boot time. Spamassassin is an email spam filter application. For a list of available options, refer to the spamd man page. By default, it configures spamd to run in daemon mode, create user preferences, and auto-create whitelists (allowed bulk senders).

For more information about Spamassassin, refer to 25.5.2.6절. "Spam Filters" .

30.1.29. /etc/sysconfig/squid

The /etc/sysconfig/squid file is used to pass arguments to the squid daemon at boot time. The squid daemon is a proxy caching server for Web client applications. For more information on configuring a squid proxy server, use a Web browser to open the /usr/share/doc/squid-<version>/ directory (replace <version> with the squid version number installed on the system). By default, this file sets squid to start in daemon mode and sets the amount of time before it shuts itself down.

30.1.30. /etc/sysconfig/system-config-securitylevel

The /etc/sysconfig/system-config-securitylevel file contains all options chosen by the user the last time the Security Level Configuration Tool (system-config-securitylevel) was run. Users should not modify this file by hand. For more information about the Security Level Configuration Tool, refer to 46.8.2절. "Basic Firewall Configuration" .

30.1.31. /etc/sysconfig/system-config-selinux

The /etc/sysconfig/system-config-selinux file contains all options chosen by the user the last time the SELinux Administration Tool (system-config-selinux) was run. Users should not modify this file by hand. For more information about the SELinux Administration Tool and SELinux in general, refer to 47.2절. "Introduction to SELinux" .

30.1.32. /etc/sysconfig/system-config-users

The /etc/sysconfig/system-config-users file is the configuration file for the graphical application, User Manager. This file is used to filter out system users such as root, daemon, or lp. This file is edited by the Preferences > Filter system users and groups pull-down menu in the User Manager application and should never be edited by hand. For more information on using this application, refer to 35.1절. "사용자와 그룹 설정" .

30.1.33. /etc/sysconfig/system-logviewer

The /etc/sysconfig/system-logviewer file is the configuration file for the graphical, interactive log viewing application, Log Viewer. This file is edited by the Edit > Preferences pull-down menu in the Log Viewer application and should not be edited by hand. For more information on using this application, refer to 38장. 로그 파일.

30.1.34. /etc/sysconfig/tux

The /etc/sysconfig/tux file is the configuration file for the Red Hat Content Accelerator (formerly known as TUX), the kernel-based Web server. For more information on configuring the Red Hat Content Accelerator, use a Web browser to open the /usr/share/doc/tux-<version>/tux/index.html file (replace <version> with the version number of TUX installed on the system). The parameters available for this file are listed in /usr/share/doc/tux-<version>/tux/parameters.html.

30.1.35. /etc/sysconfig/vncservers

The /etc/sysconfig/vncservers file configures the way the Virtual Network Computing (VNC) server starts up.

VNC는 원격 디스플레이 시스템으로서 사용자가 실행 중인 시스템뿐만 아니라 다양한 구조의 다른 네트워크상에서도 데스크톱 환경을 볼 수 있도록 구성합니다.

설정 값:

- VNCSERVERS=<value>, where <value> is set to something like "1:fred", to indicate that a VNC server should be started for user fred on display :1. User fred must have set a VNC password using the vncpasswd command before attempting to connect to the remote VNC server.

30.1.36. /etc/sysconfig/xinetd

The /etc/sysconfig/xinetd file is used to pass arguments to the xinetd daemon at boot time. The xinetd daemon starts programs that provide Internet services when a request to the port for that service is received. For more information about available parameters for this file, refer to the xinetd man page. For more information on the xinetd service, refer to 46.5.3절. “xinetd” .

30.2. Directories in the /etc/sysconfig/ Directory

The following directories are normally found in /etc/sysconfig/.

apm-scripts/
: This directory contains the APM suspend/resume script. Do not edit the files directly. If customization is necessary, create a file called /etc/sysconfig/apm-scripts/apmcontinue which is called at the end of the script. It is also possible to control the script by editing /etc/sysconfig/apmd.

cbq/
: This directory contains the configuration files needed to do Class Based Queuing for bandwidth management on network interfaces. CBQ divides user traffic into a hierarchy of classes based on any combination of IP addresses, protocols, and application types.

networking/
: This directory is used by the Network Administration Tool (system-config-network), and its contents should not be edited manually. For more information about configuring network interfaces using the Network Administration Tool, refer to 16장. 네트워크 설정.

network-scripts/
: This directory contains the following network-related configuration files:

- Network configuration files for each configured network interface, such as ifcfg-eth0 for the eth0 Ethernet interface.
- Scripts used to bring network interfaces up and down, such as ifup and ifdown.
- Scripts used to bring ISDN interfaces up and down, such as ifup-isdn and ifdown-isdn.
- 직접 수정할 수 없는 다양한 공유 네트워크 기능 스크립트.

For more information on the network-scripts directory, refer to 15장. 네트워크 인터페이스.

rhn/

Deprecated. This directory contains the configuration files and GPG keys used by the RHN Classic content service. No files in this directory should be edited by hand.

This directory is available for legacy systems which are still managed by RHN Classic. Systems which are registered against the Certificate-Based Red Hat Network do not use this directory.

30.3. 추가 자료

This chapter is only intended as an introduction to the files in the /etc/sysconfig/ directory. The following source contains more comprehensive information.

30.3.1. 설치된 문서자료

- /usr/share/doc/initscripts-<version-number>/sysconfig.txt — This file contains a more authoritative listing of the files found in the /etc/sysconfig/ directory and the configuration options available for them. The <version-number> in the path to this file corresponds to the version of the initscripts package installed.

날짜와 시간 설정

The Time and Date Properties Tool allows the user to change the system date and time, to configure the time zone used by the system, and to setup the Network Time Protocol (NTP) daemon to synchronize the system clock with a time server.

이 도구를 사용하려면 X Window System을 실행해야 하며 루트 권한을 가지고 있어야 합니다. 이 어플리케이션을 시작하는 데 세 가지 방법이 있습니다:

- From the desktop, go to Applications (the main menu on the panel) > System Settings > Date & Time
- From the desktop, right-click on the time in the toolbar and select Adjust Date and Time.
- Type the command system-config-date, system-config-time, or dateconfig at a shell prompt (for example, in an XTerm or a GNOME terminal).

31.1. Time and Date Properties

As shown in 그림 31.1. “Time and Date Properties” , the first tabbed window that appears is for configuring the system date and time.

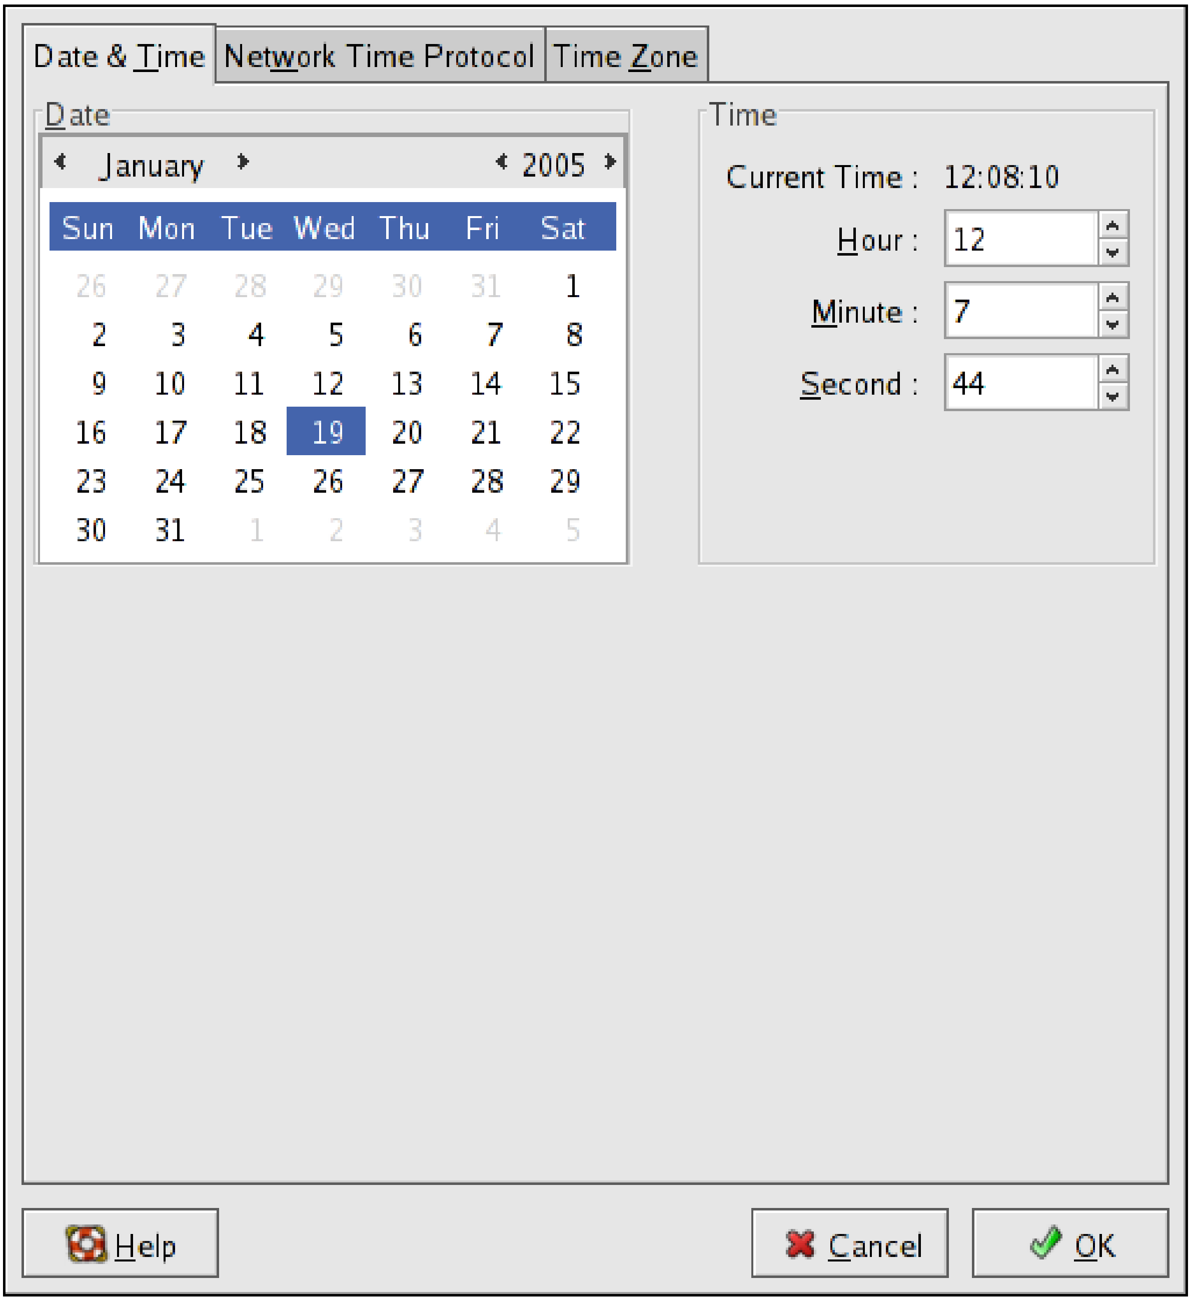

January 2005

Sun	Mon	Tue	Wed	Thu	Fri	Sat
26	27	28	29	30	31	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Current Time : 12:08:10

Hour : 12

Minute : 7

Second : 44

그림 31.1. Time and Date Properties

날짜를 변경하려면, 월 양쪽에 위치한 왼쪽/오른쪽 화살표를 사용하여 월수를 변경하시고, 연도수 양쪽에 위치한 화살표를 사용하여 연도수를 변경합니다. 또한, 요일수를 변경하려면 해당 요일에 클릭하면 됩니다.

To change the time, use the up and down arrow buttons beside the Hour, Minute, and Second in the Time section.

Clicking the OK button applies any changes made to the date and time, the NTP daemon settings, and the time zone settings. It also exits the program.

31.2. 네트워크 시간 프로토콜(NTP) 등록정보

As shown in 그림 31.2. “NTP Properties” , the second tabbed window that appears is for configuring NTP.

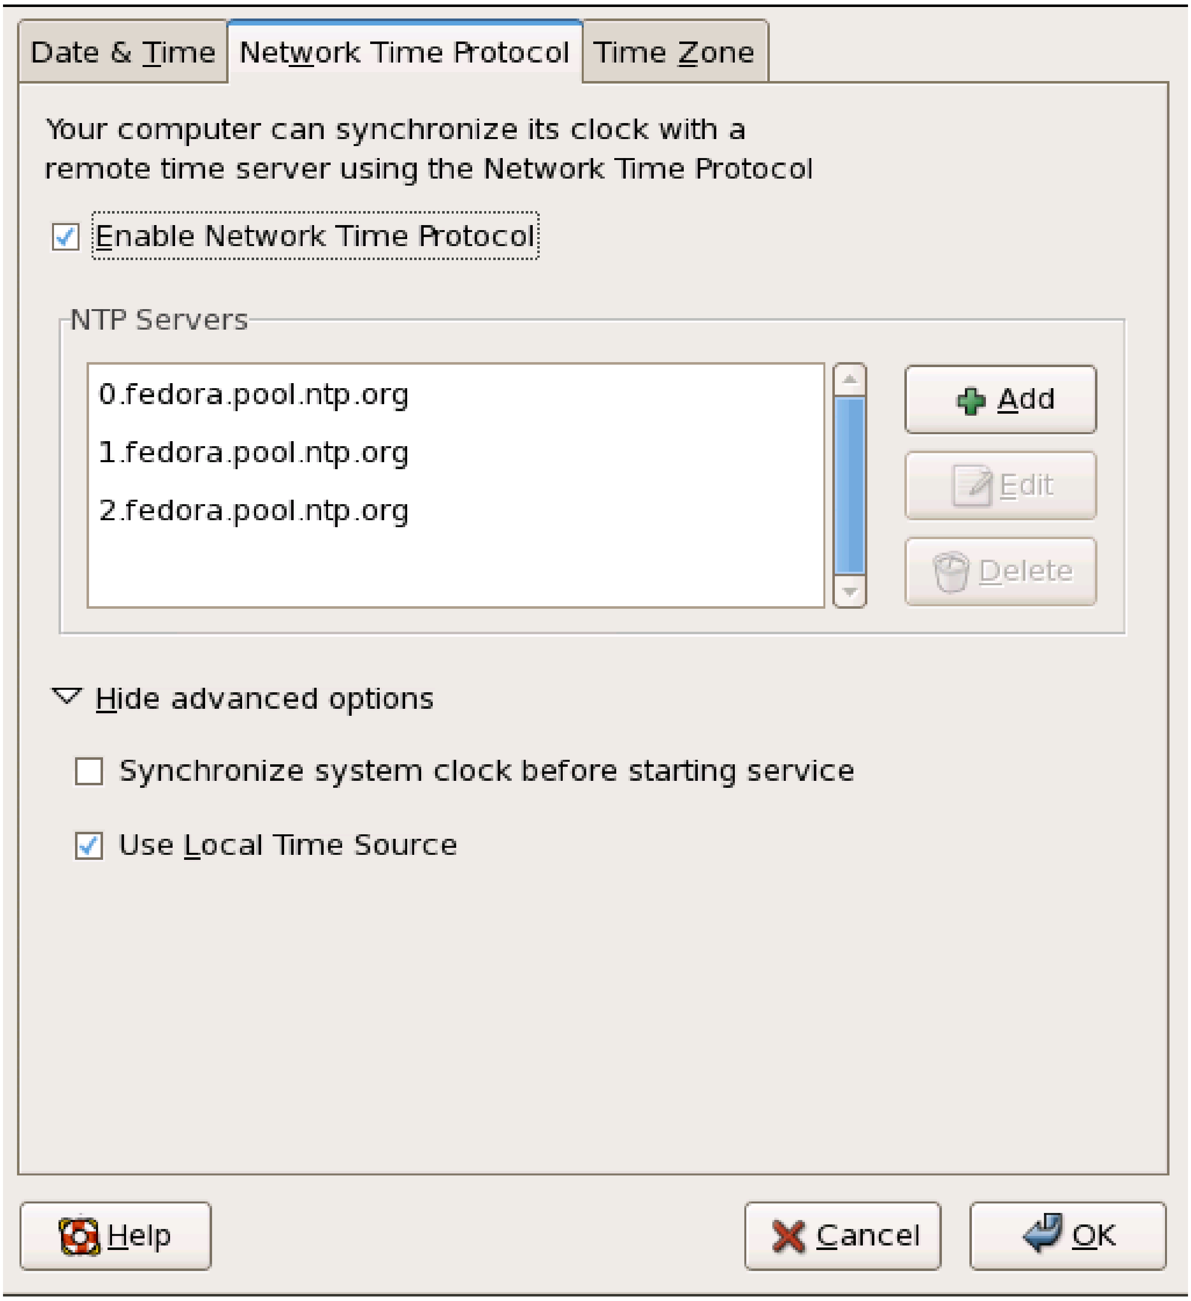

Your computer can synchronize its clock with a remote time server using the Network Time Protocol

☑ Enable Network Time Protocol

0.fedora.pool.ntp.org
1.fedora.pool.ntp.org
2.fedora.pool.ntp.org

□ Synchronize system clock before starting service

☑ Use Local Time Source

그림 31.2. NTP Properties

The Network Time Protocol (NTP) daemon synchronizes the system clock with a remote time server or time source. The application allows you to configure an NTP daemon to synchronize your system clock with a remote server. To enable this feature, select Enable Network Time Protocol. This enables the NTP Servers list and other options. You can choose one of the predefined servers, edit a predefined server by clicking the Edit or add a new server name by clicking Add. Your system does not start synchronizing with the NTP server until you click OK. After clicking OK, the configuration is saved and the NTP daemon is started (or restarted if it is already running).

Clicking the OK button applies any changes made to the date and time, the NTP daemon settings, and the time zone settings. It also exits the program.

31.3. 시간대 설정

As shown in 그림 31.3. "Timezone Properties" , the third tabbed window that appears is for configuring the system time zone.

To configure the system time zone, click the Time Zone tab. The time zone can be changed by either using the interactive map or by choosing the desired time zone from the list below the map. To use the map, click on the desired region. The map zooms into the region selected, after which you may choose the city specific to your time zone. A red X appears and the time zone selection changes in the list below the map.

다른 방법으로 지도 아래에 있는 목록을 사용하실 수 있습니다. 동일한 방법으로 지도에서 지역 및 도시를 선택하면 선택한 지역에 있는 도시 및 국가로 그룹 지어진 시간대 목록이 나타납니다. 비지역적인 시간대도 주소에 추가되었습니다.

Click OK to apply the changes and exit the program.

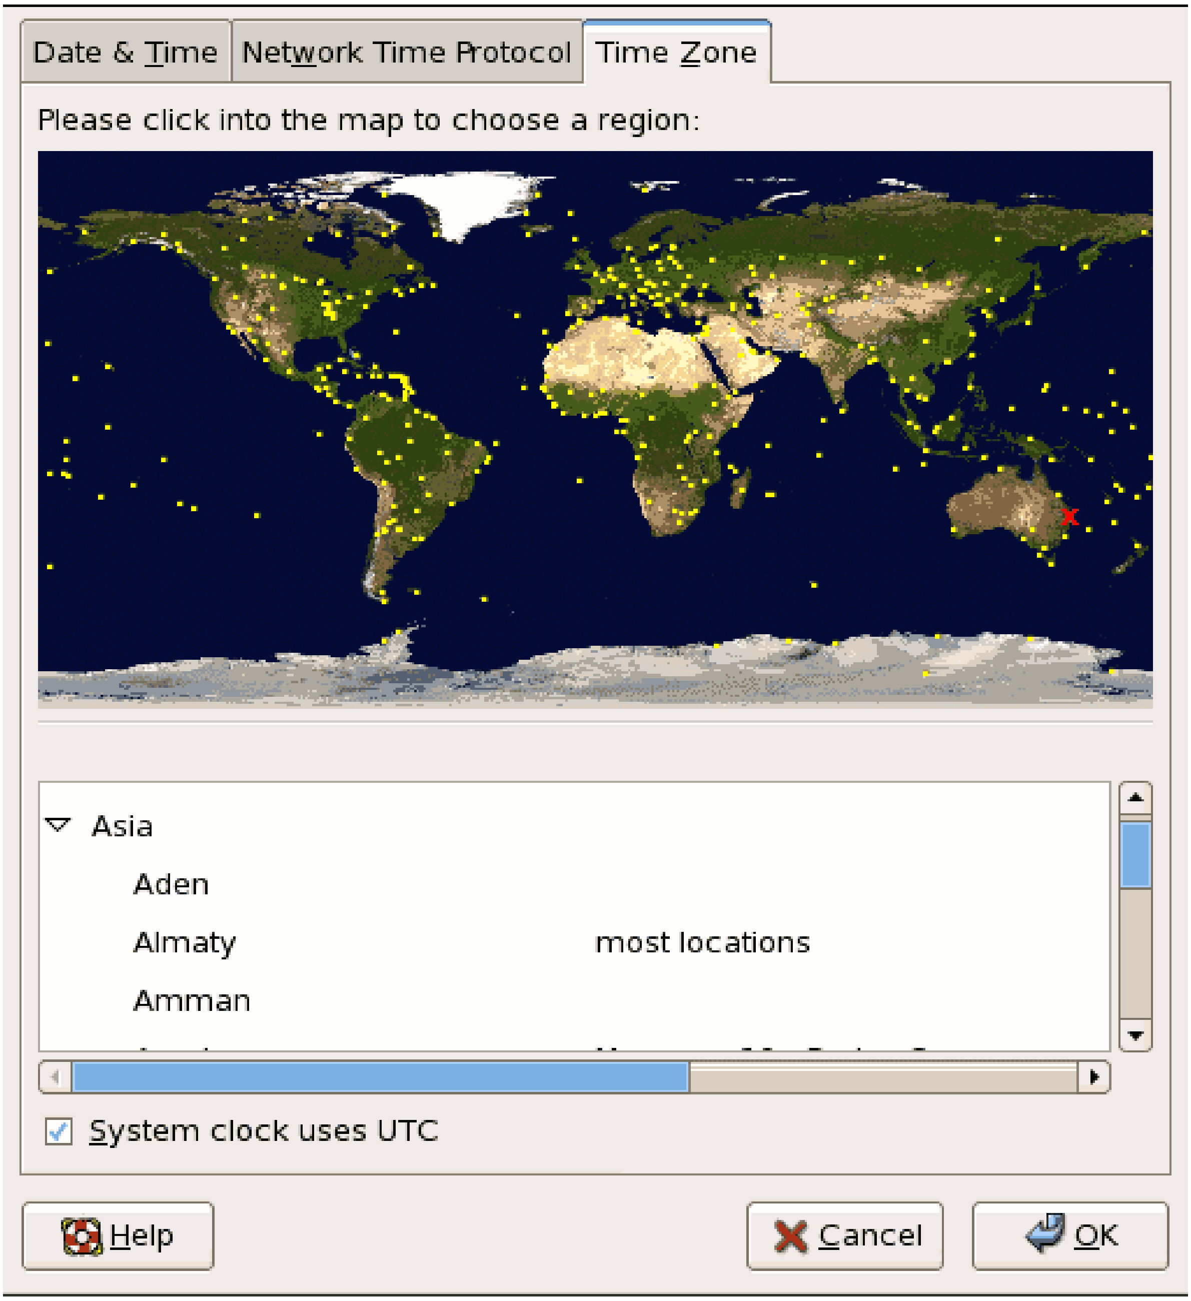

▽ Asia
Aden
Almaty most locations
Amman

☑ System clock uses UTC

그림 31.3. Timezone Properties

If your system clock is set to use UTC, select the System clock uses UTC option. UTC stands for the Universal Time, Coordinated, also known as Greenwich Mean Time (GMT). Other time zones are determined by adding or subtracting from the UTC time.

키보드 설정

The installation program allows you to configure a keyboard layout for your system. To configure a different keyboard layout after installation, use the Keyboard Configuration Tool.

To start the Keyboard Configuration Tool, select System (on the panel) > Administration > Keyboard, or type the command system-config-keyboard at a shell prompt.

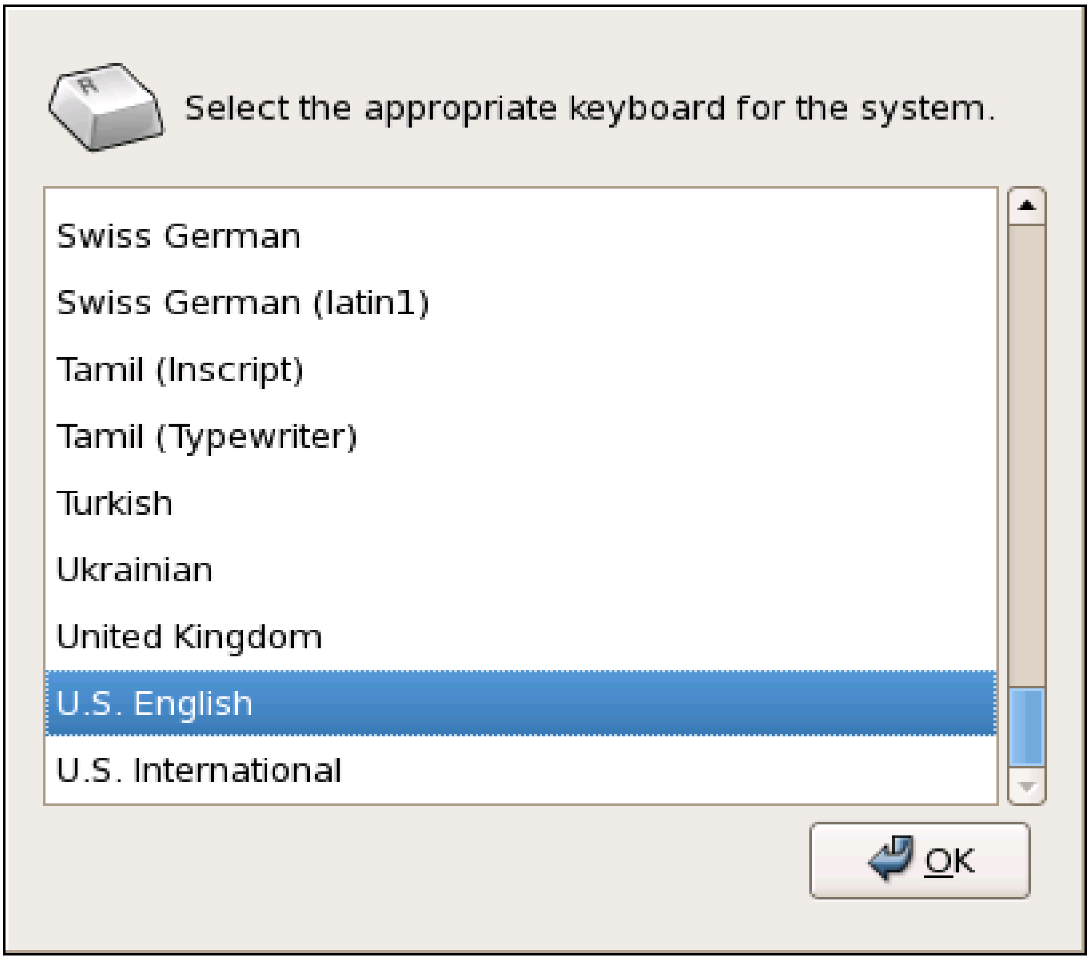

Select the appropriate keyboard for the system.

Swiss German
Swiss German (latin1)
Tamil (Inscript)
Tamil (Typewriter)
Turkish
Ukrainian
United Kingdom
U.S. English
U.S. International

그림 32.1. Keyboard Configuration Tool

Select a keyboard layout from the list (for example, U.S. English) and click OK.

변경 사항은 바로 적용됩니다.

X Window System

Red Hat Enterprise Linux의 중심부는 커널이지만 실제로 운영 체제의 외관은 X라고 부르는 X Window System에서 제공되는 그래픽형식 환경을 의미합니다.

1984년 6월에 X Window System이 배포되기 전부터 UNIX 세계에는 다른 윈도우 환경이 존재했습니다. 하지만, X는 수년 동안 Red Hat Enterprise Linux을 포함한 대부분 UNIX 운영 체제에서 기본 그래픽 환경으로 자리 잡았습니다.

Red Hat Enterprise Linux의 그래픽 환경은 X Window System과 관련 기술에 필요한 개발 및 전략을 관리하도록 창설된 오픈 소스 단체인 X.Org Foundation에서 제공됩니다. X.Org는 세계의 수 백명의 개발자가 참여하여 가장 빠르게 성장하는 대규모 프로젝트입니다. X는 다양한 하드웨어 장치와 구조에 대한 폭 넓은 지원을 갖추고 있으며 다른 운영 체제와 플랫폼에서도 실행 가능합니다. 최신 Red Hat Enterprise Linux 배포판에는 X Window System의 X11R7.1 버전 배포판이 내장되어 있습니다.

The X Window System uses a client-server architecture. The X server (the Xorg binary) listens for connections from X client applications via a network or local loopback interface. The server communicates with the hardware, such as the video card, monitor, keyboard, and mouse. X client applications exist in the user-space, creating a graphical user interface (GUI) for the user and passing user requests to the X server.

33.1. X11R7.1 배포판

Red Hat Enterprise Linux 5.8는 현재 기본 X Window System으로 새로운 비디오 드라이버, EXA, 이전 배포판에 없는 플랫폼 지원 강화 기능을 포함하는 X11R7.1 배포판을 사용합니다. 또한, X11R7.1 배포판은 X 서버에 사용되는 몇 가지 자동 구성 기능을 갖추고 있습니다.

X11R7.1 is the first release to take specific advantage of the modularization of the X Window System. This modularization, which splits X into logically distinct modules, makes it easier for open source developers to contribute code to the system.

중요

Red Hat Enterprise Linux no longer provides the XFree86™ server packages. Before upgrading a system to the latest version of Red Hat Enterprise Linux, be sure that the system's video card is compatible with the X11R7.1 release by checking the Red Hat Hardware Compatibility List located online at http://hardware.redhat.com/.

In the X11R7.1 release, all libraries, headers, and binaries now live under /usr/ instead of /usr/X11R6. The /etc/X11/ directory contains configuration files for X client and server applications. This includes configuration files for the X server itself, the xfs font server, the X display managers, and many other base components.

The configuration file for the newer Fontconfig-based font architecture is still /etc/fonts/fonts.conf. For more on configuring and adding fonts, refer to 33.4절. "Fonts" .

X 서버는 다양한 하드웨어에 고급 기능을 수행하기 때문에 작업하려는 하드웨어에 대한 자세한 정보가 필요합니다. X 서버는 자동으로 일부 하드웨어 정보를 검색하며, 다른 자세한 정보는 직접 설정돼야 합니다.

The installation program installs and configures X automatically, unless the X11R7.1 release packages are not selected for installation. However, if there are any changes to the monitor, video card or other devices managed by the X server, X must be reconfigured. The best way to do this is to use the X Configuration Tool (system-config-display), particularly for devices that are not detected manually.

In Red Hat Enterprise Linux's default graphical environment, the X Configuration Tool is available at System (on the panel) > Administration > Display.

Changes made with the X Configuration Tool take effect after logging out and logging back in.

For more information about X Configuration Tool, refer to 34장. X 윈도우 시스템 설정.

In some situations, reconfiguring the X server may require manually editing its configuration file, /etc/X11/xorg.conf. For information about the structure of this file, refer to 33.3절. "X 서버 구성 파일" .

33.2. 데스크톱 환경과 창 관리자

X 서버가 실행되면 X 클라이언트 어플리케이션은 X 서버로 연결되고 사용자에 필요한 GUI를 생성합니다. Red Hat Enterprise Linux에 사용 가능한 GUI는 기본적인 Tab Window Manager부터 거의 모든 Red Hat Enterprise Linux 사용자에게 익숙한 진보된 대화식 GNOME 데스크톱 환경에 이르기까지 다양합니다.

고급의 종합적인 GUI를 생성하려면 다음 두 개의 주요 X 클라이언트 클래스가 반드시 X 서버로 연결되어야 합니다: 데스크톱 환경과 창 관리자

33.2.1. 데스크톱 환경

데스크톱 환경은 다양한 X 클라이언트를 통합하여 일반 그래픽형식 사용자 환경과 개발 플랫폼을 생성합니다.

데스크톱 환경은 X 클라이언트와 다른 실행 프로세스가 의사소통할 수 있는 고급 기능을 갖추고 있습니다. 또한, 데스크톱 환경에서 작동하도록 쓰여진 모든 어플리케이션이 드래그-앤-드랍 기능과 같은 고급 작업을 수행할 수 있도록 구성합니다.

Red Hat Enterprise Linux은 두 가지의 데스크톱 환경을 제공합니다:

- GNOME — Red Hat Enterprise Linux에 사용되는 기본 데스크톱 환경으로서 GTK+ 2 그래픽형식 툴킷을 기반으로 합니다.

- KDE — 대안적인 데스크톱 환경으로서 Qt 3 그래픽형식 툴킷을 기반으로 합니다.

GNOME과 KDE 모두 워드 프로세서, 스프레드시트, 웹 브라우저와 같은 생산성있는 고급 어플리케이션을 갖추고 있으며, GUI 외관을 사용자 요구에 맞게 변경할 수 있는 도구도 제공합니다. 또한, GTK+ 2와 Qt 라이브러리가 모두 존재하여 GNOME에서 KDE 어플리케이션을 실행하거나 KDE에서 GNOME 어플리케이션을 실행할 수 있습니다.

33.2.2. 창 관리자

창 관리자는 X 클라이언트 프로그램으로서 데스크톱 환경의 일부분 또는 특정 상황에서 독립형으로 작동합니다. 창 관리자의 주요 목적은 그래픽형식 창이 어떻게 위치하고, 크기가 변경되고, 이

동하는지 제어하는 데 있습니다. 창 관리자는 주제 바, 창 중심 행동, 사용자 지정 키, 마우스 버튼 설정 등도 제어할 수 있습니다.

Red Hat Enterprise Linux에는 네 개의 창 관리자가 포함되어 있습니다:

kwin
: KWin은 KDE에 기본으로 사용되며 사용자 테마를 지원하는 효율적인 창 관리자입니다.

metacity
: Metacity는 GNOME에 기본으로 사용되며 사용자 테마를 지원하는 간단하고 효율적인 창 관리자입니다. Metacity 창 관리자를 실행하려면 metacity 패키지를 실행해야 합니다.

mwm
: Motif 창 관리자(mwm)는 기본 독립형 창 관리자입니다. 독립형 창 관리자로 작동하도록 설계되었기 때문에 GNOME 또는 KDE와 같이 사용될 수 없습니다. Motif 창 관리자 창 관리자를 실행하려면 openmotif 패키지를 설치해야 합니다.

twm
: Tab 창 관리자 (twm는 다른 창 관리자의 가장 기본적인 도구 모음을 제공하는 창 관리자로서 독립형과 데스크톱 환경 모두에서 사용될 수 있습니다. Tab 창 관리자는 X11R7.1 배포판의 일부분으로 설치됩니다.

To run any of the aforementioned window managers, you will first need to boot into Runlevel 3. For instructions on how to do this, refer to 17.1절. "런레벨 (runlevels)" .

Once you are logged in to Runlevel 3, you will be presented with a terminal prompt, not a graphical environment. To start a window manager, type xinit -e <path-to-window-manager> at the prompt.

<path-to-window-manager> is the location of the window manager binary file. The binary file can be located by typing which window-manager-name, where window-manager-name is the name of the window manager you want to run.

예를 들어:

```
~]# which twm
/usr/bin/twm
~]# xinit -e /usr/bin/twm
```

위의 첫 번째 명령어는 twm 창 관리자에 대한 절대 경로를 표시하며, 두 번째 명령어는 twm을 실행합니다.

창 관리자를 종료하려면 마지막 창을 닫고 Ctrl+Alt+Backspace를 누릅니다. 창 관리자를 종료하면 프롬프트에서 startx를 입력하여 런레벨 5로 돌아올 수 있습니다.

33.3. X 서버 구성 파일

X 서버는 하나의 바이너리 실행 파일(/usr/bin/Xorg)입니다. 관련 구성 파일은 /etc/X11/ 디렉토리 (심볼릭 링크 — X — 는 /usr/bin/Xorg를 가리킴)에 저장됩니다. X 서버에 필요한 구성 파일은 /etc/X11/xorg.conf입니다.

/usr/lib/xorg/modules/ 디렉토리는 실시간으로 자동 로드되는 X 서버 모듈을 포함하고 있습니다. 디폴트로 /usr/lib/xorg/modules/ 디렉토리의 일부 모듈만이 X 서버에 의해 자동으로 로드됩니다.

To load optional modules, they must be specified in the X server configuration file, /etc/X11/xorg.conf. For more information about loading modules, refer to 33.3.1.5절. "Module" .

When Red Hat Enterprise Linux 5.8 is installed, the configuration files for X are created using information gathered about the system hardware during the installation process.

33.3.1. xorg.conf

While there is rarely a need to manually edit the /etc/X11/xorg.conf file, it is useful to understand the various sections and optional parameters available, especially when troubleshooting.

33.3.1.1. 구조

The /etc/X11/xorg.conf file is comprised of many different sections which address specific aspects of the system hardware.

Each section begins with a Section "<section-name>" line (where <section-name> is the title for the section) and ends with an EndSection line. Each section contains lines that include option names and one or more option values. These are sometimes enclosed in double quotes (").

Lines beginning with a hash mark (#) are not read by the X server and are used for human-readable comments.

Some options within the /etc/X11/xorg.conf file accept a boolean switch which turns the feature on or off. Acceptable boolean values are:

- 1, on, true, or yes — Turns the option on.
- 0, off, false, or no — Turns the option off.

The following are some of the more important sections in the order in which they appear in a typical /etc/X11/xorg.conf file. More detailed information about the X server configuration file can be found in the xorg.conf man page.

33.3.1.2. ServerFlags

The optional ServerFlags section contains miscellaneous global X server settings. Any settings in this section may be overridden by options placed in the ServerLayout section (refer to 33.3.1.3절. “ServerLayout” for details).

Each entry within the ServerFlags section is on its own line and begins with the term Option followed by an option enclosed in double quotation marks (").

The following is a sample ServerFlags section:

```
Section "ServerFlags"
 Option "DontZap" "true"
EndSection
```

다음은 가장 많이 사용되는 옵션 목록을 보여줍니다:

- "DontZap" "<boolean>" — When the value of <boolean> is set to true, this setting prevents the use of the Ctrl+Alt+Backspace key combination to immediately terminate the X server.
- "DontZoom" "<boolean>" — When the value of <boolean> is set to true, this setting prevents cycling through configured video resolutions using the Ctrl+Alt+Keypad-Plus and Ctrl+Alt+Keypad-Minus key combinations.

33.3.1.3. ServerLayout

The ServerLayout section binds together the input and output devices controlled by the X server. At a minimum, this section must specify one output device and one input device. By default, a monitor (output device) and keyboard (input device) are specified.

The following example illustrates a typical ServerLayout section:

```
Section  "ServerLayout"
 Identifier     "Default Layout"
 Screen      0  "Screen0" 0 0
 InputDevice    "Mouse0" "CorePointer"
 InputDevice    "Keyboard0" "CoreKeyboard"
EndSection
```

The following entries are commonly used in the ServerLayout section:

- Identifier — Specifies a unique name for this ServerLayout section.
- Screen — Specifies the name of a Screen section to be used with the X server. More than one Screen option may be present.

  The following is an example of a typical Screen entry:

  ```
  Screen      0  "Screen0" 0 0
  ```

  The first number in this example Screen entry (0) indicates that the first monitor connector or head on the video card uses the configuration specified in the Screen section with the identifier "Screen0".

  An example of a Screen section with the identifier "Screen0" can be found in 33.3.1.9절. "Screen" .

  If the video card has more than one head, another Screen entry with a different number and a different Screen section identifier is necessary .

  The numbers to the right of "Screen0" give the absolute X and Y coordinates for the upper-left corner of the screen (0 0 by default).

- InputDevice — Specifies the name of an InputDevice section to be used with the X server.

  It is advisable that there be at least two InputDevice entries: one for the default mouse and one for the default keyboard. The options CorePointer and CoreKeyboard indicate that these are the primary mouse and keyboard.

- Option "<option-name>" — An optional entry which specifies extra parameters for the section. Any options listed here override those listed in the ServerFlags section.

  Replace <option-name> with a valid option listed for this section in the xorg.conf man page.

/etc/X11/xorg.conf 파일에 하나 이상의 ServerLayout 섹션을 설정하는 것이 가능하지만, 서버는 디폴트로 첫 번째 설정된 섹션만을 읽게 됩니다.

다른 ServerLayout 섹션이 있는 경우, X 세션을 시작할 때 명령 행 인수로 지정할 수 있습니다.

33.3.1.4. Files

The Files section sets paths for services vital to the X server, such as the font path. This is an optional section, these paths are normally detected automatically. This section may be used to override any automatically detected defaults.

The following example illustrates a typical Files section:

```
Section "Files"
 RgbPath       "/usr/share/X11/rgb.txt"
 FontPath      "unix/:7100"
EndSection
```

The following entries are commonly used in the Files section:

- RgbPath — Specifies the location of the RGB color database. This database defines all valid color names in X and ties them to specific RGB values.
- FontPath — Specifies where the X server must connect to obtain fonts from the xfs font server.

  By default, the FontPath is unix/:7100. This tells the X server to obtain font information using UNIX-domain sockets for inter-process communication (IPC) on port 7100.

  Refer to 33.4절. "Fonts" for more information concerning X and fonts.
- ModulePath — An optional parameter which specifies alternate directories which store X server modules.

33.3.1.5. Module

디폴트로 X 서버는 자동으로 /usr/lib/xorg/modules/ 디렉토리에서 다음 모듈을 로드합니다:

- extmod
- dbe
- glx
- freetype
- type1
- record
- dri

The default directory for loading these modules can be changed by specifying a different directory with the optional ModulePath parameter in the Files section. Refer to 33.3.1.4절. "Files" for more information on this section.

/etc/X11/xorg.conf 파일에 Module 섹션을 추가하면 X 서버가 디폴트 모듈 대신 Module 섹션에 기재된 모듈을 로드하도록 설정됩니다.

For example, the following typical Module section:

```
Section "Module"
 Load  "fbdevhw"
EndSection
```

instructs the X server to load the fbdevhw instead of the default modules.

마찬가지로, /etc/X11/xorg.conf 파일에 Module 섹션을 추가하면 로드하려는 모든 디폴트 모듈은 물론 모든 추가 모듈을 지정해야 합니다.

33.3.1.6. InputDevice

Each InputDevice section configures one input device for the X server. Systems typically have at least one InputDevice section for the keyboard. It is perfectly normal to have no entry for a mouse, as most mouse settings are automatically detected.

The following example illustrates a typical InputDevice section for a keyboard:

```
Section "InputDevice"
        Identifier  "Keyboard0"
        Driver        "kbd"
        Option        "XkbModel" "pc105"
        Option        "XkbLayout" "us"
EndSection
```

The following entries are commonly used in the InputDevice section:

- Identifier — Specifies a unique name for this InputDevice section. This is a required entry.

- Driver — Specifies the name of the device driver X must load for the device.

- Option — Specifies necessary options pertaining to the device.

  자동 검색된 디폴트 마우스 설정을 무시하고 새로운 마우스를 설정할 수 있습니다. xorg.conf 파일에 마우스를 추가할 때 주로 다음 옵션을 사용합니다:

  - Protocol — Specifies the protocol used by the mouse, such as IMPS/2.

  - Device — Specifies the location of the physical device.

  - Emulate3Buttons — Specifies whether to allow a two-button mouse to act like a three-button mouse when both mouse buttons are pressed simultaneously.

  Consult the xorg.conf man page for a list of valid options for this section.

33.3.1.7. Monitor

Each Monitor section configures one type of monitor used by the system. This is an optional entry as well, as most monitors are now automatically detected.

The easiest way to configure a monitor is to configure X during the installation process or by using the X Configuration Tool. For more information about using the X Configuration Tool, refer to 34장. X 윈도우 시스템 설정.

This example illustrates a typical Monitor section for a monitor:

```
Section "Monitor"
 Identifier     "Monitor0"
 VendorName     "Monitor Vendor"
 ModelName      "DDC Probed Monitor - ViewSonic G773-2"
 DisplaySize   320 240
 HorizSync      30.0 - 70.0
 VertRefresh   50.0 - 180.0
```

```
EndSection
```

경고

Be careful when manually editing values in the Monitor section of /etc/X11/xorg.conf. Inappropriate values can damage or destroy a monitor. Consult the monitor's documentation for a listing of safe operating parameters.

The following are commonly entries used in the Monitor section:

- Identifier — Specifies a unique name for this Monitor section. This is a required entry.
- VendorName — An optional parameter which specifies the vendor of the monitor.
- ModelName — An optional parameter which specifies the monitor's model name.
- DisplaySize — An optional parameter which specifies, in millimeters, the physical size of the monitor's picture area.
- HorizSync — Specifies the range of horizontal sync frequencies compatible with the monitor in kHz. These values help the X server determine the validity of built-in or specified Modeline entries for the monitor.
- VertRefresh — Specifies the range of vertical refresh frequencies supported by the monitor, in kHz. These values help the X server determine the validity of built in or specified Modeline entries for the monitor.
- Modeline — An optional parameter which specifies additional video modes for the monitor at particular resolutions, with certain horizontal sync and vertical refresh resolutions. Refer to the xorg.conf man page for a more detailed explanation of Modeline entries.
- Option "<option-name>" — An optional entry which specifies extra parameters for the section. Replace <option-name> with a valid option listed for this section in the xorg.conf man page.

33.3.1.8. Device

Each Device section configures one video card on the system. While one Device section is the minimum, additional instances may occur for each video card installed on the machine.

The best way to configure a video card is to configure X during the installation process or by using the X Configuration Tool. For more about using the X Configuration Tool, refer to 34장. X 윈도우 시스템 설정.

The following example illustrates a typical Device section for a video card:

```
Section "Device"
 Identifier    "Videocard0"
 Driver        "mga"
 VendorName    "Videocard vendor"
 BoardName     "Matrox Millennium G200"
 VideoRam      8192
 Option        "dpms"
EndSection
```

The following entries are commonly used in the Device section:

- Identifier — Specifies a unique name for this Device section. This is a required entry.
- Driver — Specifies which driver the X server must load to utilize the video card. A list of drivers can be found in /usr/share/hwdata/videodrivers, which is installed with the hwdata package.
- VendorName — An optional parameter which specifies the vendor of the video card.
- BoardName — An optional parameter which specifies the name of the video card.
- VideoRam — An optional parameter which specifies the amount of RAM available on the video card in kilobytes. This setting is only necessary for video cards the X server cannot probe to detect the amount of video RAM.
- BusID — An entry which specifies the bus location of the video card. On systems with only one video card a BusID entry is optional and may not even be present in the default /etc/X11/xorg.conf file. On systems with more than one video card, however, a BusID entry must be present.
- Screen — An optional entry which specifies which monitor connector or head on the video card the Device section configures. This option is only useful for video cards with multiple heads.

  If multiple monitors are connected to different heads on the same video card, separate Device sections must exist and each of these sections must have a different Screen value.

  Values for the Screen entry must be an integer. The first head on the video card has a value of 0. The value for each additional head increments this value by one.
- Option "<option-name>" — An optional entry which specifies extra parameters for the section. Replace <option-name> with a valid option listed for this section in the xorg.conf man page.

  One of the more common options is "dpms" (for Display Power Management Signaling, a VESA standard), which activates the Service Star energy compliance setting for the monitor.

33.3.1.9. Screen

Each Screen section binds one video card (or video card head) to one monitor by referencing the Device section and the Monitor section for each. While one Screen section is the minimum, additional instances may occur for each video card and monitor combination present on the machine.

The following example illustrates a typical Screen section:

```
Section "Screen"
 Identifier "Screen0"
 Device     "Videocard0"
 Monitor    "Monitor0"
 DefaultDepth     16
 SubSection "Display"
  Depth     24
  Modes     "1280x1024" "1280x960" "1152x864" "1024x768" "800x600" "640x480"
 EndSubSection
 SubSection "Display"
  Depth     16
  Modes     "1152x864" "1024x768" "800x600" "640x480"
 EndSubSection
EndSection
```

The following entries are commonly used in the Screen section:

- Identifier — Specifies a unique name for this Screen section. This is a required entry.
- Device — Specifies the unique name of a Device section. This is a required entry.
- Monitor — Specifies the unique name of a Monitor section. This is only required if a specific Monitor section is defined in the xorg.conf file. Normally, monitors are automatically detected.
- DefaultDepth — Specifies the default color depth in bits. In the previous example, 16 (which provides thousands of colors) is the default. Only one DefaultDepth is permitted, although this can be overridden with the Xorg command line option -depth <n>,where <n> is any additional depth specified.
- SubSection "Display" — Specifies the screen modes available at a particular color depth. The Screen section can have multiple Display subsections, which are entirely optional since screen modes are automatically detected.

  이 하위 섹션은 주로 자동으로 검색된 모드를 덮어쓰는 데 사용됩니다.
- Option "<option-name>" — An optional entry which specifies extra parameters for the section. Replace <option-name> with a valid option listed for this section in the xorg.conf man page.

33.3.1.10. DRI

The optional DRI section specifies parameters for the Direct Rendering Infrastructure (DRI). DRI is an interface which allows 3D software applications to take advantage of 3D hardware acceleration capabilities built into most modern video hardware. In addition, DRI can improve 2D performance via hardware acceleration, if supported by the video card driver.

DRI 그룹과 모드가 디폴트 값으로 자동으로 초기화되기 때문에 이 섹션은 거의 나타나지 않습니다. 다른 그룹 또는 모드를 설정하려면 xorg.conf 파일에 이 섹션을 추가하여 디폴트 값을 덮어쓸 수 있습니다.

The following example illustrates a typical DRI section:

```
Section "DRI"
 Group        0
 Mode         0666
EndSection
```

비디오 카드는 서로 다른 방식으로 DRI를 사용하므로 이 섹션을 추가할 때 http://dri.sourceforge.net/을 참조하시기 바랍니다.

33.4. Fonts

Red Hat Enterprise Linux uses two subsystems to manage and display fonts under X: Fontconfig and xfs.

새로운 Fontconfig 폰트 하위 시스템은 폰트 관리를 간단화하고 안티 앨리어싱(Anti-aliasing)과 같은 고급 화면 표시 기능을 제공합니다. 이 시스템은 Qt 3 또는 GTK+ 2 그래픽형식 툴킷을 사용하도록 설계된 어플리케이션에 자동으로 사용됩니다.

호환 목적으로 Red Hat Enterprise Linux는 코어 X 폰트 하위 시스템이라는 기존의 폰트 하위 시스템을 갖추고 있습니다. 15년이 넘은 코어 X 폰트 하위 시스템은 X Font Server(xfs)를 기반으로 합니다.

이번 섹션은 이러한 두 개의 시스템을 사용하여 X에서 어떻게 폰트를 설정하는지 소개합니다.

33.4.1. Fontconfig

Fontconfig 폰트 하위 시스템은 어플리케이션이 시스템의 폰트를 직접 사용하고 Xft 또는 다른 렌더링 기술을 사용하여 고급 안티 앨리어싱으로 Fontconfig 폰트를 렌더링하도록 설정합니다. 그래픽형식 어플리케이션은 Fontconfig으로 Xft 라이브러리를 사용하여 화면에 문자를 표시할 수 있습니다.

이제 Fontconfg/Xft 하위 시스템은 코어 X 폰트 하위 시스템을 대체합니다.

중요

The Fontconfig font subsystem does not yet work for OpenOffice.org, which uses its own font rendering technology.

It is important to note that Fontconfig uses the /etc/fonts/fonts.conf configuration file, which should not be edited by hand.

Tip

Due to the transition to the new font system, GTK+ 1.2 applications are not affected by any changes made via the Font Preferences dialog (accessed by selecting System (on the panel) > Preferences > Fonts). For these applications, a font can be configured by adding the following lines to the file ~/.gtkrc.mine:

```
style "user-font" {
 fontset = "<font-specification>"
}

widget_class "*" style "user-font"
```

Replace <font-specification> with a font specification in the style used by traditional X applications, such as -adobe-helvetica-medium-r-normal--*-120-*-*-*-*-*-*. A full list of core fonts can be obtained by running xlsfonts or created interactively using the xfontsel command.

33.4.1.1. Fontconfig에 폰트 추가

Fontconfig 하위 시스템에 새로운 폰트를 추가하는 것이 가장 일반적인 방식입니다.

1. To add fonts system-wide, copy the new fonts into the /usr/share/fonts/ directory. It is a good idea to create a new subdirectory, such as local/ or similar, to help distinguish between user-installed and default fonts.

   To add fonts for an individual user, copy the new fonts into the .fonts/ directory in the user's home directory.

2. Use the fc-cache command to update the font information cache, as in the following example:

```
fc-cache <path-to-font-directory>
```

In this command, replace <path-to-font-directory> with the directory containing the new fonts (either /usr/share/fonts/local/ or /home/<user>/.fonts/).

Tip

Individual users may also install fonts graphically, by typing fonts:/// into the Nautilus address bar, and dragging the new font files there.

중요

If the font file name ends with a .gz extension, it is compressed and cannot be used until uncompressed. To do this, use the gunzip command or double-click the file and drag the font to a directory in Nautilus.

33.4.2. 코어 X 폰트 시스템

For compatibility, Red Hat Enterprise Linux provides the core X font subsystem, which uses the X Font Server (xfs) to provide fonts to X client applications.

The X server looks for a font server specified in the FontPath directive within the Files section of the /etc/X11/xorg.conf configuration file. Refer to 33.3.1.4절. "Files" for more information about the FontPath entry.

The X server connects to the xfs server on a specified port to acquire font information. For this reason, the xfs service must be running for X to start. For more about configuring services for a particular runlevel, refer to 17장. 서비스로의 접근 통제.

33.4.2.1. xfs Configuration

The /etc/rc.d/init.d/xfs script starts the xfs server. Several options can be configured within its configuration file, /etc/X11/fs/config.

다음은 흔히 사용되는 옵션 목록을 보여줍니다:

- alternate-servers — Specifies a list of alternate font servers to be used if this font server is not available. A comma must separate each font server in a list.
- catalogue — Specifies an ordered list of font paths to use. A comma must separate each font path in a list.

  Use the string :unscaled immediately after the font path to make the unscaled fonts in that path load first. Then specify the entire path again, so that other scaled fonts are also loaded.
- client-limit — Specifies the maximum number of clients the font server services. The default is 10.

- clone-self — Allows the font server to clone a new version of itself when the client-limit is hit. By default, this option is on.
- default-point-size — Specifies the default point size for any font that does not specify this value. The value for this option is set in decipoints. The default of 120 corresponds to a 12 point font.
- default-resolutions — Specifies a list of resolutions supported by the X server. Each resolution in the list must be separated by a comma.
- deferglyphs — Specifies whether to defer loading glyphs (the graphic used to visually represent a font). To disable this feature use none, to enable this feature for all fonts use all, or to turn this feature on only for 16-bit fonts use 16.
- error-file — Specifies the path and file name of a location where xfs errors are logged.
- no-listen — Prevents xfs from listening to particular protocols. By default, this option is set to tcp to prevent xfs from listening on TCP ports for security reasons.

> **Tip**
>
> If xfs is used to serve fonts over the network, remove this line.

- port — Specifies the TCP port that xfs listens on if no-listen does not exist or is commented out.
- use-syslog — Specifies whether to use the system error log.

33.4.2.2. Adding Fonts to xfs

To add fonts to the core X font subsystem (xfs), follow these steps:

1. If it does not already exist, create a directory called /usr/share/fonts/local/ using the following command as root:

   ```
   mkdir /usr/share/fonts/local/
   ```

   If creating the /usr/share/fonts/local/ directory is necessary, it must be added to the xfs path using the following command as root:

   ```
   chkfontpath --add /usr/share/fonts/local/
   ```

2. Copy the new font file into the /usr/share/fonts/local/ directory
3. 루트로 다음 명령어를 실행하여 폰트 정보를 업데이트합니다:

   ```
   ttmkfdir -d /usr/share/fonts/local/ -o /usr/share/fonts/local/fonts.scale
   ```

4. Reload the xfs font server configuration file by issuing the following command as root:

   ```
   service xfs reload
   ```

33.5. 런레벨과 X

Red Hat Enterprise Linux 설치 도구는 주로 시스템을 Runlevel 5라고 알려진 그래픽형식 로그인 환경으로 부팅되도록 설정합니다. 하지만, Runlevel 3라는 텍스트-전용 다중-사용자 모드로 부팅하여 X 세션을 시작할 수도 있습니다.

For more information about runlevels, refer to 17.1절. "런레벨 (runlevels)" .

다음 섹션은 X가 런레벨 3와 런레벨 5에서 어떻게 시작하는지 보여줍니다.

33.5.1. 런레벨 3

When in runlevel 3, the best way to start an X session is to log in and type startx. The startx command is a front-end to the xinit command, which launches the X server (Xorg) and connects X client applications to it. Because the user is already logged into the system at runlevel 3, startx does not launch a display manager or authenticate users. Refer to 33.5.2절. "런레벨 5" for more information about display managers.

When the startx command is executed, it searches for the .xinitrc file in the user's home directory to define the desktop environment and possibly other X client applications to run. If no .xinitrc file is present, it uses the system default /etc/X11/xinit/xinitrc file instead.

The default xinitrc script then searches for user-defined files and default system files, including .Xresources, .Xmodmap, and .Xkbmap in the user's home directory, and Xresources, Xmodmap, and Xkbmap in the /etc/X11/ directory. The Xmodmap and Xkbmap files, if they exist, are used by the xmodmap utility to configure the keyboard. The Xresources file is read to assign specific preference values to applications.

After setting these options, the xinitrc script executes all scripts located in the /etc/X11/xinit/xinitrc.d/ directory. One important script in this directory is xinput.sh, which configures settings such as the default language.

Next, the xinitrc script attempts to execute .Xclients in the user's home directory and turns to /etc/X11/xinit/Xclients if it cannot be found. The purpose of the Xclients file is to start the desktop environment or, possibly, just a basic window manager. The .Xclients script in the user's home directory starts the user-specified desktop environment in the .Xclients-default file. If .Xclients does not exist in the user's home directory, the standard /etc/X11/xinit/Xclients script attempts to start another desktop environment, trying GNOME first and then KDE followed by twm.

런레벨 3에서 사용자는 X 세션이 끝나면 텍스트 모드 사용자 세션으로 돌아갑니다.

33.5.2. 런레벨 5

시스템이 런레벨 5로 부팅될 때 화면 관리자라는 특수한 X 클라이언트 어플리케이션이 실행됩니다. 사용자는 데스크톱 환경 또는 창 관리자가 실행되기 전에 반드시 화면 관리자를 사용하여 인증해야 합니다.

시스템에 설치된 데스크톱 환경에 따라 세 가지 다른 화면 관리자를 사용하여 사용자 인증을 수행할 수 있습니다.

- GNOME — The default display manager for Red Hat Enterprise Linux, GNOME allows the user to configure language settings, shutdown, restart or log in to the system.

- KDE — KDE's display manager which allows the user to shutdown, restart or log in to the system.

- xdm — A very basic display manager which only lets the user log in to the system.

When booting into runlevel 5, the prefdm script determines the preferred display manager by referencing the /etc/sysconfig/desktop file. A list of options for this file is available in this file:

```
/usr/share/doc/initscripts-<version-number>/sysconfig.txt
```

where <version-number> is the version number of the initscripts package.

Each of the display managers reference the /etc/X11/xdm/Xsetup_0 file to set up the login screen. Once the user logs into the system, the /etc/X11/xdm/GiveConsole script runs to assign ownership of the console to the user. Then, the /etc/X11/xdm/Xsession script runs to accomplish many of the tasks normally performed by the xinitrc script when starting X from runlevel 3, including setting system and user resources, as well as running the scripts in the /etc/X11/xinit/xinitrc.d/ directory.

Users can specify which desktop environment they want to utilize when they authenticate using the GNOME or KDE display managers by selecting it from the Sessions menu item (accessed by selecting System (on the panel) > Preferences > More Preferences > Sessions). If the desktop environment is not specified in the display manager, the /etc/X11/xdm/Xsession script checks the .xsession and .Xclients files in the user's home directory to decide which desktop environment to load. As a last resort, the /etc/X11/xinit/Xclients file is used to select a desktop environment or window manager to use in the same way as runlevel 3.

When the user finishes an X session on the default display (:0) and logs out, the /etc/X11/xdm/TakeConsole script runs and reassigns ownership of the console to the root user. The original display manager, which continues running after the user logged in, takes control by spawning a new display manager. This restarts the X server, displays a new login window, and starts the entire process over again.

사용자는 런레벨 5의 X에서 로그아웃하면 화면 관리자로 돌아옵니다.

For more information on how display managers control user authentication, refer to the /usr/share/doc/gdm-<version-number>/README (where <version-number> is the version number for the gdm package installed) and the xdm man page.

33.6. 추가 자료

X 서버, X 서버로 연결되는 X 클라이언트, 데스크톱 환경과 창 관리자 종류에 관한 상세하고 방대한 정보가 제공됩니다.

33.6.1. 설치된 문서 자료

- /usr/share/X11/doc/ — contains detailed documentation on the X Window System architecture, as well as how to get additional information about the Xorg project as a new user.
- man xorg.conf — Contains information about the xorg.conf configuration files, including the meaning and syntax for the different sections within the files.
- man Xorg — Describes the Xorg display server.

33.6.2. 유용한 웹사이트

- http://www.X.org/ — X Window System의 X11R7.1 배포판을 생산하는 X.Org 단체의 홈 페이지입니다. X11R7.1 배포판은 Red Hat Enterprise Linux에 내장되어 필수 하드웨어를 제어하고 GUI 환경을 제공합니다.
- http://dri.sourceforge.net/ — DRI(Direct Rendering Infrastructure) 프로젝트 홈 페이지입니다. DRI는 X의 주요 하드웨어 3D 가속 구성 요소입니다.
- http://www.gnome.org/[1] — GNOME 프로젝트의 홈 페이지입니다.
- http://www.kde.org/[2] — KDE 데스크톱 환경의 홈 페이지입니다.

[1] http://www.gnome.org

[2] http://www.kde.org

X 윈도우 시스템 설정

During installation, the system's monitor, video card, and display settings are configured. To change any of these settings after installation, use the X Configuration Tool.

To start the X Configuration Tool, go to System (on the panel) > Administration > Display, or type the command system-config-display at a shell prompt (for example, in an XTerm or GNOME terminal). If the X Window System is not running, a small version of X is started to run the program.

설정을 변경한 후에는 그래픽 데스크탑에서 로그아웃한 후 다시 로그인하면 변경 사항이 적용됩니다.

34.1. Display Settings

The Settings tab allows users to change the resolution and color depth. The display of a monitor consists of tiny dots called pixels. The number of pixels displayed at one time is called the resolution. For example, the resolution 1024x768 means that 1024 horizontal pixels and 768 vertical pixels are used. The higher the resolution values, the more images the monitor can display at one time.

화면 색상수에 따라서 화면에 표시되는 색상의 수가 달라집니다. 색상수가 높을수록 색상 대비가 더 뚜렷해집니다.

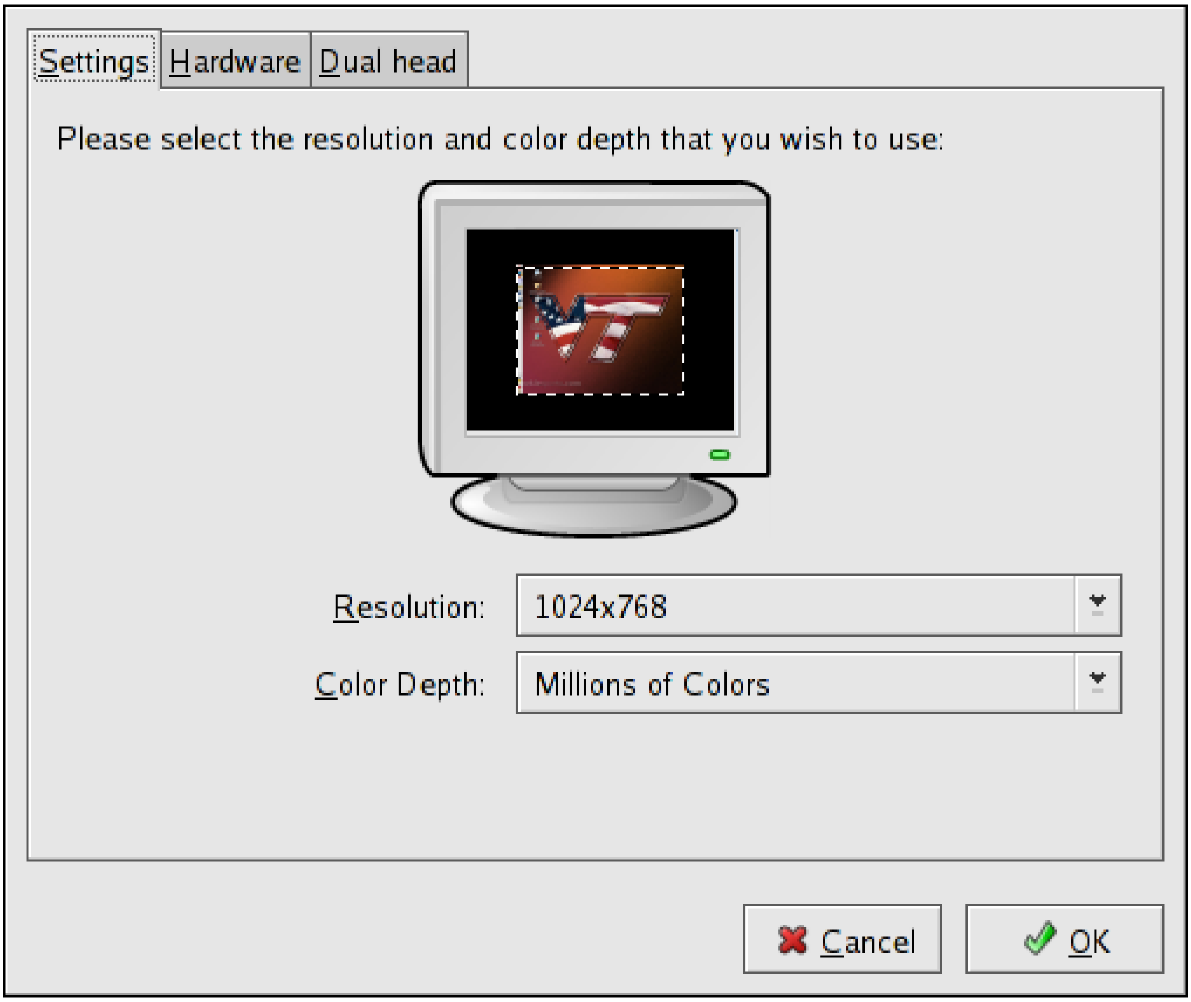

Please select the resolution and color depth that you wish to use:

Resolution: 1024x768

Color Depth: Millions of Colors

그림 34.1. Display Settings

34.2. Display Hardware Settings

When the X Configuration Tool is started, it probes the monitor and video card. If the hardware is probed properly, the information for it is shown on the Hardware tab as shown in 그림 34.2. "Display Hardware Settings" .

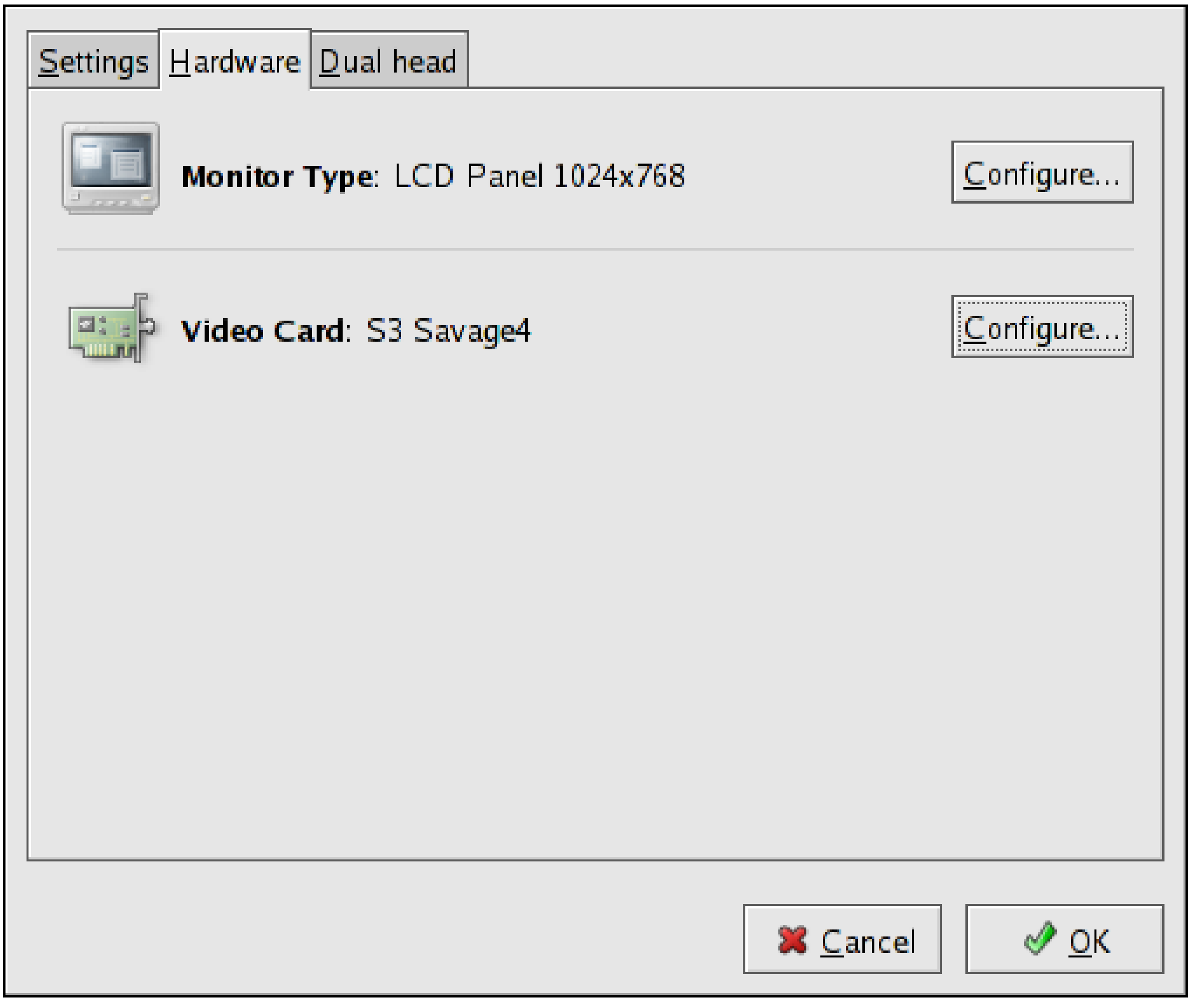

Monitor Type: LCD Panel 1024x768

Video Card: S3 Savage4

그림 34.2. Display Hardware Settings

To change the monitor type or any of its settings, click the corresponding Configure button. To change the video card type or any of its settings, click the Configure button beside its settings.

34.3. Dual Head Display Settings

If multiple video cards are installed on the system, dual head monitor support is available and is configured via the Dual head tab, as shown in 그림 34.3. "Dual Head Display Settings" .

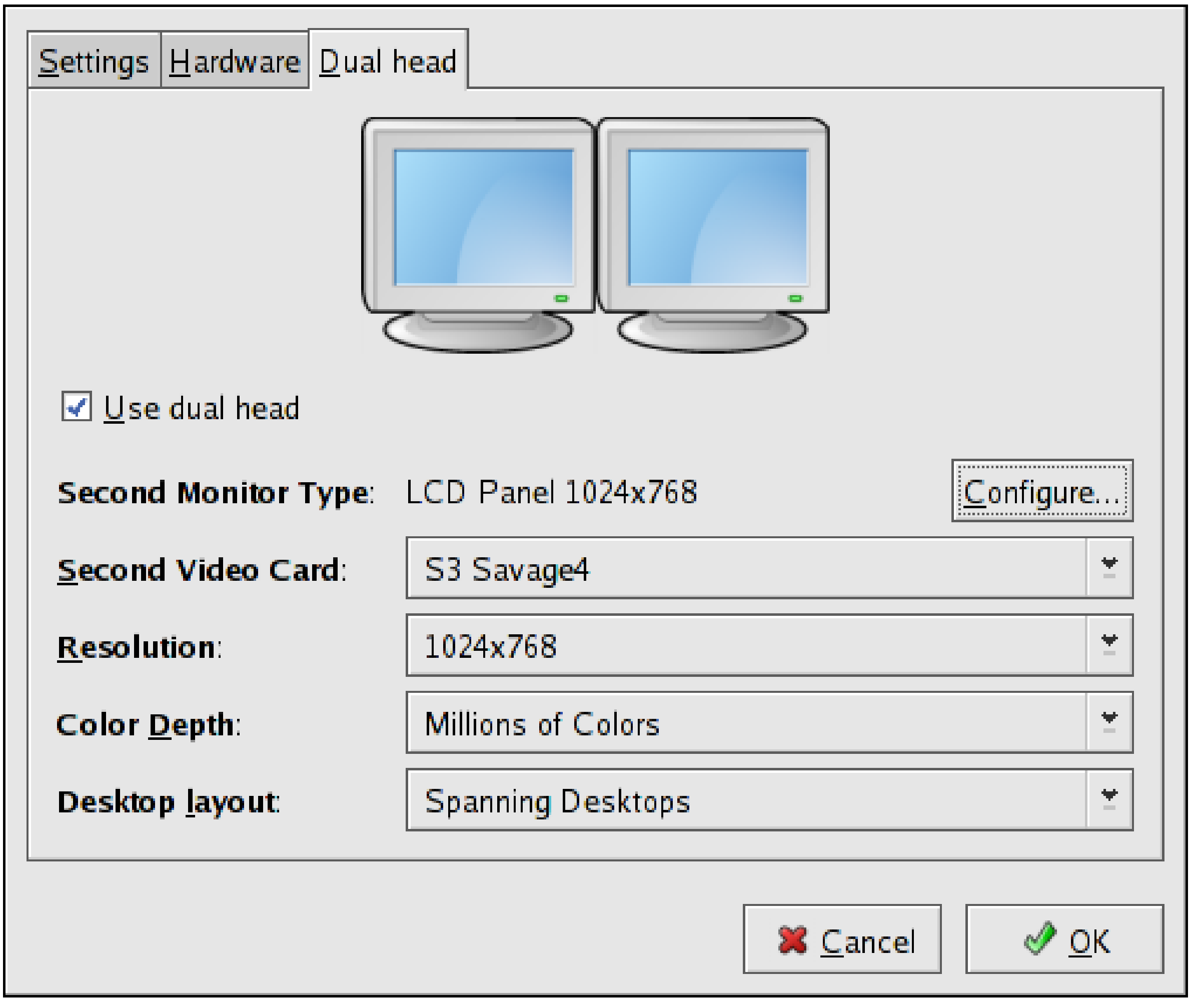

☑ Use dual head

Second Monitor Type: LCD Panel 1024x768

Second Video Card: S3 Savage4

Resolution: 1024x768

Color Depth: Millions of Colors

Desktop layout: Spanning Desktops

그림 34.3. Dual Head Display Settings

To enable use of Dual head, check the Use dual head checkbox.

To configure the second monitor type, click the corresponding Configure button. You can also configure the other Dual head settings by using the corresponding drop-down list.

For the Desktop layout option, selecting Spanning Desktops allows both monitors to use an enlarged usable workspace. Selecting Individual Desktops shares the mouse and keyboard among the displays, but restricts windows to a single display.

사용자 및 그룹

사용자 및 그룹 관리는 Red Hat Enterprise Linux 시스템 관리에 반드시 필요한 중요 요소입니다.

사용자는 사람(실제 사용자에 해당하는 계정 의미) 또는 사용할 특정 어플리케이션에 필요한 계정을 말합니다.

그룹은 일반 목적으로 사용자를 한데 묶는 논리적 표현 방식의 단체를 뜻합니다. 그룹에 속한 사용자는 그룹이 소유한 파일을 읽고, 쓰고, 실행할 수 있습니다.

모든 사용자와 그룹은 각각 userid(UID)와 groupid(GID)라고 부르는 고유한 숫자 식별 번호를 가지고 있습니다.

파일을 생성하는 사용자는 파일의 소유자와 그룹 소유자가 됩니다. 파일에는 소유자, 그룹, 기타 사용자에게 개별적인 읽기, 쓰기, 실행 권한이 할당됩니다. 파일 소유자는 루트에 의해서만 변경 가능하며, 파일 권한은 루트와 파일 소유자에 의해서 변경 가능합니다.

Red Hat Enterprise Linux also supports access control lists (ACLs) for files and directories which allow permissions for specific users outside of the owner to be set. For more information about ACLs, refer to 9장. Access Control Lists.

35.1. 사용자와 그룹 설정

The User Manager allows you to view, modify, add, and delete local users and groups.

To use the User Manager, you must be running the X Window System, have root privileges, and have the system-config-users RPM package installed. To start the User Manager from the desktop, go to System (on the panel) > Administration > Users & Groups. You can also type the command system-config-users at a shell prompt (for example, in an XTerm or a GNOME terminal).

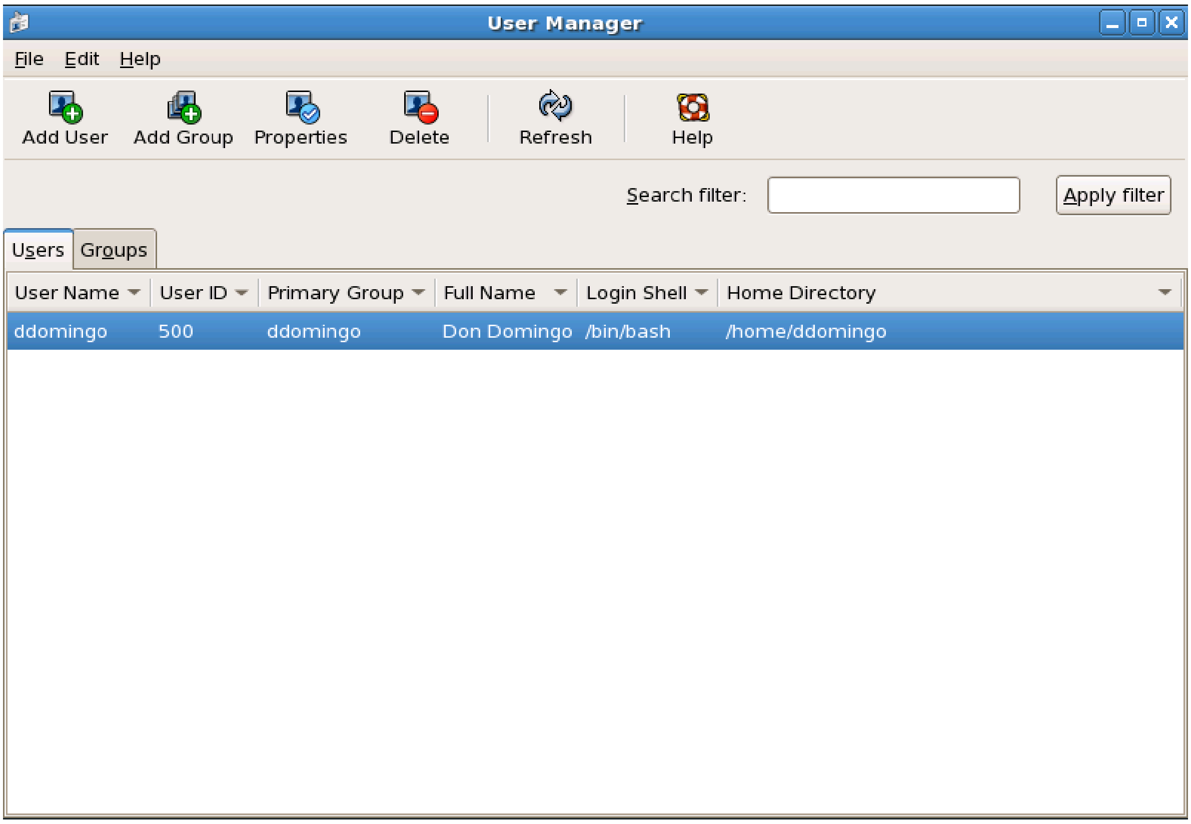

User Name	User ID	Primary Group	Full Name	Login Shell	Home Directory
ddomingo	500	ddomingo	Don Domingo	/bin/bash	/home/ddomingo

그림 35.1. User Manager

To view a list of local users on the system, click the Users tab. To view a list of local groups on the system, click the Groups tab.

To find a specific user or group, type the first few letters of the name in the Search filter field. Press Enter or click the Apply filter button. The filtered list is displayed.

사용자와 그룹을 정렬하려면 열 이름을 클릭합니다. 열 값에 맞게 사용자 또는 그룹이 정렬되어 표시됩니다.

Red Hat Enterprise Linux reserves user IDs below 500 for system users. By default, User Manager does not display system users. To view all users, including the system users, go to Edit > Preferences and uncheck Hide system users and groups from the dialog box.

35.1.1. 새로운 사용자 추가

To add a new user, click the Add User button. A window as shown in 그림 35.2. "새로운 사용자" appears. Type the username and full name for the new user in the appropriate fields. Type the user's password in the Password and Confirm Password fields. The password must be at least six characters.

Tip

권한이 주어지지 않은 상태에서 계정에 접근하기 쉽지 않도록 추측하기 어려운 최대한 긴 암호를 사용하는 것이 좋습니다. 또한, 사전 용어를 기반으로 한 암호 대신에 문자, 숫자, 특수 문자를 조합한 암호를 사용하는 것이 좋습니다.

Select a login shell. If you are not sure which shell to select, accept the default value of /bin/bash. The default home directory is /home/<username>/. You can change the home directory that is created for the user, or you can choose not to create the home directory by unselecting Create home directory.

If you select to create the home directory, default configuration files are copied from the /etc/skel/ directory into the new home directory.

Red Hat Enterprise Linux uses a user private group (UPG) scheme. The UPG scheme does not add or change anything in the standard UNIX way of handling groups; it offers a new convention. Whenever you create a new user, by default, a unique group with the same name as the user is created. If you do not want to create this group, unselect Create a private group for the user.

To specify a user ID for the user, select Specify user ID manually. If the option is not selected, the next available user ID above 500 is assigned to the new user. Because Red Hat Enterprise Linux reserves user IDs below 500 for system users, it is not advisable to manually assign user IDs 1-499.

Click OK to create the user.

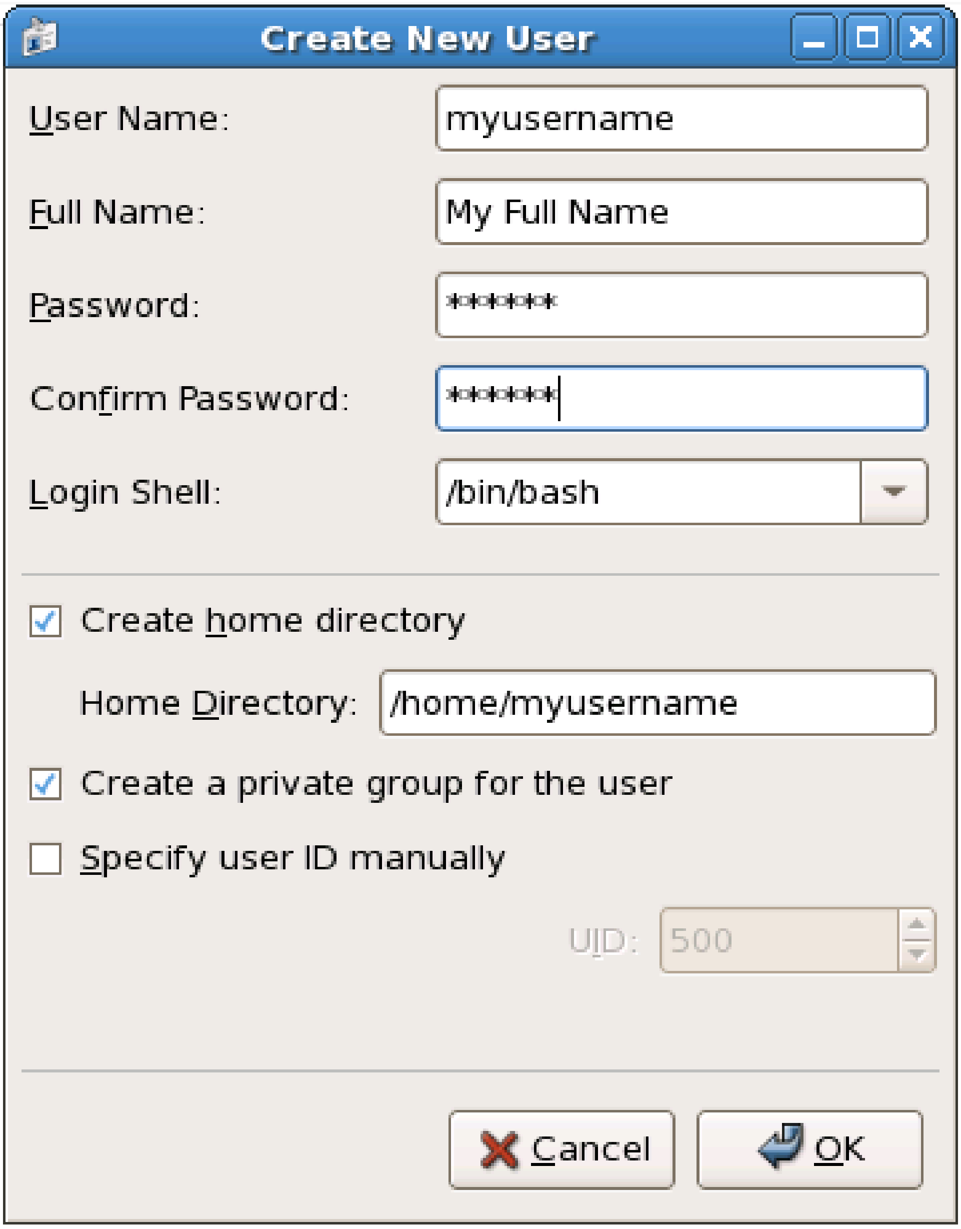
Create New User

User Name: myusername

Full Name: My Full Name

Password: *******

Confirm Password: *******

Login Shell: /bin/bash

☑ Create home directory

Home Directory: /home/myusername

☑ Create a private group for the user

□ Specify user ID manually

UID: 500

그림 35.2. 새로운 사용자

To configure more advanced user properties, such as password expiration, modify the user's properties after adding the user. Refer to 35.1.2절. "사용자 등록 정보 변경" for more information.

35.1.2. 사용자 등록 정보 변경

To view the properties of an existing user, click on the Users tab, select the user from the user list, and click Properties from the menu (or choose File > Properties from the pulldown menu). A window similar to 그림 35.3. "사용자 등록 정보" appears.

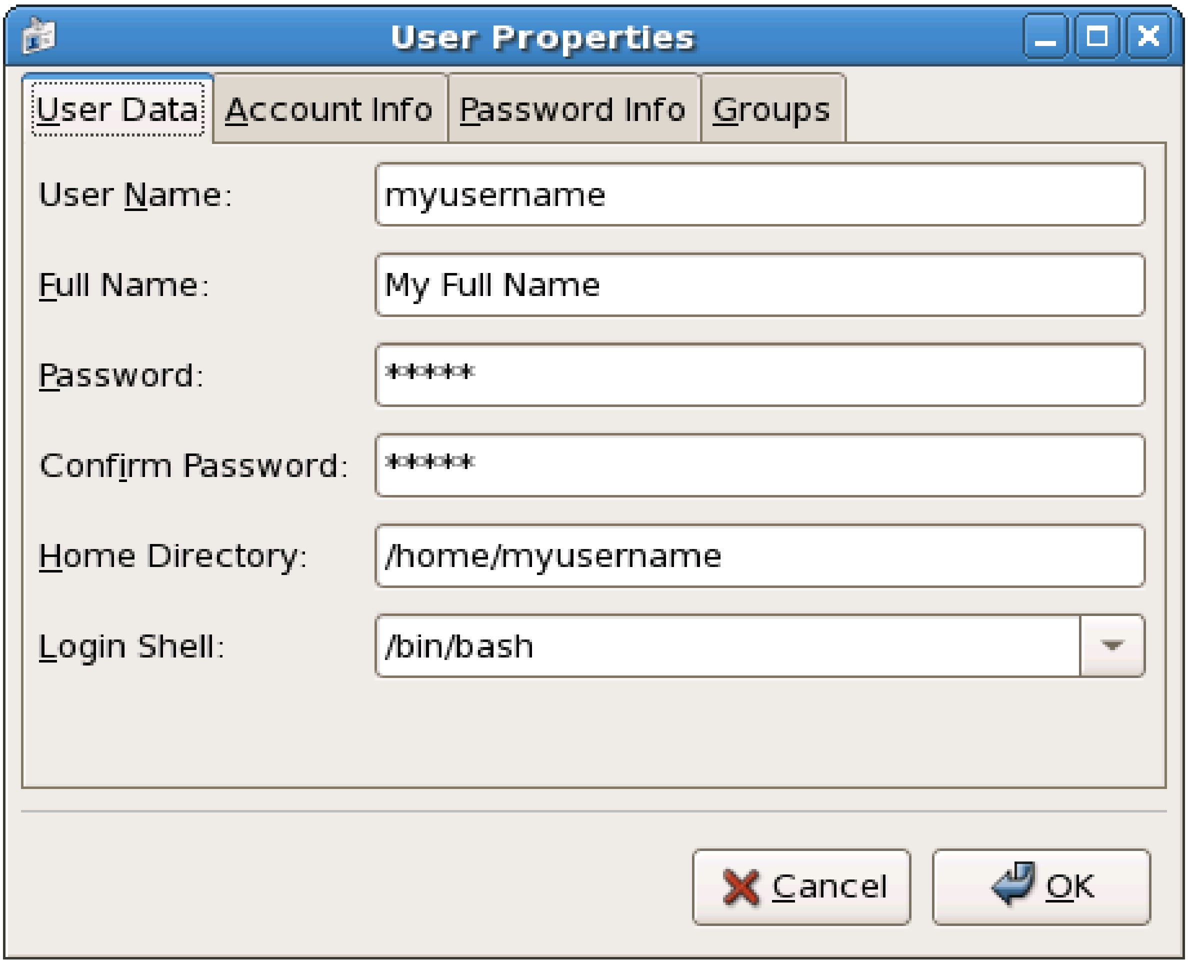

User Name: myusername

Full Name: My Full Name

Password: ******

Confirm Password: ******

Home Directory: /home/myusername

Login Shell: /bin/bash

그림 35.3. 사용자 등록 정보

The User Properties window is divided into multiple tabbed pages:

- User Data — Shows the basic user information configured when you added the user. Use this tab to change the user's full name, password, home directory, or login shell.
- Account Info — Select Enable account expiration if you want the account to expire on a certain date. Enter the date in the provided fields. Select Local password is locked to lock the user account and prevent the user from logging into the system.
- Password Info — Displays the date that the user's password last changed. To force the user to change passwords after a certain number of days, select Enable password expiration and enter a desired value in the Days before change required: field. The number of days before the user's password expires, the number of days before the user is warned to change passwords, and days before the account becomes inactive can also be changed.
- Groups — Allows you to view and configure the Primary Group of the user, as well as other groups that you want the user to be a member of.

35.1.3. 새로운 그룹 추가

To add a new user group, click the Add Group button. A window similar to 그림 35.4. "새로운 그룹" appears. Type the name of the new group to create. To specify a group ID for the new group, select Specify group ID manually and select the GID. Note that Red Hat Enterprise Linux also reserves group IDs lower than 500 for system groups.

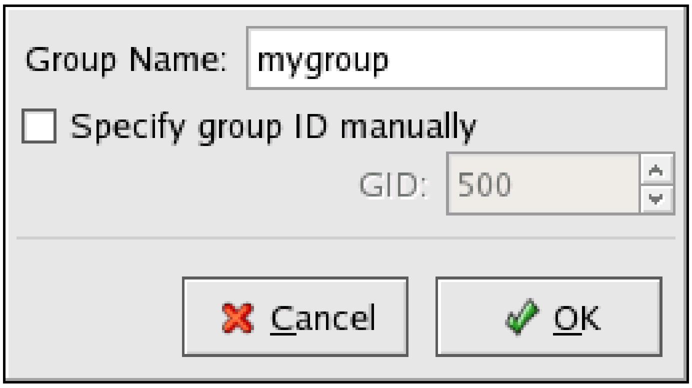

Group Name: mygroup

□ Specify group ID manually

GID: 500

그림 35.4. 새로운 그룹

Click OK to create the group. The new group appears in the group list.

35.1.4. 그룹 등록 정보 변경

To view the properties of an existing group, select the group from the group list and click Properties from the menu (or choose File > Properties from the pulldown menu). A window similar to 그림 35.5. "그룹 등록 정보" appears.

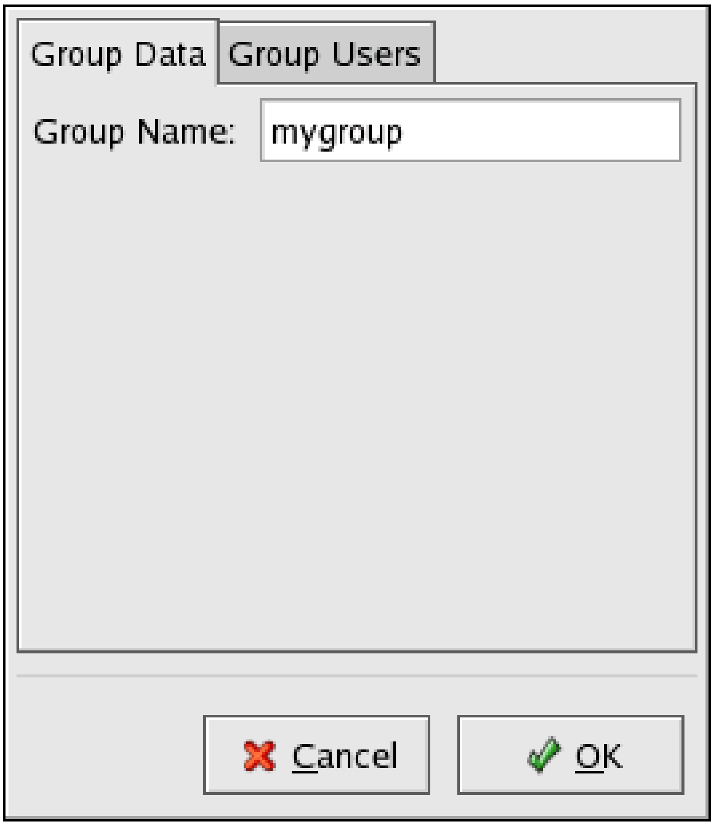

Group Name: mygroup

그림 35.5. 그룹 등록 정보

The Group Users tab displays which users are members of the group. Use this tab to add or remove users from the group. Click OK to save your changes.

35.2. 사용자 및 그룹 관리 도구

사용자와 그룹을 관리하는 것은 지루한 작업일 수 있습니다. 따라서, Red Hat Enterprise Linux는 사용자와 그룹을 관리하는 데 간편한 도구 및 방식을 제공합니다.

The easiest way to manage users and groups is through the graphical application, User Manager (system-config-users). For more information on User Manager, refer to 35.1절. "사용자와 그룹 설정" .

또한, 다음과 같은 명령행 도구를 사용하여 사용자와 그룹을 관리할 수도 있습니다:

- useradd, usermod, and userdel — Industry-standard methods of adding, deleting and modifying user accounts
- groupadd, groupmod, and groupdel — Industry-standard methods of adding, deleting, and modifying user groups
- gpasswd — Industry-standard method of administering the /etc/group file
- pwck, grpck — Tools used for the verification of the password, group, and associated shadow files
- pwconv, pwunconv — Tools used for the conversion of passwords to shadow passwords and back to standard passwords

35.2.1. 명령행 설정

명령행 도구를 사용하고자 하거나 X Window System이 설치되지 않았을 때 이번 섹션을 통해 사용자와 그룹을 구성할 수 있습니다.

35.2.2. 사용자 추가

시스템에 사용자 추가:

1. Issue the useradd command to create a locked user account:

   ```
   useradd <username>
   ```

2. Unlock the account by issuing the passwd command to assign a password and set password aging guidelines:

   ```
   passwd <username>
   ```

Command line options for useradd are detailed in 표 35.1. "useradd Command Line Options" .

표 35.1. useradd Command Line Options

옵션	내용
-c '<comment>'	<comment> can be replaced with any string. This option is generally used to specify the full name of a user.
-d <home-dir>	Home directory to be used instead of default /home/<username>/
-e <date>	계정이 비활성화되는 날짜를 YYYY-MM-DD 형식으로 지정합니다.

옵션	내용
-f <days>	Number of days after the password expires until the account is disabled. If 0 is specified, the account is disabled immediately after the password expires. If -1 is specified, the account is not be disabled after the password expires.
-g <group-name>	Group name or group number for the user's default group. The group must exist prior to being specified here.
-G <group-list>	사용자가 속할 추가 그룹 이름 또는 그룹 번호를 쉼표로 구분해 써넣을 수 있습니다. 이 옵션을 사용할 때 그룹이 이미 존재해야 합니다.
-m	홈 디렉토리가 존재하지 않을 때 생성합니다.
-M	홈 디렉토리를 생성하지 않습니다.
-n	사용자에 대한 UPG(사용자 개인 그룹)을 생성하지 않습니다.
-r	홈 디렉토리가 없고 500보다 작은 UID의 시스템 계정을 생성합니다.
-p <password>	The password encrypted with crypt
-s	User's login shell, which defaults to /bin/bash
-u <uid>	사용자에 대한 사용자 ID를 지정합니다. 사용자 ID는 499보다 큰 고유 숫자로 지정됩니다.

35.2.3. Adding a Group

To add a group to the system, use the command groupadd:

```
groupadd <group-name>
```

Command line options for groupadd are detailed in 표 35.2. “groupadd Command Line Options” .

표 35.2. groupadd Command Line Options

옵션	내용
-g <gid>	Group ID for the group, which must be unique and greater than 499
-r	Create a system group with a GID less than 500
-f	When used with -g <gid> and <gid> already exists, groupadd will choose another unique <gid> for the group.

35.2.4. Password Aging

For security reasons, it is advisable to require users to change their passwords periodically. This can be done when adding or editing a user on the Password Info tab of the User Manager.

To configure password expiration for a user from a shell prompt, use the chage command with an option from 표 35.3. “chage Command Line Options” , followed by the username.

Important

Shadow passwords must be enabled to use the chage command. For more information, see 35.6절. “Shadow Passwords” .

표 35.3. chage Command Line Options

옵션	내용
-m <days>	Specifies the minimum number of days between which the user must change passwords. If the value is 0, the password does not expire.
-M <days>	Specifies the maximum number of days for which the password is valid. When the number of days specified by this option plus the number of days specified with the -d option is less than the current day, the user must change passwords before using the account.
-d <days>	Specifies the number of days since January 1, 1970 the password was changed
-I <days>	Specifies the number of inactive days after the password expiration before locking the account. If the value is 0, the account is not locked after the password expires.
-E <date>	Specifies the date on which the account is locked, in the format YYYY-MM-DD. Instead of the date, the number of days since January 1, 1970 can also be used.
-W <days>	Specifies the number of days before the password expiration date to warn the user.
-l	Lists current account aging settings.

Tip

If the chage command is followed directly by a username (with no options), it displays the current password aging values and allows them to be changed interactively.

You can configure a password to expire the first time a user logs in. This forces users to change passwords immediately.

1. Set up an initial password — There are two common approaches to this step. The administrator can assign a default password or assign a null password.

   To assign a default password, use the following steps:

   - Start the command line Python interpreter with the python command. It displays the following:

```
Python 2.4.3 (#1, Jul 21 2006, 08:46:09)
[GCC 4.1.1 20060718 (Red Hat 4.1.1-9)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

   - At the prompt, type the following commands. Replace <password> with the password to encrypt and <salt> with a random combination of at least 2 of the following: any alphanumeric character, the slash (/) character or a dot (.):

     ```
     import crypt
     print crypt.crypt("<password>","<salt>")
     ```

     The output is the encrypted password, similar to '12CsGd8FRcMSM'.

   - Press Ctrl-D to exit the Python interpreter.

   - At the shell, enter the following command (replacing <encrypted-password> with the encrypted output of the Python interpreter):

     ```
     usermod -p "<encrypted-password>" <username>
     ```

   Alternatively, you can assign a null password instead of an initial password. To do this, use the following command:

   ```
   usermod -p "" username
   ```

   Caution

   Using a null password, while convenient, is a highly unsecure practice, as any third party can log in first an access the system using the unsecure username. Always make sure that the user is ready to log in before unlocking an account with a null password.

2. Force immediate password expiration — Type the following command:

   ```
   chage -d 0 username
   ```

   This command sets the value for the date the password was last changed to the epoch (January 1, 1970). This value forces immediate password expiration no matter what password aging policy, if any, is in place.

Upon the initial log in, the user is now prompted for a new password.

35.2.5. Explaining the Process

The following steps illustrate what happens if the command useradd juan is issued on a system that has shadow passwords enabled:

1. A new line for juan is created in /etc/passwd. The line has the following characteristics:

   - It begins with the username juan.

   - There is an x for the password field indicating that the system is using shadow passwords.

- A UID greater than 499 is created. (Under Red Hat Enterprise Linux, UIDs and GIDs below 500 are reserved for system use.)
- A GID greater than 499 is created.
- The optional GECOS information is left blank.
- The home directory for juan is set to /home/juan/.
- The default shell is set to /bin/bash.

2. A new line for juan is created in /etc/shadow. The line has the following characteristics:

   - It begins with the username juan.
   - Two exclamation points (!!) appear in the password field of the /etc/shadow file, which locks the account.

   > **Note**
   >
   > If an encrypted password is passed using the -p flag, it is placed in the /etc/shadow file on the new line for the user.

   - The password is set to never expire.

3. A new line for a group named juan is created in /etc/group. A group with the same name as a user is called a user private group. For more information on user private groups, refer to 35.1.1 절. "새로운 사용자 추가" .

   The line created in /etc/group has the following characteristics:

   - It begins with the group name juan.
   - An x appears in the password field indicating that the system is using shadow group passwords.
   - The GID matches the one listed for user juan in /etc/passwd.

4. A new line for a group named juan is created in /etc/gshadow. The line has the following characteristics:

   - It begins with the group name juan.
   - An exclamation point (!) appears in the password field of the /etc/gshadow file, which locks the group.
   - All other fields are blank.

5. A directory for user juan is created in the /home/ directory. This directory is owned by user juan and group juan. However, it has read, write, and execute privileges only for the user juan. All other permissions are denied.

6. The files within the /etc/skel/ directory (which contain default user settings) are copied into the new /home/juan/ directory.

At this point, a locked account called juan exists on the system. To activate it, the administrator must next assign a password to the account using the passwd command and, optionally, set password aging guidelines.

35.3. Standard Users

표 35.4. "Standard Users" lists the standard users configured in the /etc/passwd file by an Everything installation. The groupid (GID) in this table is the primary group for the user. See 35.4절. "Standard Groups" for a listing of standard groups.

표 35.4. Standard Users

User	UID	GID	Home Directory	Shell
root	0	0	/root	/bin/bash
bin	1	1	/bin	/sbin/nologin
daemon	2	2	/sbin	/sbin/nologin
adm	3	4	/var/adm	/sbin/nologin
lp	4	7	/var/spool/lpd	/sbin/nologin
sync	5	0	/sbin	/bin/sync
shutdown	6	0	/sbin	/sbin/shutdown
halt	7	0	/sbin	/sbin/halt
mail	8	12	/var/spool/mail	/sbin/nologin
news	9	13	/etc/news	
uucp	10	14	/var/spool/uucp	/sbin/nologin
operator	11	0	/root	/sbin/nologin
games	12	100	/usr/games	/sbin/nologin
gopher	13	30	/var/gopher	/sbin/nologin
ftp	14	50	/var/ftp	/sbin/nologin
nobody	99	99	/	/sbin/nologin
rpm	37	37	/var/lib/rpm	/sbin/nologin
vcsa	69	69	/dev	/sbin/nologin
dbus	81	81	/	/sbin/nologin
ntp	38	38	/etc/ntp	/sbin/nologin
canna	39	39	/var/lib/canna	/sbin/nologin
nscd	28	28	/	/sbin/nologin
rpc	32	32	/	/sbin/nologin
postfix	89	89	/var/spool/postfix	/sbin/nologin
mailman	41	41	/var/mailman	/sbin/nologin
named	25	25	/var/named	/bin/false
amanda	33	6	var/lib/amanda/	/bin/bash
postgres	26	26	/var/lib/pgsql	/bin/bash
exim	93	93	/var/spool/exim	/sbin/nologin
sshd	74	74	/var/empty/sshd	/sbin/nologin

User	UID	GID	Home Directory	Shell
rpcuser	29	29	/var/lib/nfs	/sbin/nologin
nsfnobody	65534	65534	/var/lib/nfs	/sbin/nologin
pvm	24	24	/usr/share/pvm3	/bin/bash
apache	48	48	/var/www	/sbin/nologin
xfs	43	43	/etc/X11/fs	/sbin/nologin
gdm	42	42	/var/gdm	/sbin/nologin
htt	100	101	/usr/lib/im	/sbin/nologin
mysql	27	27	/var/lib/mysql	/bin/bash
webalizer	67	67	/var/www/usage	/sbin/nologin
mailnull	47	47	/var/spool/mqueue	/sbin/nologin
smmsp	51	51	/var/spool/mqueue	/sbin/nologin
squid	23	23	/var/spool/squid	/sbin/nologin
ldap	55	55	/var/lib/ldap	/bin/false
netdump	34	34	/var/crash	/bin/bash
pcap	77	77	/var/arpwatch	/sbin/nologin
radiusd	95	95	/	/bin/false
radvd	75	75	/	/sbin/nologin
quagga	92	92	/var/run/quagga	/sbin/login
wnn	49	49	/var/lib/wnn	/sbin/nologin
dovecot	97	97	/usr/libexec/dovecot	/sbin/nologin

35.4. Standard Groups

표 35.5. "Standard Groups" lists the standard groups configured by an Everything installation. Groups are stored in the /etc/group file.

표 35.5. Standard Groups

Group	GID	Members
root	0	root
bin	1	root, bin, daemon
daemon	2	root, bin, daemon
sys	3	root, bin, adm
adm	4	root, adm, daemon
tty	5	
disk	6	root
lp	7	daemon, lp
mem	8	
kmem	9	
wheel	10	root
mail	12	mail, postfix, exim

Group	GID	Members
news	13	news
uucp	14	uucp
man	15	
games	20	
gopher	30	
dip	40	
ftp	50	
lock	54	
nobody	99	
users	100	
rpm	37	
utmp	22	
floppy	19	
vcsa	69	
dbus	81	
ntp	38	
canna	39	
nscd	28	
rpc	32	
postdrop	90	
postfix	89	
mailman	41	
exim	93	
named	25	
postgres	26	
sshd	74	
rpcuser	29	
nfsnobody	65534	
pvm	24	
apache	48	
xfs	43	
gdm	42	
htt	101	
mysql	27	
webalizer	67	
mailnull	47	
smmsp	51	
squid	23	

Group	GID	Members
ldap	55	
netdump	34	
pcap	77	
quaggavt	102	
quagga	92	
radvd	75	
slocate	21	
wnn	49	
dovecot	97	
radiusd	95	

35.5. User Private Groups

Red Hat Enterprise Linux uses a user private group (UPG) scheme, which makes UNIX groups easier to manage.

A UPG is created whenever a new user is added to the system. A UPG has the same name as the user for which it was created and that user is the only member of the UPG.

UPGs make it safe to set default permissions for a newly created file or directory, allowing both the user and the group of that user to make modifications to the file or directory.

The setting which determines what permissions are applied to a newly created file or directory is called a umask and is configured in the /etc/bashrc file. Traditionally on UNIX systems, the umask is set to 022, which allows only the user who created the file or directory to make modifications. Under this scheme, all other users, including members of the creator's group, are not allowed to make any modifications. However, under the UPG scheme, this "group protection" is not necessary since every user has their own private group.

35.5.1. Group Directories

Many IT organizations like to create a group for each major project and then assign people to the group if they need to access that project's files. Using this traditional scheme, managing files has been difficult; when someone creates a file, it is associated with the primary group to which they belong. When a single person works on multiple projects, it is difficult to associate the right files with the right group. Using the UPG scheme, however, groups are automatically assigned to files created within a directory with the setgid bit set. The setgid bit makes managing group projects that share a common directory very simple because any files a user creates within the directory are owned by the group which owns the directory.

Let us say, for example, that a group of people need to work on files in the /usr/share/emacs/site-lisp/ directory. Some people are trusted to modify the directory, but certainly not everyone is trusted. First create an emacs group, as in the following command:

```
groupadd emacs
```

To associate the contents of the directory with the emacs group, type:

```
chown -R root.emacs /usr/share/emacs/site-lisp
```

Now, it is possible to add the proper users to the group with the gpasswd command:

```
gpasswd -a <username> emacs
```

To allow users to create files within the directory, use the following command:

```
chmod 775 /usr/share/emacs/site-lisp
```

When a user creates a new file, it is assigned the group of the user's default private group. Next, set the setgid bit, which assigns everything created in the directory the same group permission as the directory itself (emacs). Use the following command:

```
chmod 2775 /usr/share/emacs/site-lisp
```

At this point, because the default umask of each user is 002, all members of the emacs group can create and edit files in the /usr/share/emacs/site-lisp/ directory without the administrator having to change file permissions every time users write new files.

35.6. Shadow Passwords

In multiuser environments it is very important to use shadow passwords (provided by the shadow-utils package). Doing so enhances the security of system authentication files. For this reason, the installation program enables shadow passwords by default.

The following lists the advantages pf shadow passwords have over the traditional way of storing passwords on UNIX-based systems:

- Improves system security by moving encrypted password hashes from the world-readable /etc/passwd file to /etc/shadow, which is readable only by the root user.
- Stores information about password aging.
- Allows the use the /etc/login.defs file to enforce security policies.

Most utilities provided by the shadow-utils package work properly whether or not shadow passwords are enabled. However, since password aging information is stored exclusively in the /etc/shadow file, any commands which create or modify password aging information do not work.

The following is a list of commands which do not work without first enabling shadow passwords:

- chage
- gpasswd
- /usr/sbin/usermod -e or -f options
- /usr/sbin/useradd -e or -f options

35.7. Additional Resources

For more information about users and groups, and tools to manage them, refer to the following resources.

35.7.1. Installed Documentation

- Related man pages — There are a number of man pages for the various applications and configuration files involved with managing users and groups. Some of the more important man pages have been listed here:

  User and Group Administrative Applications

  - man chage — A command to modify password aging policies and account expiration.
  - man gpasswd — A command to administer the /etc/group file.
  - man groupadd — A command to add groups.
  - man grpck — A command to verify the /etc/group file.
  - man groupdel — A command to remove groups.
  - man groupmod — A command to modify group membership.
  - man pwck — A command to verify the /etc/passwd and /etc/shadow files.
  - man pwconv — A tool to convert standard passwords to shadow passwords.
  - man pwunconv — A tool to convert shadow passwords to standard passwords.
  - man useradd — A command to add users.
  - man userdel — A command to remove users.
  - man usermod — A command to modify users.

  Configuration Files

  - man 5 group — The file containing group information for the system.
  - man 5 passwd — The file containing user information for the system.
  - man 5 shadow — The file containing passwords and account expiration information for the system.

Printer Configuration

Printer Configuration Tool allows users to configure a printer. This tool helps maintain the printer configuration file, print spool directories, print filters, and printer classes.

Red Hat Enterprise Linux 5.8 uses the Common Unix Printing System (CUPS). If a system was upgraded from a previous Red Hat Enterprise Linux version that used CUPS, the upgrade process preserves the configured queues.

Using Printer Configuration Tool requires root privileges. To start the application, select System (on the panel) > Administration > Printing, or type the command system-config-printer at a shell prompt.

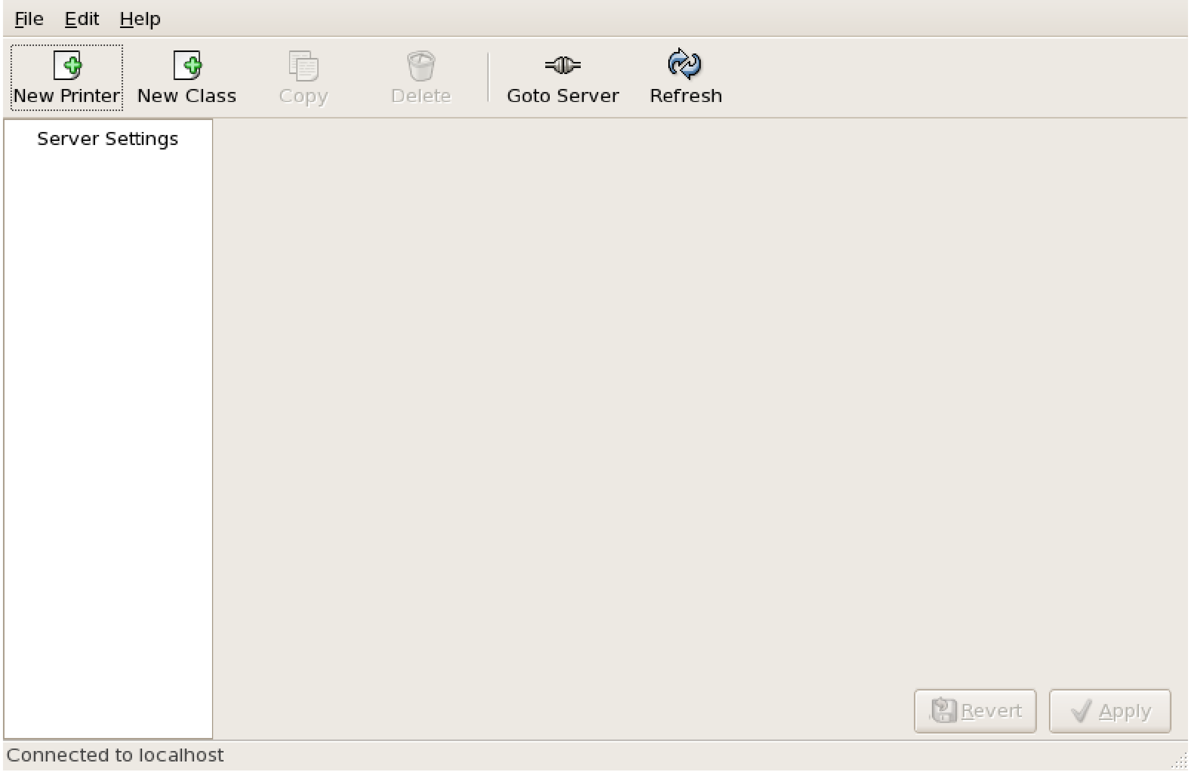

그림 36.1. Printer Configuration Tool

다음과 같은 유형의 인쇄 대기열을 설정 가능합니다:

- AppSocket/HP JetDirect — a printer connected directly to the network through HP JetDirect or Appsocket interface instead of a computer.
- Internet Printing Protocol (IPP) — a printer that can be accessed over a TCP/IP network via the Internet Printing Protocol (for example, a printer attached to another Red Hat Enterprise Linux system running CUPS on the network).
- LPD/LPR Host or Printer — a printer attached to a different UNIX system that can be accessed over a TCP/IP network (for example, a printer attached to another Red Hat Enterprise Linux system running LPD on the network).
- Networked Windows (SMB) — a printer attached to a different system which is sharing a printer over an SMB network (for example, a printer attached to a Microsoft Windows™ machine).
- Networked JetDirect — a printer connected directly to the network through HP JetDirect instead of a computer.

중요

만일 새로운 인쇄 대기열을 추가하시거나 현재의 것을 변경하실 때, 변경 사항을 적용하셔야 효력을 발생합니다.

Clicking the Apply button prompts the printer daemon to restart with the changes you have configured.

Clicking the Revert button discards unapplied changes.

36.1. Adding a Local Printer

To add a local printer, such as one attached through a parallel port or USB port on your computer, click the New Printer button in the main Printer Configuration Tool window to display the window in 그림 36.2. “Adding a Printer” .

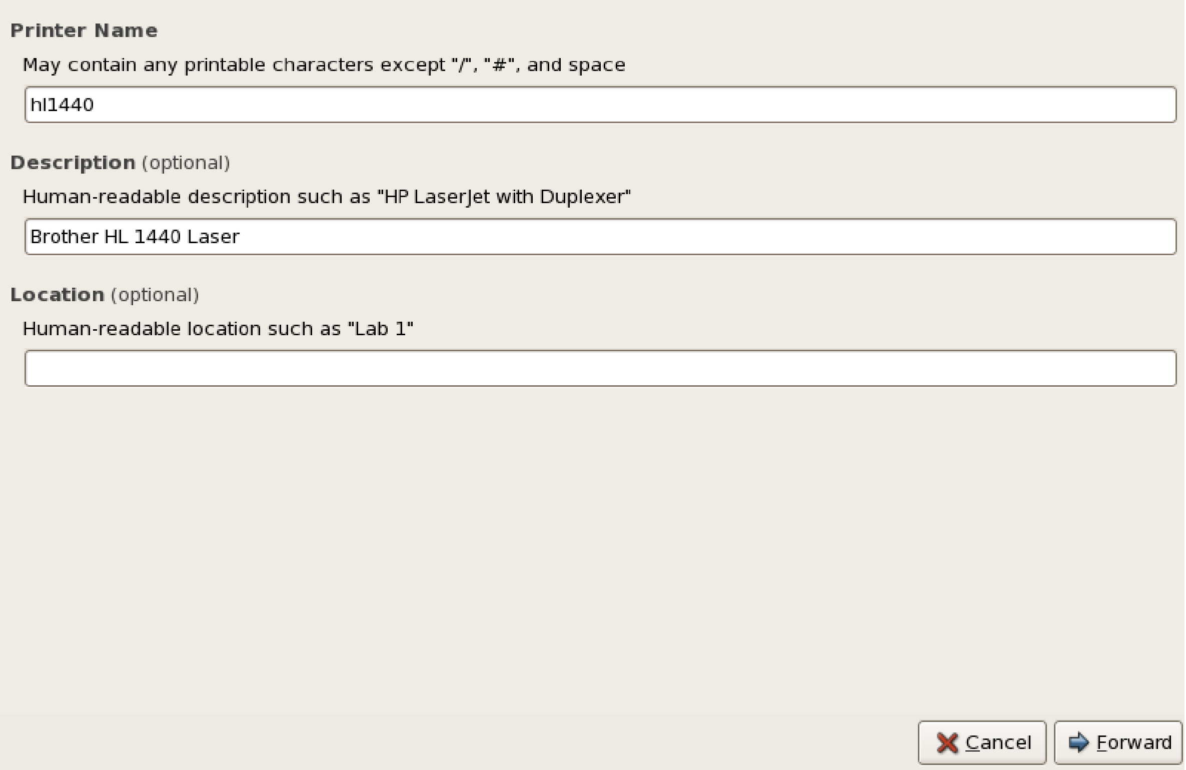

Printer Name

May contain any printable characters except "/", "#", and space

hl1440

Description (optional)

Human-readable description such as "HP LaserJet with Duplexer"

Brother HL 1440 Laser

Location (optional)

Human-readable location such as "Lab 1"

그림 36.2. Adding a Printer

Click Forward to proceed.

Enter a unique name for the printer in the Printer Name field. The printer name can contain letters, numbers, dashes (-), and underscores (_); it must not contain any spaces.

You can also use the Description and Location fields to further distinguish this printer from others that may be configured on your system. Both of these fields are optional, and may contain spaces.

Click Forward to open the New Printer dialogue (refer to 그림 36.3. "Adding a Local Printer"). If the printer has been automatically detected, the printer model appears in Select Connection. Select the printer model and click Forward to continue.

If the device does not automatically appear, select the device to which the printer is connected (such as LPT #1 or Serial Port #1) in Select Connection.

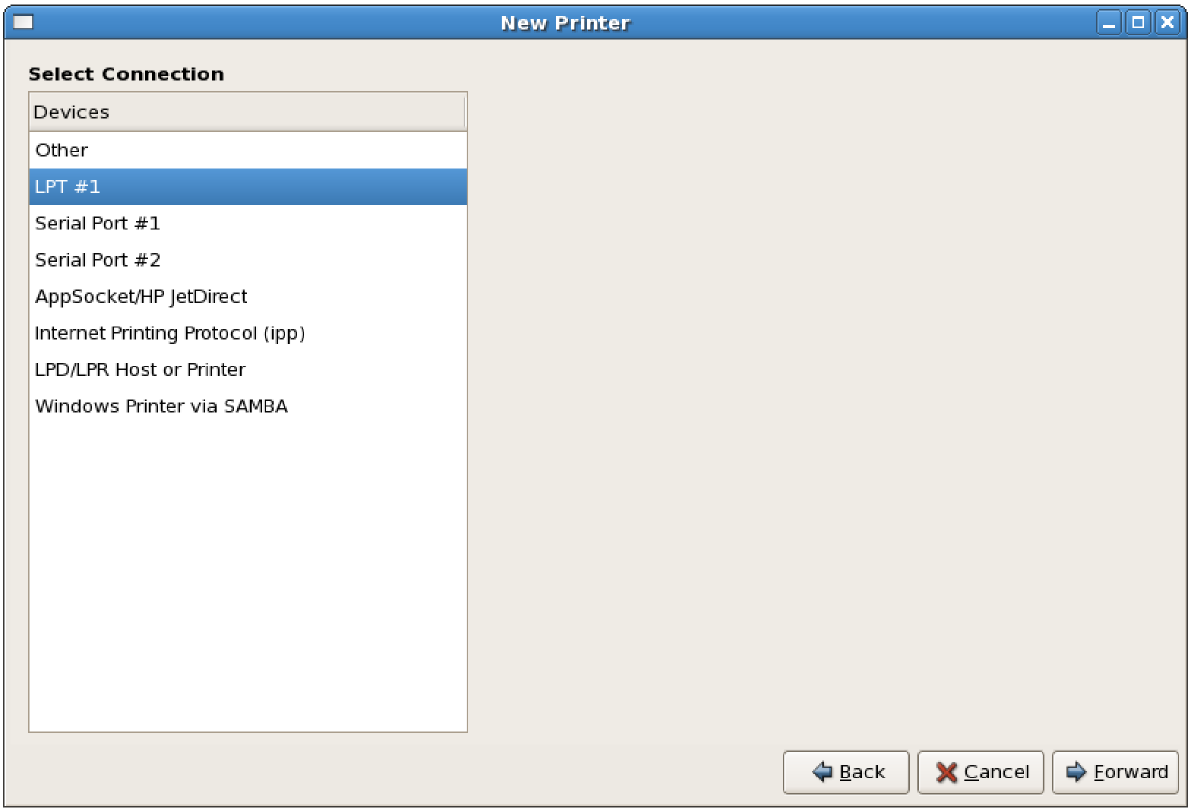

New Printer

Select Connection

Devices
- Other
- LPT #1
- Serial Port #1
- Serial Port #2
- AppSocket/HP JetDirect
- Internet Printing Protocol (ipp)
- LPD/LPR Host or Printer
- Windows Printer via SAMBA

그림 36.3. Adding a Local Printer

Next, select the printer type. Refer to 36.5절. "프린터 모델 선택 후 완료하기" for details.

36.2. Adding an IPP Printer

An IPP printer is a printer attached to a different system on the same TCP/IP network. The system this printer is attached to may either be running CUPS or simply configured to use IPP.

If a firewall is enabled on the printer server, then the firewall should be configured to allow send / receive connections on the incoming UDP port 631. If a firewall is enabled on the client (the system sending the print request) then the firewall should be configured to allow accept and create connections through port 631.

You can add a networked IPP printer by clicking the New Printer button in the main Printer Configuration Tool window to display the window in 그림 36.2. "Adding a Printer" . Enter the Printer Name (printer names cannot contain spaces and may contain letters, numbers, dashes (-), and underscores (_)), Description, and Location to distinguish this printer from others that you may configure on your system. Click Forward to proceed.

In the window shown in 그림 36.4. "Adding an IPP Printer" , enter the hostname of the IPP printer in the Hostname field as well as a unique name for the printer in the Printername field.

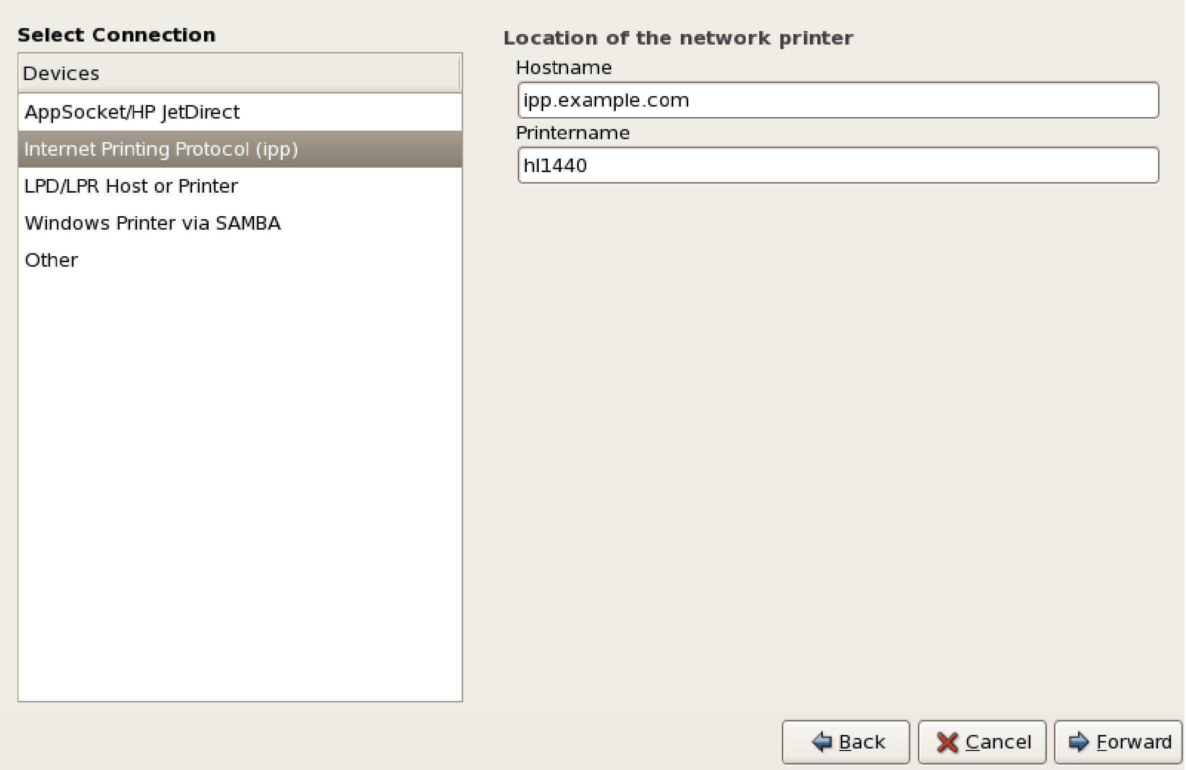

Location of the network printer

Hostname: ipp.example.com

Printername: hl1440

그림 36.4. Adding an IPP Printer

Click Forward to continue.

Next, select the printer type. Refer to 36.5절. "프린터 모델 선택 후 완료하기" for details.

36.3. Adding a Samba (SMB) Printer

You can add a Samba (SMB) based printer share by clicking the New Printer button in the main Printer Configuration Tool window to display the window in 그림 36.2. "Adding a Printer" . Enter a unique name for the printer in the Printer Name field. The printer name can contain letters, numbers, dashes (-), and underscores (_); it must not contain any spaces.

You can also use the Description and Location fields to further distinguish this printer from others that may be configured on your system. Both of these fields are optional, and may contain spaces.

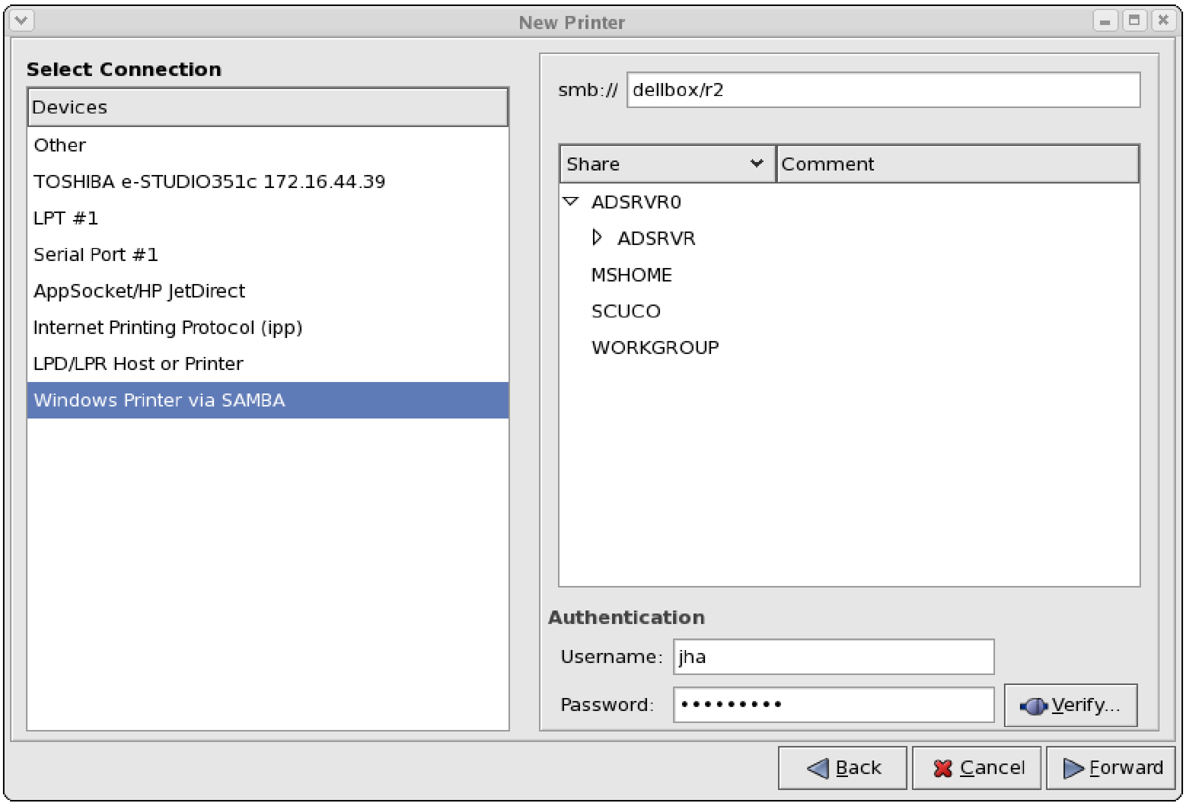

Select Connection: Windows Printer via SAMBA

smb:// dellbox/r2

Authentication

Username: jha

Password: •••••••••

그림 36.5. Adding a SMB Printer

As shown in 그림 36.5. “Adding a SMB Printer” , available SMB shares are automatically detected and listed in the Share column. Click the arrow (▷) beside a Workgroup to expand it. From the expanded list, select a printer.

If the printer you are looking for does not appear in the list, enter the SMB address in the smb:// field. Use the format computer name/printer share. In 그림 36.5. “Adding a SMB Printer” , the computer name is dellbox, while the printer share is r2.

In the Username field, enter the username to access the printer. This user must exist on the SMB system, and the user must have permission to access the printer. The default user name is typically guest for Windows servers, or nobody for Samba servers.

Enter the Password (if required) for the user specified in the Username field.

You can then test the connection by clicking Verify. Upon successful verification, a dialog box appears confirming printer share accessibility.

Next, select the printer type. Refer to 36.5절. “프린터 모델 선택 후 완료하기” for details.

경고

Samba printer usernames and passwords are stored in the printer server as unencrypted files readable by root and lpd. Thus, other users that have root access to the printer server can view the username and password you use to access the Samba printer.

As such, when you choose a username and password to access a Samba printer, it is advisable that you choose a password that is different from what you use to access your local Red Hat Enterprise Linux system.

If there are files shared on the Samba print server, it is recommended that they also use a password different from what is used by the print queue.

36.4. Adding a JetDirect Printer

To add a JetDirect or AppSocket connected printer share, click the New Printer button in the main Printer Configuration Tool window to display the window in 그림 36.2. "Adding a Printer" . Enter a unique name for the printer in the Printer Name field. The printer name can contain letters, numbers, dashes (-), and underscores (_); it must not contain any spaces.

You can also use the Description and Location fields to further distinguish this printer from others that may be configured on your system. Both of these fields are optional, and may contain spaces.

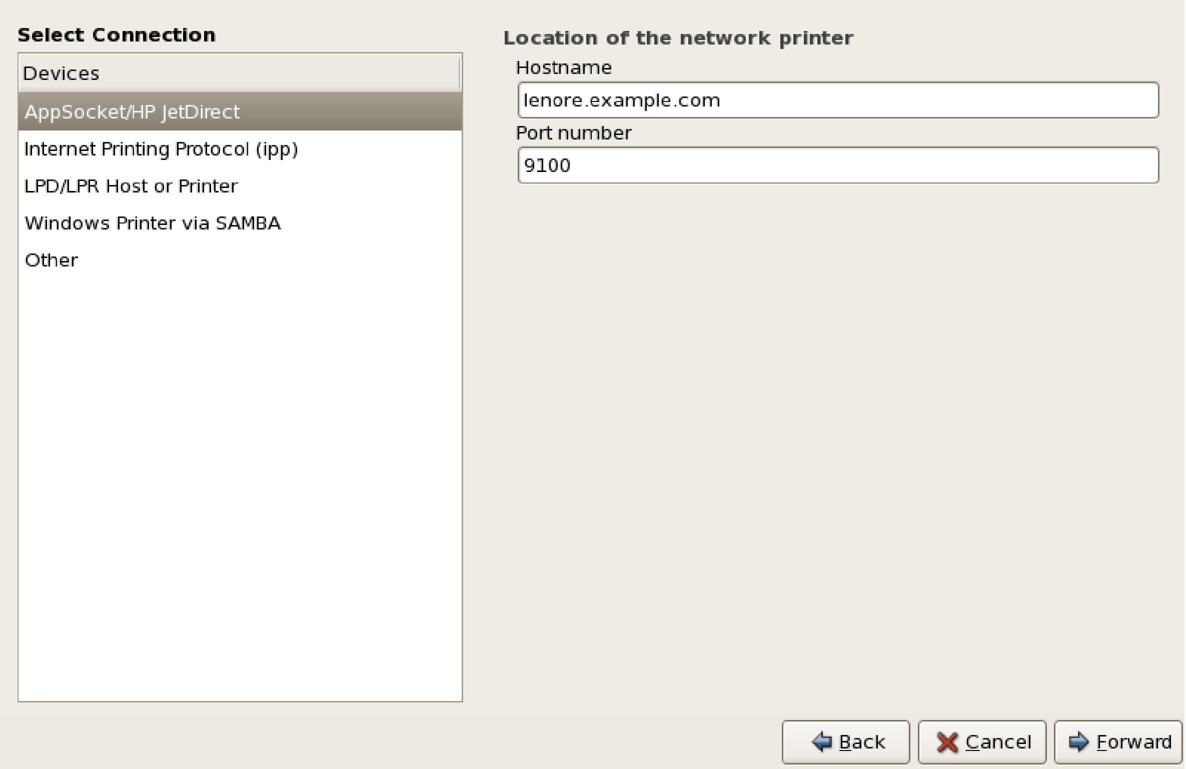

Hostname: lenore.example.com

Port number: 9100

그림 36.6. Adding a JetDirect Printer

Click Forward to continue.

다음과 같은 옵션에 대한 입력란이 나타납니다:

- Hostname — The hostname or IP address of the JetDirect printer.

- Port Number — The port on the JetDirect printer that is listening for print jobs. The default port is 9100.

Next, select the printer type. Refer to 36.5절. "프린터 모델 선택 후 완료하기" for details.

36.5. 프린터 모델 선택 후 완료하기

Once you have properly selected a printer queue type, you can choose either option:

- Select a Printer from database - If you select this option, choose the make of your printer from the list of Makes. If your printer make is not listed, choose Generic.

- Provide PPD file - A PostScript Printer Description (PPD) file may also be provided with your printer. This file is normally provided by the manufacturer. If you are provided with a PPD file, you can choose this option and use the browser bar below the option description to select the PPD file.

Refer to 그림 36.7. "Selecting a Printer Model" .

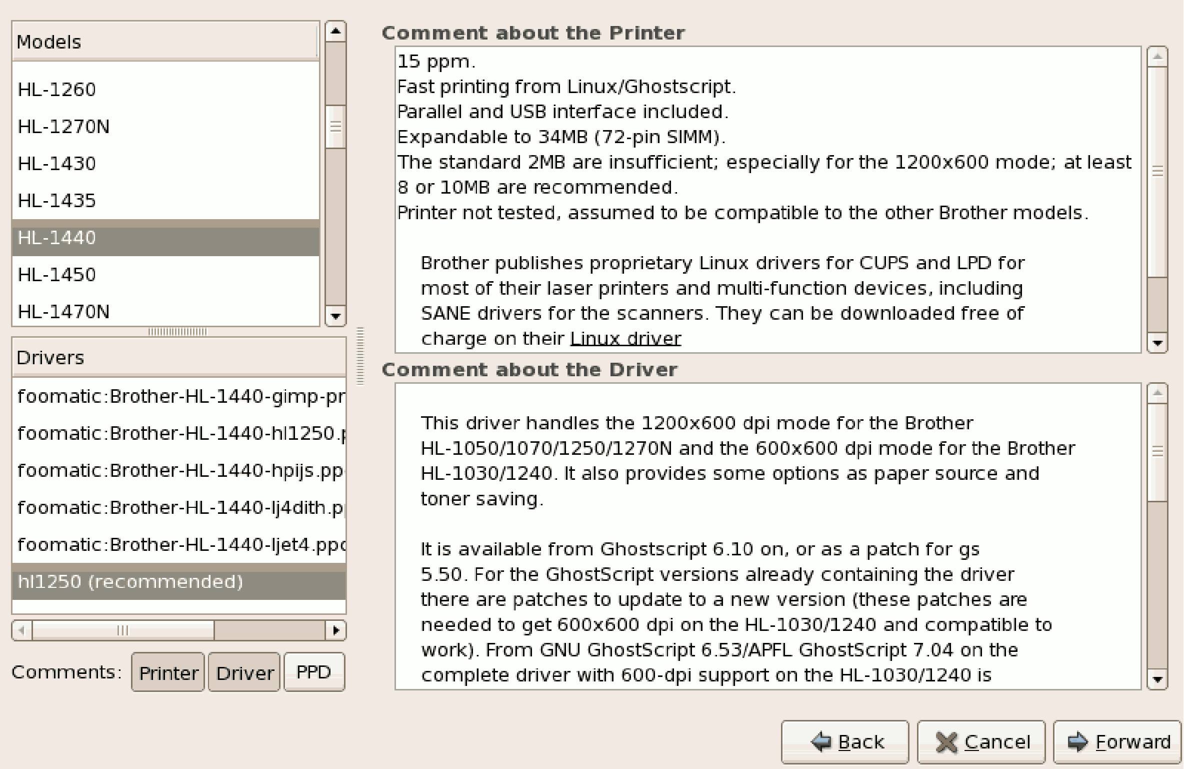

Models

HL-1260
HL-1270N
HL-1430
HL-1435
HL-1440
HL-1450
HL-1470N

Drivers

foomatic:Brother-HL-1440-gimp-pr
foomatic:Brother-HL-1440-hl1250.
foomatic:Brother-HL-1440-hpijs.pp
foomatic:Brother-HL-1440-lj4dith.p
foomatic:Brother-HL-1440-ljet4.ppc
hl1250 (recommended)

Comment about the Printer

15 ppm.
Fast printing from Linux/Ghostscript.
Parallel and USB interface included.
Expandable to 34MB (72-pin SIMM).
The standard 2MB are insufficient; especially for the 1200x600 mode; at least 8 or 10MB are recommended.
Printer not tested, assumed to be compatible to the other Brother models.

Brother publishes proprietary Linux drivers for CUPS and LPD for most of their laser printers and multi-function devices, including SANE drivers for the scanners. They can be downloaded free of charge on their Linux driver

Comment about the Driver

This driver handles the 1200x600 dpi mode for the Brother HL-1050/1070/1250/1270N and the 600x600 dpi mode for the Brother HL-1030/1240. It also provides some options as paper source and toner saving.

It is available from Ghostscript 6.10 on, or as a patch for gs 5.50. For the GhostScript versions already containing the driver there are patches to update to a new version (these patches are needed to get 600x600 dpi on the HL-1030/1240 and compatible to work). From GNU GhostScript 6.53/APFL GhostScript 7.04 on the complete driver with 600-dpi support on the HL-1030/1240 is

그림 36.7. Selecting a Printer Model

After choosing an option, click Forward to continue. 그림 36.7. "Selecting a Printer Model" appears. You now have to choose the corresponding model and driver for the printer.

The recommended printed driver is automatically selected based on the printer model you chose. The print driver processes the data that you want to print into a format the printer can understand. Since a local printer is attached directly to your computer, you need a printer driver to process the data that is sent to the printer.

If you have a PPD file for the device (usually provided by the manufacturer), you can select it by choosing Provide PPD file. You can then browse the filesystem for the PPD file by clicking Browse.

36.5.1. Confirming Printer Configuration

The last step is to confirm your printer configuration. Click Apply to add the print queue if the settings are correct. Click Back to modify the printer configuration.

After applying the changes, print a test page to ensure the configuration is correct. Refer to 36.6절. "테스트 페이지 인쇄하기" for details.

36.6. 테스트 페이지 인쇄하기

After you have configured your printer, you should print a test page to make sure the printer is functioning properly. To print a test page, select the printer that you want to try out from the printer list, then click Print Test Page from the printer's Settings tab.

프린트 드라이버나 드라이버 옵션을 변경하신다면, 다른 설정을 테스트하기 위해 테스트 페이지를 인쇄해보셔야 합니다.

36.7. 기존 프린터 수정하기

To delete an existing printer, select the printer and click the Delete button on the toolbar. The printer is removed from the printer list once you confirm deletion of the printer configuration.

To set the default printer, select the printer from the printer list and click the Make Default Printer button in the Settings tab.

36.7.1. The Settings Tab

To change printer driver configuration, click the corresponding name in the Printer list and click the Settings tab.

You can modify printer settings such as make and model, make a printer the default, print a test page, change the device location (URI), and more.

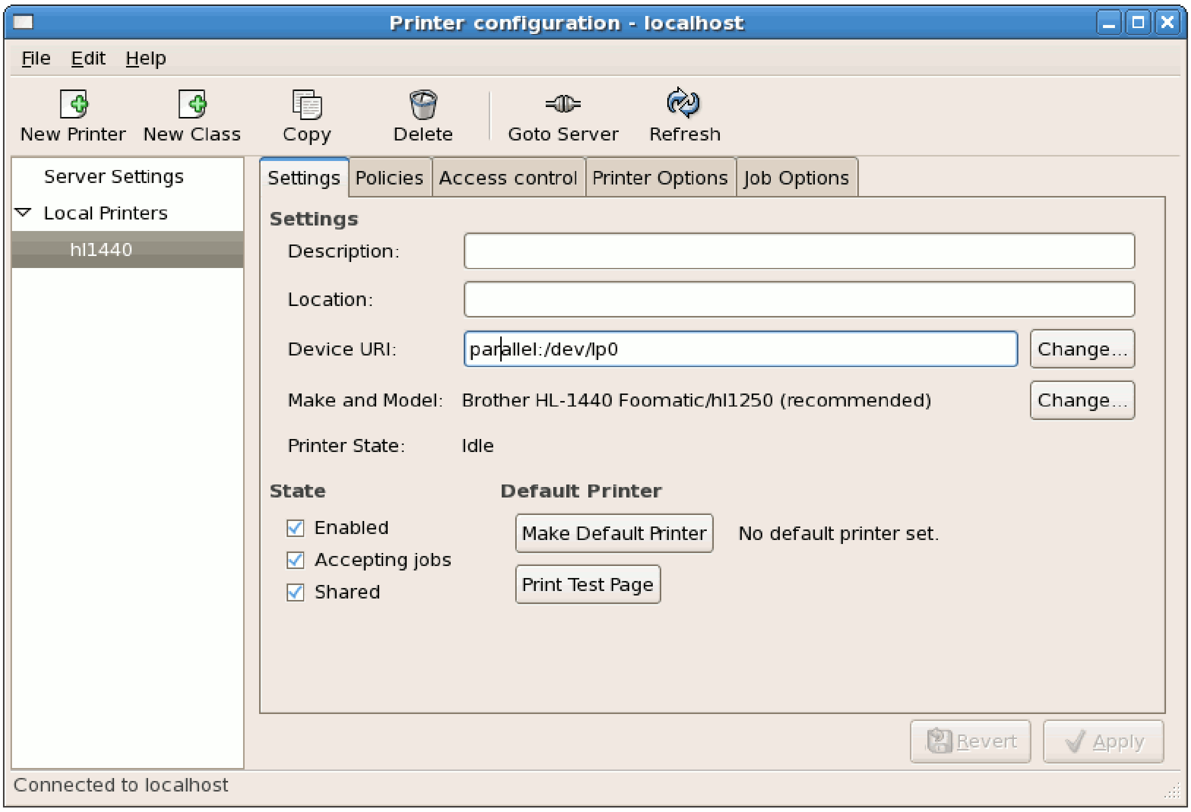

Settings

Description:

Location:

Device URI: parallel:/dev/lp0

Make and Model: Brother HL-1440 Foomatic/hl1250 (recommended)

Printer State: Idle

State

☑ Enabled

☑ Accepting jobs

☑ Shared

Default Printer

No default printer set.

Connected to localhost

그림 36.8. Settings Tab

36.7.2. The Policies Tab

To change settings in print output, click the Policies tab.

For example, to create a banner page (a page that describes aspects of the print job such as the originating printer, the username from the which the job originated, and the security status of the document being printed) click the Starting Banner or Ending Banner drop-menu and choose the option that best describes the nature of the print jobs (such as topsecret, classified, or confidential).

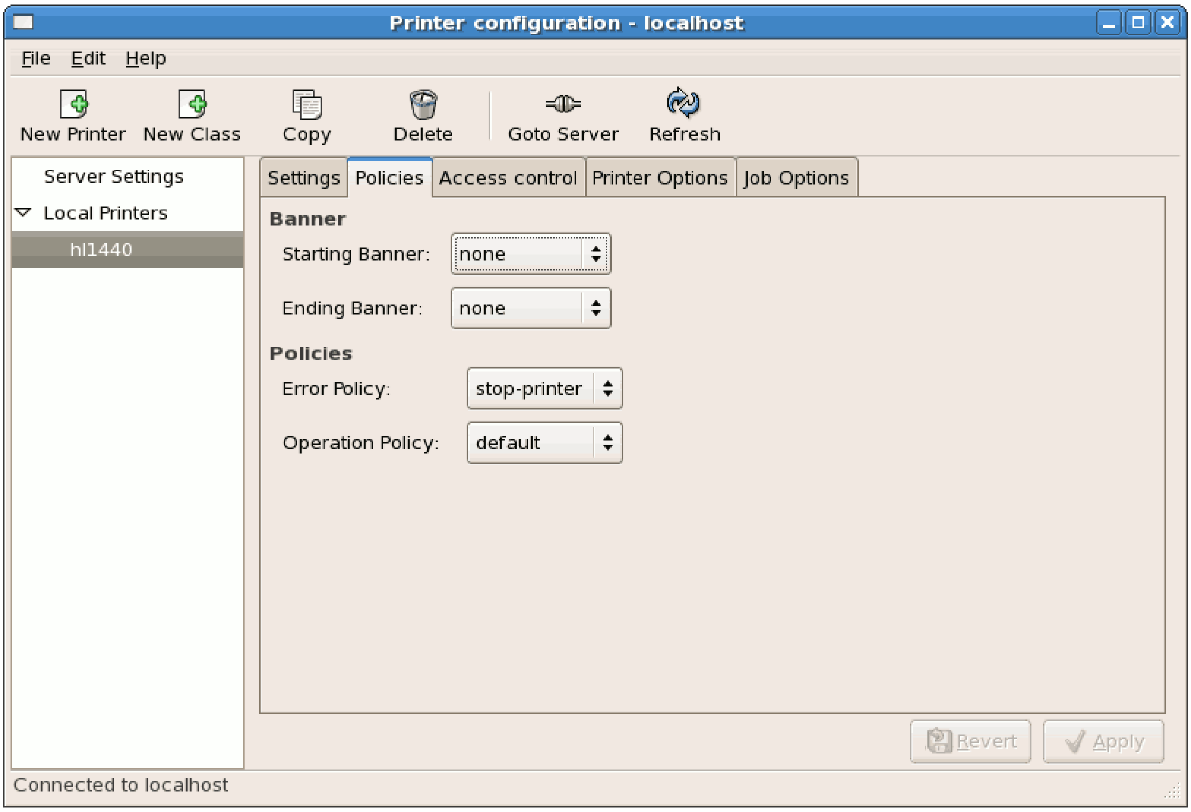

Banner

Starting Banner: none

Ending Banner: none

Policies

Error Policy: stop-printer

Operation Policy: default

그림 36.9. Policies Tab

You can also configure the Error Policy of the printer, by choosing an option from the drop-down menu. You can choose to abort the print job, retry, or stop it.

36.7.3. The Access Control Tab

You can change user-level access to the configured printer by clicking the Access Control tab.

Add users using the text box and click the Add button beside it. You can then choose to only allow use of the printer to that subset of users or deny use to those users.

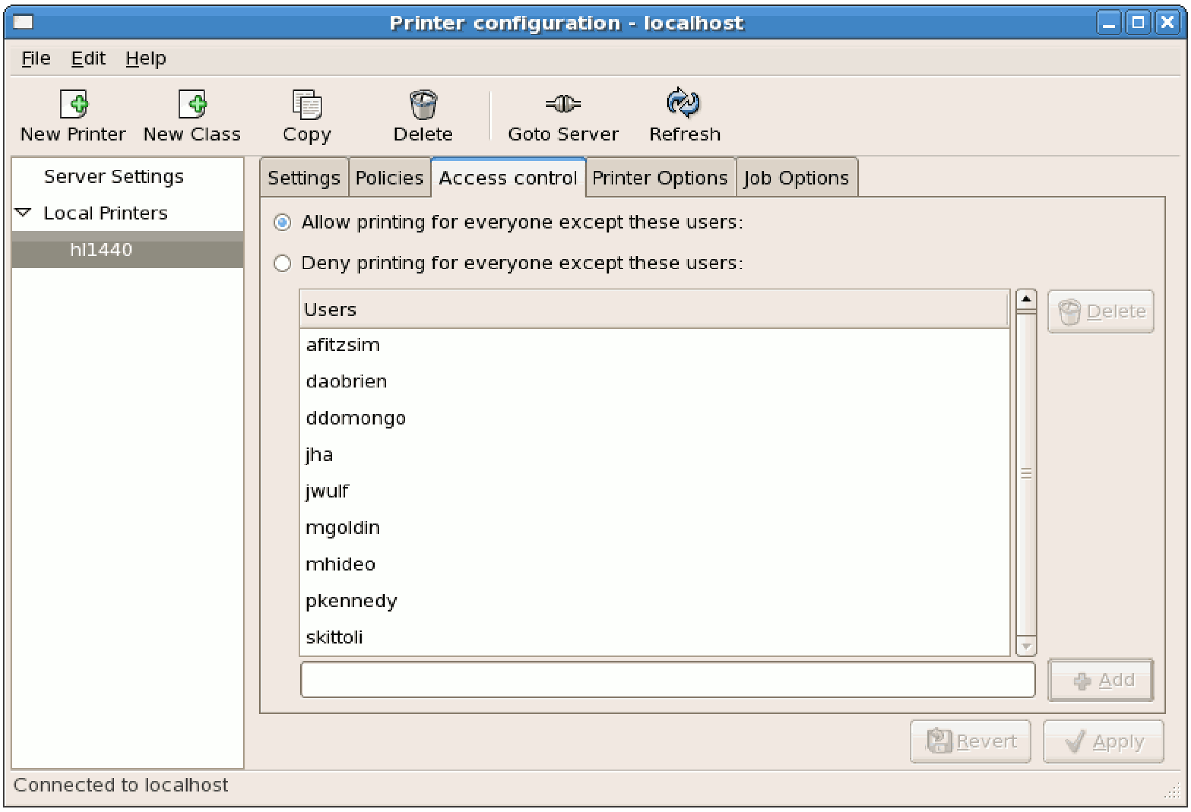

● Allow printing for everyone except these users:

○ Deny printing for everyone except these users:

Users
afitzsim
daobrien
ddomongo
jha
jwulf
mgoldin
mhideo
pkennedy
skittoli

그림 36.10. Access Control Tab

36.7.4. The Printer and Job OptionsTab

The Printer Options tab contains various configuration options for the printer media and output.

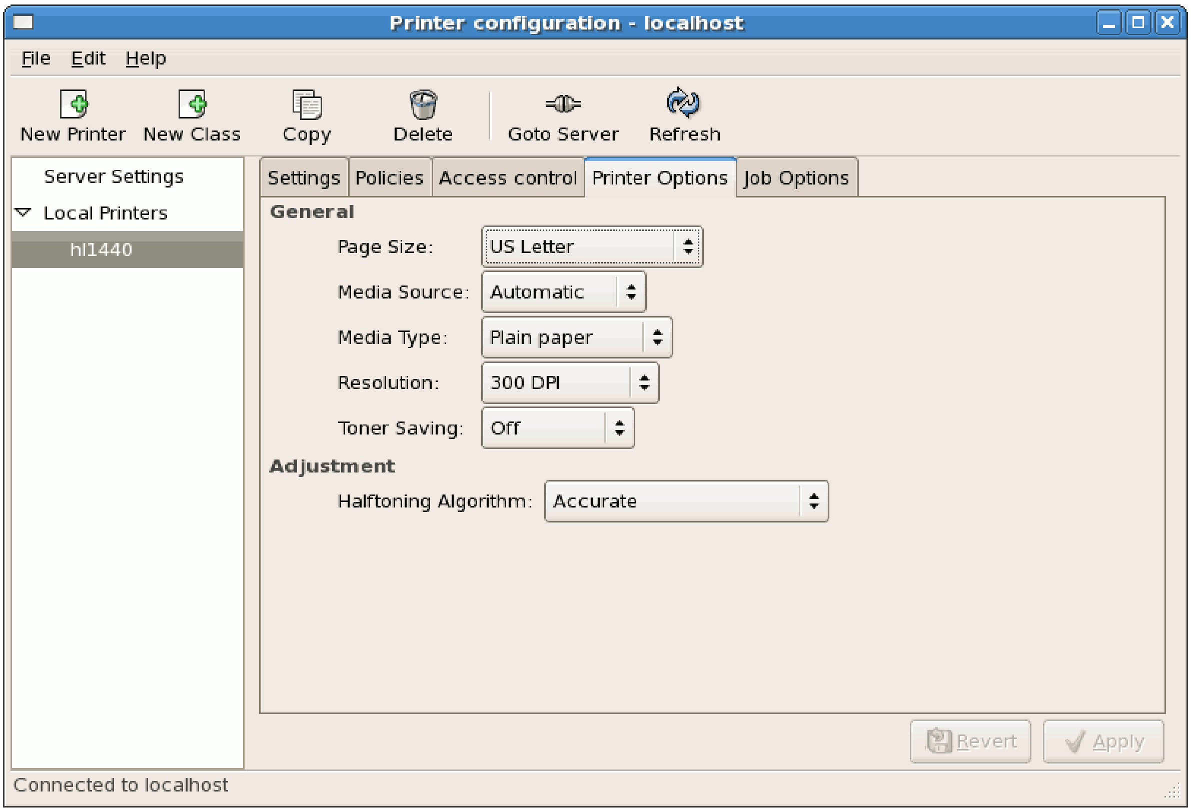

General

Page Size: US Letter

Media Source: Automatic

Media Type: Plain paper

Resolution: 300 DPI

Toner Saving: Off

Adjustment

Halftoning Algorithm: Accurate

그림 36.11. Printer Options Tab

- Page Size — Allows the paper size to be selected. The options include US Letter, US Legal, A3, and A4
- Media Source — set to Automatic by default. Change this option to use paper from a different tray.
- Media Type — Allows you to change paper type. Options include: Plain, thick, bond, and transparency.
- Resolution — Configure the quality and detail of the printout. Default is 300 dots per inch (dpi).
- Toner Saving — Choose whether the printer uses less toner to conserve resources.

You can also configure printer job options using the Job Options tab. Use the drop-menu and choose the job options you wish to use, such as Landscape modes (horizontal or vertical printout), copies, or scaling (increase or decrease the size of the printable area, which can be used to fit an oversize print area onto a smaller physical sheet of print medium).

36.8. 인쇄 작업 관리하기

When you send a print job to the printer daemon, such as printing a text file from Emacs or printing an image from The GIMP, the print job is added to the print spool queue. The print spool queue is a list of print jobs that have been sent to the printer and information about each print request, such as the status of the request, the job number, and more.

During the printing process, the Printer Status icon appears in the Notification Area on the panel. To check the status of a print job, double click the Printer Status, which displays a window similar to 그림 36.12. “GNOME Print Status” .

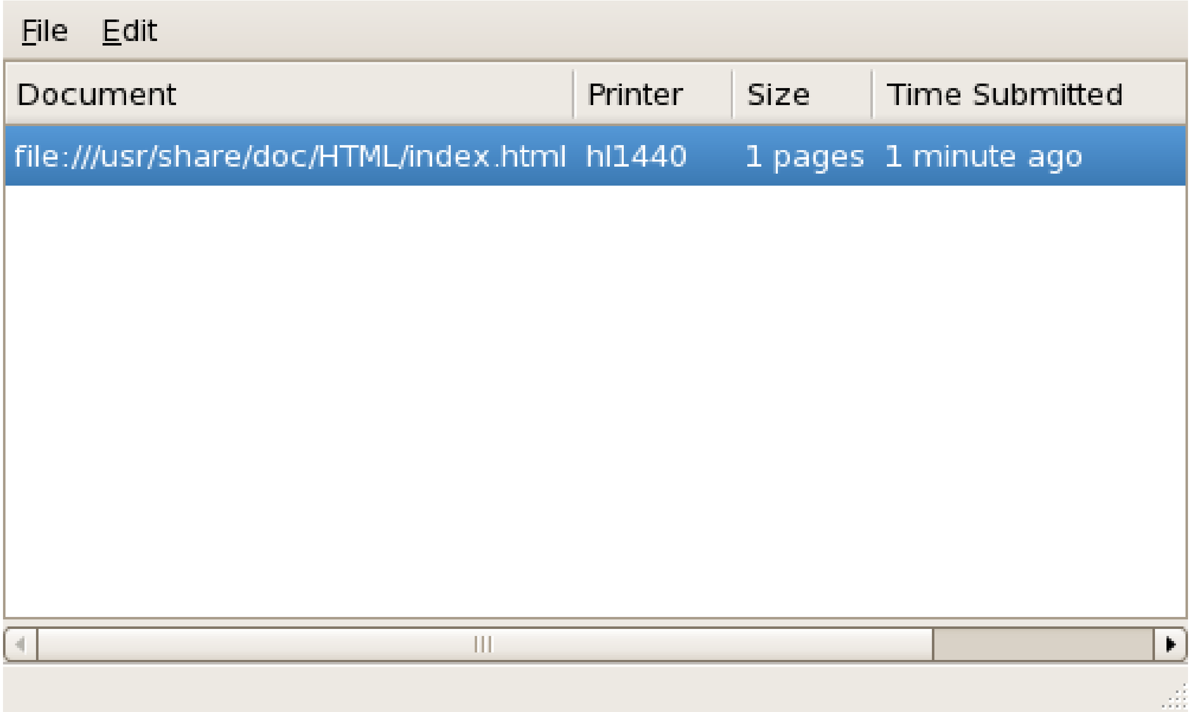

Document	Printer	Size	Time Submitted
file:///usr/share/doc/HTML/index.html	hl1440	1 pages	1 minute ago

그림 36.12. GNOME Print Status

To cancel a specific print job listed in the GNOME Print Status, select it from the list and select Edit > Cancel Documents from the pulldown menu.

To view the list of print jobs in the print spool from a shell prompt, type the command lpq. The last few lines look similar to the following:

예 36.1. Example of lpq output

```
Rank    Owner/ID                Class   Job  Files          Size  Time
active  user@localhost+902      A       902  sample.txt     2050  01:20:46
```

If you want to cancel a print job, find the job number of the request with the command lpq and then use the command lprm job number. For example, lprm 902 would cancel the print job in 예 36.1. “Example of lpq output” . You must have proper permissions to cancel a print job. You can not cancel print jobs that were started by other users unless you are logged in as root on the machine to which the printer is attached.

You can also print a file directly from a shell prompt. For example, the command lpr sample.txt prints the text file sample.txt. The print filter determines what type of file it is and converts it into a format the printer can understand.

36.9. 추가 자료

To learn more about printing on Red Hat Enterprise Linux, refer to the following resources.

36.9.1. 설치된 문서 자료

- map lpr — The manual page for the lpr command that allows you to print files from the command line.

- man lprm — The manual page for the command line utility to remove print jobs from the print queue.

- man mpage — The manual page for the command line utility to print multiple pages on one sheet of paper.

- man cupsd — The manual page for the CUPS printer daemon.

- man cupsd.conf — The manual page for the CUPS printer daemon configuration file.

- man classes.conf — The manual page for the class configuration file for CUPS.

36.9.2. 유용한 웹사이트

- http://www.linuxprinting.org — GNU/Linux Printing에는 리눅스에서 인쇄하기와 관련된 많은 정보가 포함되어 있습니다.

- http://www.cups.org/ — 문서 자료, FAQ와 CUPS 관련 뉴스 그룹.

Automated Tasks

In Linux, tasks can be configured to run automatically within a specified period of time, on a specified date, or when the system load average is below a specified number. Red Hat Enterprise Linux is pre-configured to run important system tasks to keep the system updated. For example, the slocate database used by the locate command is updated daily. A system administrator can use automated tasks to perform periodic backups, monitor the system, run custom scripts, and more.

Red Hat Enterprise Linux comes with several automated tasks utilities: cron, at, and batch.

37.1. Cron

Cron은 정해진 시간, 일, 월, 주마다 반복적인 작업을 실행하도록 스케줄하는데 사용되는 데몬입니다.

Cron assumes that the system is on continuously. If the system is not on when a task is scheduled, it is not executed. To schedule one-time tasks, refer to 37.2절. “At와 Batch” .

To use the cron service, the vixie-cron RPM package must be installed and the crond service must be running. To determine if the package is installed, use the rpm -q vixie-cron command. To determine if the service is running, use the command /sbin/service crond status.

37.1.1. Cron 작업 설정하기

The main configuration file for cron, /etc/crontab, contains the following lines:

```
SHELL=/bin/bash
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root HOME=/
# run-parts
01 * * * * root run-parts /etc/cron.hourly
02 4 * * * root run-parts /etc/cron.daily
22 4 * * 0 root run-parts /etc/cron.weekly
42 4 1 * * root run-parts /etc/cron.monthly
```

The first four lines are variables used to configure the environment in which the cron tasks are run. The SHELL variable tells the system which shell environment to use (in this example the bash shell), while the PATH variable defines the path used to execute commands. The output of the cron tasks are emailed to the username defined with the MAILTO variable. If the MAILTO variable is defined as an empty string (MAILTO=""), email is not sent. The HOME variable can be used to set the home directory to use when executing commands or scripts.

Each line in the /etc/crontab file represents a task and has the following format:

```
minute   hour   day   month   dayofweek   command
```

- minute — any integer from 0 to 59
- hour — any integer from 0 to 23
- day — any integer from 1 to 31 (must be a valid day if a month is specified)
- month — any integer from 1 to 12 (or the short name of the month such as jan or feb)

- dayofweek — any integer from 0 to 7, where 0 or 7 represents Sunday (or the short name of the week such as sun or mon)
- command — the command to execute (the command can either be a command such as ls /proc >> /tmp/proc or the command to execute a custom script)

앞에서 언급된 값에서 별표 (*)를 사용하시면 모든 유효값을 지정합니다. 예를 들어, 월 대신 별표를 지정한다면 다른 값의 범위 안에서 매달마다 명령을 실행하게 됩니다.

A hyphen (-) between integers specifies a range of integers. For example, 1-4 means the integers 1, 2, 3, and 4.

A list of values separated by commas (,) specifies a list. For example, 3, 4, 6, 8 indicates those four specific integers.

The forward slash (/) can be used to specify step values. The value of an integer can be skipped within a range by following the range with /<integer>. For example, 0-59/2 can be used to define every other minute in the minute field. Step values can also be used with an asterisk. For instance, the value */3 can be used in the month field to run the task every third month.

우물정자 표시 (#)로 시작하는 줄은 모두 주석 처리되어 실행되지 않습니다.

As shown in the /etc/crontab file, the run-parts script executes the scripts in the /etc/cron.hourly/, /etc/cron.daily/, /etc/cron.weekly/, and /etc/cron.monthly/ directories on an hourly, daily, weekly, or monthly basis respectively. The files in these directories should be shell scripts.

If a cron task is required to be executed on a schedule other than hourly, daily, weekly, or monthly, it can be added to the /etc/cron.d/ directory. All files in this directory use the same syntax as /etc/crontab. Refer to 예 37.1. “Crontab 예시” for examples.

예 37.1. Crontab 예시

```
# record the memory usage of the system every monday
# at 3:30AM in the file /tmp/meminfo
30 3 * * mon cat /proc/meminfo >> /tmp/meminfo
# run custom script the first day of every month at 4:10AM
10 4 1 * * /root/scripts/backup.sh
```

Users other than root can configure cron tasks by using the crontab utility. All user-defined crontabs are stored in the /var/spool/cron/ directory and are executed using the usernames of the users that created them. To create a crontab as a user, login as that user and type the command crontab -e to edit the user's crontab using the editor specified by the VISUAL or EDITOR environment variable. The file uses the same format as /etc/crontab. When the changes to the crontab are saved, the crontab is stored according to username and written to the file /var/spool/cron/username.

The cron daemon checks the /etc/crontab file, the /etc/cron.d/ directory, and the /var/spool/cron/ directory every minute for any changes. If any changes are found, they are loaded into memory. Thus, the daemon does not need to be restarted if a crontab file is changed.

37.1.2. Cron으로 접근을 통제하기

The /etc/cron.allow and /etc/cron.deny files are used to restrict access to cron. The format of both access control files is one username on each line. Whitespace is not permitted in either file. The cron daemon (crond) does not have to be restarted if the access control files are modified. The access control files are read each time a user tries to add or delete a cron task.

루트 사용자는 접근 통제 파일의 목록에 사용자명이 포함되지 않아도 항상 cron을 사용할 수 있습니다.

If the file cron.allow exists, only users listed in it are allowed to use cron, and the cron.deny file is ignored.

If cron.allow does not exist, users listed in cron.deny are not allowed to use cron.

37.1.3. 서비스 시작과 정지

To start the cron service, use the command /sbin/service crond start. To stop the service, use the command /sbin/service crond stop. It is recommended that you start the service at boot time. Refer to 17장. 서비스로의 접근 통제 for details on starting the cron service automatically at boot time.

37.2. At와 Batch

While cron is used to schedule recurring tasks, the at command is used to schedule a one-time task at a specific time and the batch command is used to schedule a one-time task to be executed when the systems load average drops below 0.8.

To use at or batch, the at RPM package must be installed, and the atd service must be running. To determine if the package is installed, use the rpm -q at command. To determine if the service is running, use the command /sbin/service atd status.

37.2.1. At 작업 설정하기

To schedule a one-time job at a specific time, type the command at time, where time is the time to execute the command.

time에는 다음 중 한가지 인자를 사용할 수 있습니다:

- HH:MM format — For example, 04:00 specifies 4:00 a.m. If the time is already past, it is executed at the specified time the next day.
- midnight — Specifies 12:00 a.m.
- noon — Specifies 12:00 p.m.
- teatime — Specifies 4:00 p.m.
- month-name day year 형식 — 예로 들면, January 15 2002는 2002년 1월의 15번째 날을 의미합니다. 년도수는 옵션입니다.
- MMDDYY, MM/DD/YY, 또는 MM.DD.YY 형식 — 예로 들면, 011502는 2002년 1월의 15번째 날을 의미합니다.
- now + time — 시간은 분, 시, 일 또는 주 단위입니다. 예를 들어, now + 5 days 라고 지정하시면 이 명령은 5일 후 같은 시간에 실행됩니다.

The time must be specified first, followed by the optional date. For more information about the time format, read the /usr/share/doc/at-<version>/timespec text file.

After typing the at command with the time argument, the at> prompt is displayed. Type the command to execute, press Enter, and type Ctrl+D . Multiple commands can be specified by typing each command followed by the Enter key. After typing all the commands, press Enter to go to a blank line and type Ctrl+D . Alternatively, a shell script can be entered at the prompt, pressing Enter after each line in the script, and typing Ctrl+D on a blank line to exit. If a script is entered, the shell used is the shell set in the user's SHELL environment, the user's login shell, or /bin/sh (whichever is found first).

명령어나 스크립트로 정보를 표준 출력하면, 출력 결과는 사용자에게 이메일로 전달됩니다.

Use the command atq to view pending jobs. Refer to 37.2.3절. “이후 실행할 작업 보기” for more information.

Usage of the at command can be restricted. For more information, refer to 37.2.5절. “At와 Batch로의 접근 통제하기” for details.

37.2.2. Batch 작업 설정하기

To execute a one-time task when the load average is below 0.8, use the batch command.

After typing the batch command, the at> prompt is displayed. Type the command to execute, press Enter, and type Ctrl+D . Multiple commands can be specified by typing each command followed by the Enter key. After typing all the commands, press Enter to go to a blank line and type Ctrl+D . Alternatively, a shell script can be entered at the prompt, pressing Enter after each line in the script, and typing Ctrl+D on a blank line to exit. If a script is entered, the shell used is the shell set in the user's SHELL environment, the user's login shell, or /bin/sh (whichever is found first). As soon as the load average is below 0.8, the set of commands or script is executed.

명령어나 스크립트로 정보를 표준 출력하면, 출력 결과는 사용자에게 이메일로 전달됩니다.

Use the command atq to view pending jobs. Refer to 37.2.3절. “이후 실행할 작업 보기” for more information.

Usage of the batch command can be restricted. For more information, refer to 37.2.5절. “At와 Batch로의 접근 통제하기” for details.

37.2.3. 이후 실행할 작업 보기

To view pending at and batch jobs, use the atq command. The atq command displays a list of pending jobs, with each job on a line. Each line follows the job number, date, hour, job class, and username format. Users can only view their own jobs. If the root user executes the atq command, all jobs for all users are displayed.

37.2.4. 추가 명령행 옵션

Additional command line options for at and batch include:

표 37.1. at and batch Command Line Options

옵션	설명
-f	명령어나 쉘 스크립트를 프롬프트에서 지정하지 않고 대신 파일에서 읽어옴.
-m	작업이 완료되면 사용자에게 이메일을 보냄.
-v	Display the time that the job is executed.

37.2.5. At와 Batch로의 접근 통제하기

The /etc/at.allow and /etc/at.deny files can be used to restrict access to the at and batch commands. The format of both access control files is one username on each line. Whitespace is not permitted in either file. The at daemon (atd) does not have to be restarted if the access control files are modified. The access control files are read each time a user tries to execute the at or batch commands.

The root user can always execute at and batch commands, regardless of the access control files.

If the file at.allow exists, only users listed in it are allowed to use at or batch, and the at.deny file is ignored.

If at.allow does not exist, users listed in at.deny are not allowed to use at or batch.

37.2.6. 서비스 시작과 정지

To start the at service, use the command /sbin/service atd start. To stop the service, use the command /sbin/service atd stop. It is recommended that you start the service at boot time. Refer to 17장. 서비스로의 접근 통제 for details on starting the cron service automatically at boot time.

37.3. 추가 자료

자동화 작업을 설정하는 방법에 대한 보다 많은 정보를 원하신다면, 다음의 자료를 참조하시기 바랍니다.

37.3.1. 설치된 문서 자료

- cron man page — overview of cron.
- crontab man pages in sections 1 and 5 — The man page in section 1 contains an overview of the crontab file. The man page in section 5 contains the format for the file and some example entries.
- /usr/share/doc/at-<version>/timespec contains more detailed information about the times that can be specified for cron jobs.
- at man page — description of at and batch and their command line options.

로그 파일

로그 파일 (Log file)이란 커널, 서비스 및 실행 중인 응용 프로그램과 같은 시스템 관련 메시지를 포함하는 파일을 의미합니다. 정보의 종류에 따라서 다른 로그 파일이 존재합니다. 예로 들면, 기본 시스템 로그 파일 및 보안 메시지 용 로그 파일, 크론(cron) 작업 용 로그 파일과 같은 여러 가지 로그 파일이 있습니다.

로그 파일은 시스템 상의 문제를 해결하려고 하실 때 (예, 커널 드라이버를 로드하거나 시스템으로 허가되지 않은 로그인 시도에 대한 로그를 찾으실 때) 유용하게 사용됩니다. 이 장에서는 로그 파일이 저장된 위치와 로그 파일을 보는 방법 및 로그 파일에서 중요한 정보를 찾는 방법에 대하여 다루어 보겠습니다.

Some log files are controlled by a daemon called syslogd. A list of log messages maintained by syslogd can be found in the /etc/syslog.conf configuration file.

38.1. 로그 파일 찾기

Most log files are located in the /var/log/ directory. Some applications such as httpd and samba have a directory within /var/log/ for their log files.

You may notice multiple files in the log file directory with numbers after them. These are created when the log files are rotated. Log files are rotated so their file sizes do not become too large. The logrotate package contains a cron task that automatically rotates log files according to the /etc/logrotate.conf configuration file and the configuration files in the /etc/logrotate.d/ directory. By default, it is configured to rotate every week and keep four weeks worth of previous log files.

38.2. 로그 파일 보기

Most log files are in plain text format. You can view them with any text editor such as Vi or Emacs. Some log files are readable by all users on the system; however, root privileges are required to read most log files.

To view system log files in an interactive, real-time application, use the System Log Viewer. To start the application, go to Applications (the main menu on the panel) > System > System Logs, or type the command gnome-system-log at a shell prompt.

The application only displays log files that exist; thus, the list might differ from the one shown in 그림 38.1. “System Log Viewer” .

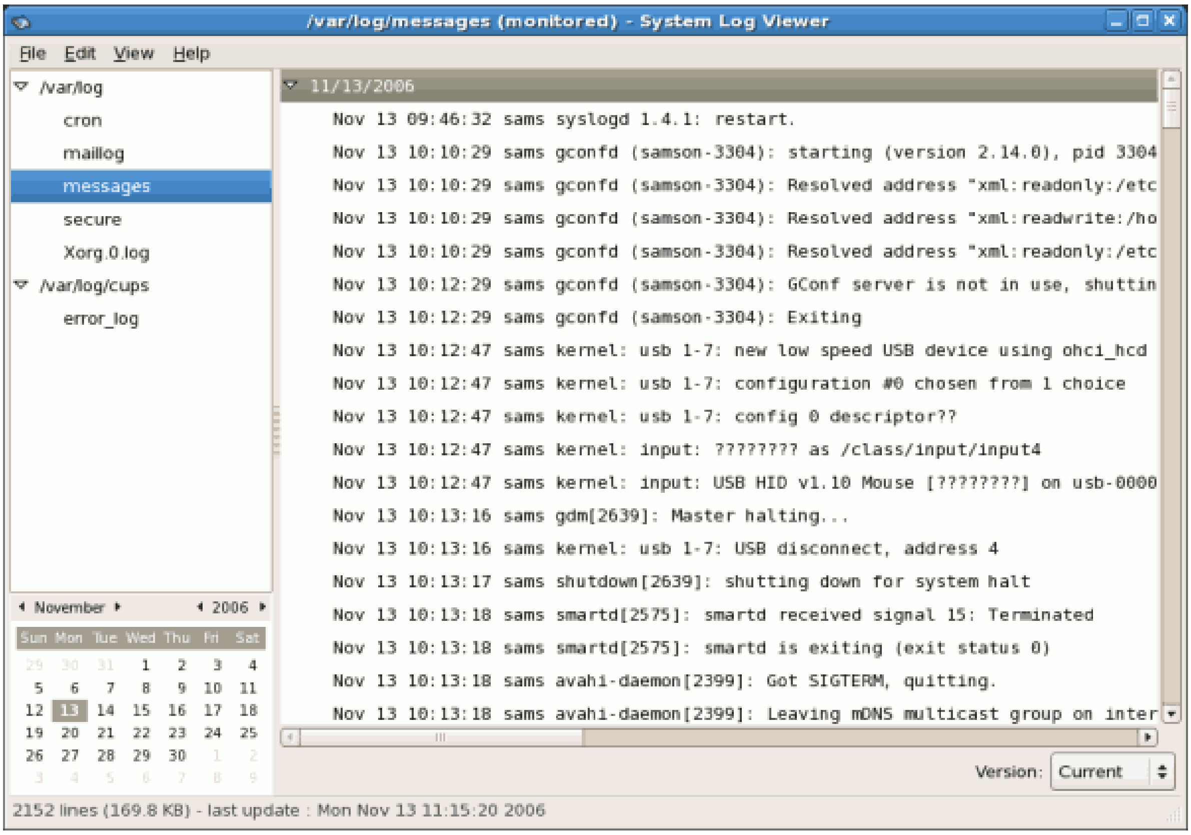

```
Nov 13 09:46:32 sams syslogd 1.4.1: restart.
Nov 13 10:10:29 sams gconfd (samson-3304): starting (version 2.14.0), pid 3304
Nov 13 10:10:29 sams gconfd (samson-3304): Resolved address "xml:readonly:/etc
Nov 13 10:10:29 sams gconfd (samson-3304): Resolved address "xml:readwrite:/ho
Nov 13 10:10:29 sams gconfd (samson-3304): Resolved address "xml:readonly:/etc
Nov 13 10:12:29 sams gconfd (samson-3304): GConf server is not in use, shuttin
Nov 13 10:12:29 sams gconfd (samson-3304): Exiting
Nov 13 10:12:47 sams kernel: usb 1-7: new low speed USB device using ohci_hcd
Nov 13 10:12:47 sams kernel: usb 1-7: configuration #0 chosen from 1 choice
Nov 13 10:12:47 sams kernel: usb 1-7: config 0 descriptor??
Nov 13 10:12:47 sams kernel: input: ???????? as /class/input/input4
Nov 13 10:12:47 sams kernel: input: USB HID v1.10 Mouse [????????] on usb-0000
Nov 13 10:13:16 sams gdm[2639]: Master halting...
Nov 13 10:13:16 sams kernel: usb 1-7: USB disconnect, address 4
Nov 13 10:13:17 sams shutdown[2639]: shutting down for system halt
Nov 13 10:13:18 sams smartd[2575]: smartd received signal 15: Terminated
Nov 13 10:13:18 sams smartd[2575]: smartd is exiting (exit status 0)
Nov 13 10:13:18 sams avahi-daemon[2399]: Got SIGTERM, quitting.
Nov 13 10:13:18 sams avahi-daemon[2399]: Leaving mDNS multicast group on inter
```

2152 lines (169.8 KB) - last update : Mon Nov 13 11:15:20 2006

그림 38.1. System Log Viewer

To filter the contents of the selected log file, click on View from the menu and select Filter as illustrated below.

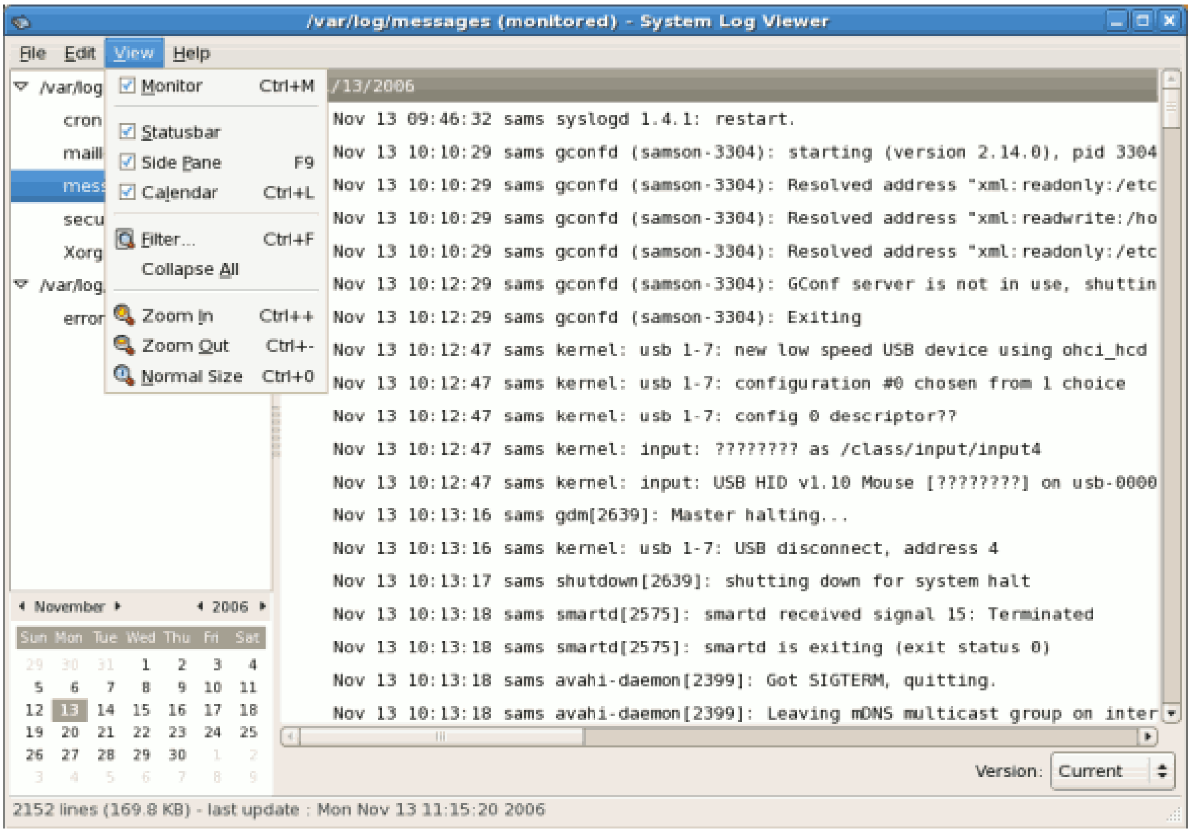

```
/13/2006
Nov 13 09:46:32 sams syslogd 1.4.1: restart.
Nov 13 10:10:29 sams gconfd (samson-3304): starting (version 2.14.0), pid 3304
Nov 13 10:10:29 sams gconfd (samson-3304): Resolved address "xml:readonly:/etc
Nov 13 10:10:29 sams gconfd (samson-3304): Resolved address "xml:readwrite:/ho
Nov 13 10:10:29 sams gconfd (samson-3304): Resolved address "xml:readonly:/etc
Nov 13 10:12:29 sams gconfd (samson-3304): GConf server is not in use, shuttin
Nov 13 10:12:29 sams gconfd (samson-3304): Exiting
Nov 13 10:12:47 sams kernel: usb 1-7: new low speed USB device using ohci_hcd
Nov 13 10:12:47 sams kernel: usb 1-7: configuration #0 chosen from 1 choice
Nov 13 10:12:47 sams kernel: usb 1-7: config 0 descriptor??
Nov 13 10:12:47 sams kernel: input: ???????? as /class/input/input4
Nov 13 10:12:47 sams kernel: input: USB HID v1.10 Mouse [????????] on usb-0000
Nov 13 10:13:16 sams gdm[2639]: Master halting...
Nov 13 10:13:16 sams kernel: usb 1-7: USB disconnect, address 4
Nov 13 10:13:17 sams shutdown[2639]: shutting down for system halt
Nov 13 10:13:18 sams smartd[2575]: smartd received signal 15: Terminated
Nov 13 10:13:18 sams smartd[2575]: smartd is exiting (exit status 0)
Nov 13 10:13:18 sams avahi-daemon[2399]: Got SIGTERM, quitting.
Nov 13 10:13:18 sams avahi-daemon[2399]: Leaving mDNS multicast group on inter
```

2152 lines (169.8 KB) - last update : Mon Nov 13 11:15:20 2006

그림 38.2. System Log Viewer - View Menu

Selecting the Filter menu item will display the Filter text field where you can type the keywords you wish to use for your filter. To clear your filter click on the Clear button.The figure below illustrates a sample filter.

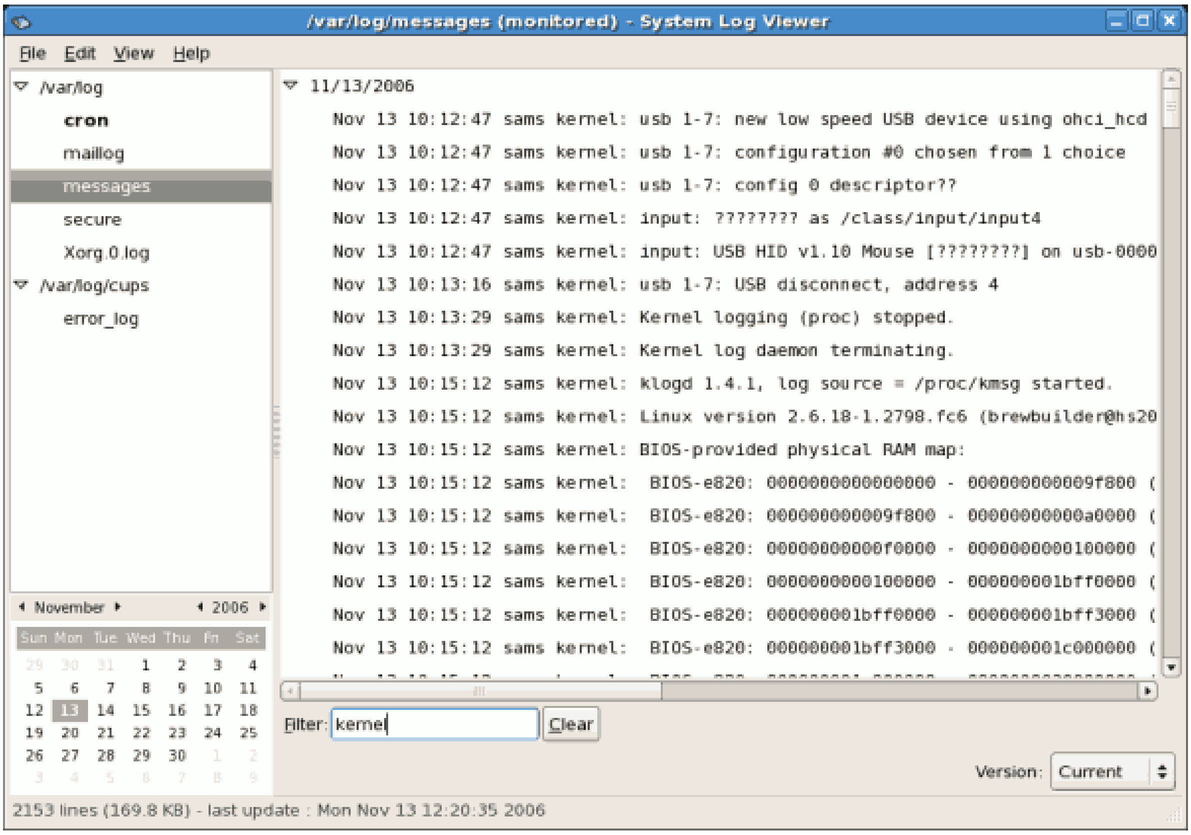

```
11/13/2006
    Nov 13 10:12:47 sams kernel: usb 1-7: new low speed USB device using ohci_hcd
    Nov 13 10:12:47 sams kernel: usb 1-7: configuration #0 chosen from 1 choice
    Nov 13 10:12:47 sams kernel: usb 1-7: config 0 descriptor??
    Nov 13 10:12:47 sams kernel: input: ???????? as /class/input/input4
    Nov 13 10:12:47 sams kernel: input: USB HID v1.10 Mouse [????????] on usb-0000
    Nov 13 10:13:16 sams kernel: usb 1-7: USB disconnect, address 4
    Nov 13 10:13:29 sams kernel: Kernel logging (proc) stopped.
    Nov 13 10:13:29 sams kernel: Kernel log daemon terminating.
    Nov 13 10:15:12 sams kernel: klogd 1.4.1, log source = /proc/kmsg started.
    Nov 13 10:15:12 sams kernel: Linux version 2.6.18-1.2798.fc6 (brewbuilder@hs20
    Nov 13 10:15:12 sams kernel: BIOS-provided physical RAM map:
    Nov 13 10:15:12 sams kernel:  BIOS-e820: 0000000000000000 - 000000000009f800 (
    Nov 13 10:15:12 sams kernel:  BIOS-e820: 000000000009f800 - 00000000000a0000 (
    Nov 13 10:15:12 sams kernel:  BIOS-e820: 00000000000f0000 - 0000000000100000 (
    Nov 13 10:15:12 sams kernel:  BIOS-e820: 0000000000100000 - 000000001bff0000 (
    Nov 13 10:15:12 sams kernel:  BIOS-e820: 000000001bff0000 - 000000001bff3000 (
    Nov 13 10:15:12 sams kernel:  BIOS-e820: 000000001bff3000 - 000000001c000000 (
```

Filter: kernel

Version: Current

2153 lines (169.8 KB) - last update : Mon Nov 13 12:20:35 2006

그림 38.3. System Log Viewer - Filter

38.3. Adding a Log File

To add a log file you wish to view in the list, select File > Open. This will display the Open Log window where you can select the directory and filename of the log file you wish to view.The figure below illustrates the Open Log window.

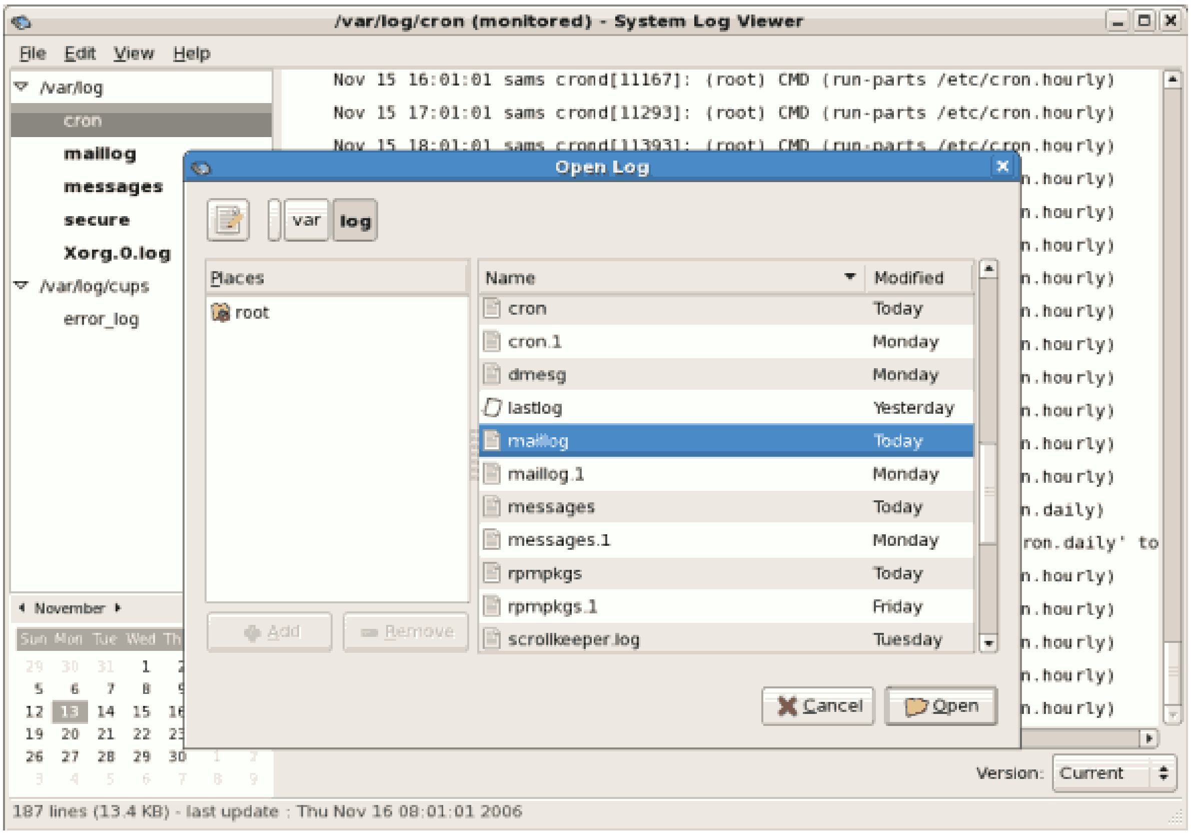

```
Nov 15 16:01:01 sams crond[11167]: (root) CMD (run-parts /etc/cron.hourly)
Nov 15 17:01:01 sams crond[11293]: (root) CMD (run-parts /etc/cron.hourly)
Nov 15 18:01:01 sams crond[11393]: (root) CMD (run-parts /etc/cron.hourly)
```

Name	Modified
cron	Today
cron.1	Monday
dmesg	Monday
lastlog	Yesterday
maillog	Today
maillog.1	Monday
messages	Today
messages.1	Monday
rpmpkgs	Today
rpmpkgs.1	Friday
scrollkeeper.log	Tuesday

187 lines (13.4 KB) - last update : Thu Nov 16 08:01:01 2006

그림 38.4. Adding a Log File

Click on the Open button to open the file. The file is immediately added to the viewing list where you can select it and view the contents.

Please also note that the System Log Viewer also allows you to open zipped logs whose filenames end in ".gz".

38.4. Monitoring Log Files

System Log Viewer monitors all opened logs by default. If a new line is added to a monitored log file, the log name appears in bold in the log list. If the log file is selected or displayed, the new lines appear in bold at the bottom of the log file and after five seconds are displayed in normal format. This is illustrated in the figures below. The figure below illustrates a new alert in the messages log file. The log file is listed in bold text.

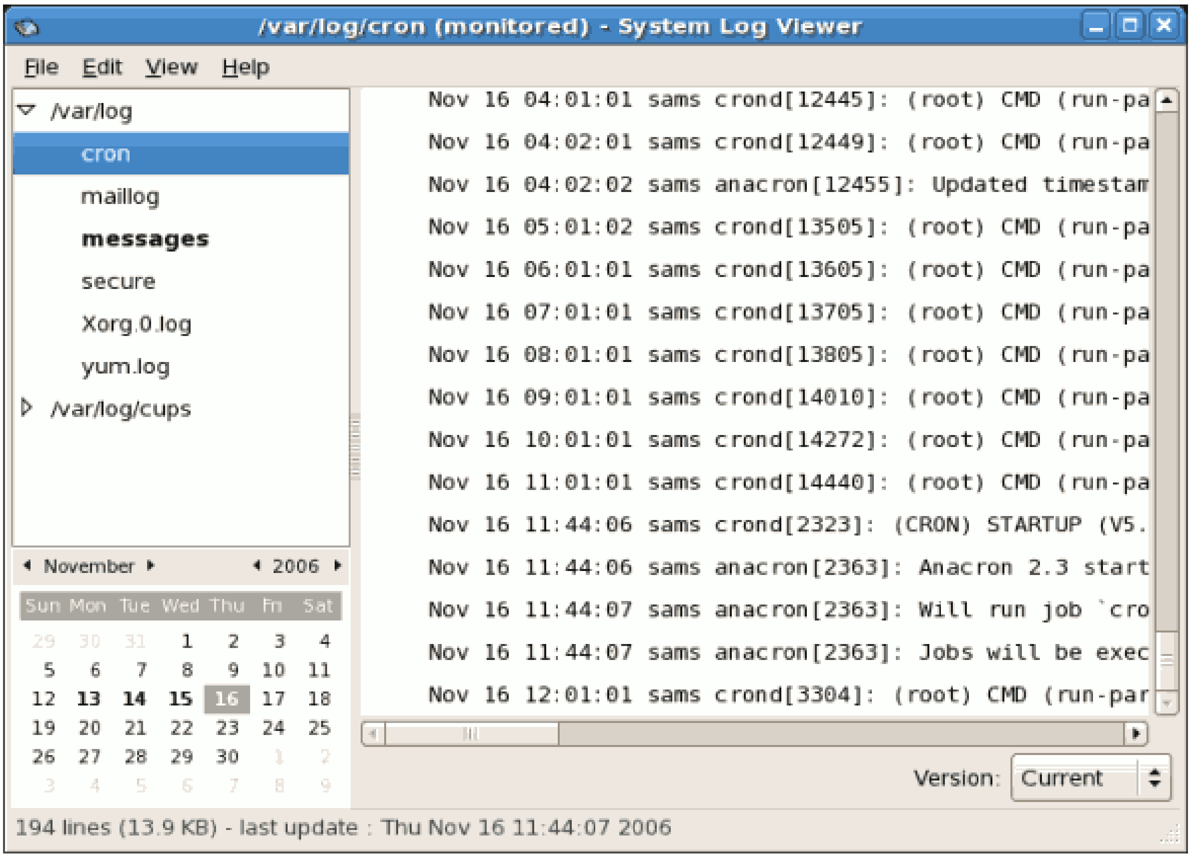

```
Nov 16 04:01:01 sams crond[12445]: (root) CMD (run-pa
Nov 16 04:02:01 sams crond[12449]: (root) CMD (run-pa
Nov 16 04:02:02 sams anacron[12455]: Updated timestam
Nov 16 05:01:02 sams crond[13505]: (root) CMD (run-pa
Nov 16 06:01:01 sams crond[13605]: (root) CMD (run-pa
Nov 16 07:01:01 sams crond[13705]: (root) CMD (run-pa
Nov 16 08:01:01 sams crond[13805]: (root) CMD (run-pa
Nov 16 09:01:01 sams crond[14010]: (root) CMD (run-pa
Nov 16 10:01:01 sams crond[14272]: (root) CMD (run-pa
Nov 16 11:01:01 sams crond[14440]: (root) CMD (run-pa
Nov 16 11:44:06 sams crond[2323]: (CRON) STARTUP (V5.
Nov 16 11:44:06 sams anacron[2363]: Anacron 2.3 start
Nov 16 11:44:07 sams anacron[2363]: Will run job `cro
Nov 16 11:44:07 sams anacron[2363]: Jobs will be exec
Nov 16 12:01:01 sams crond[3304]: (root) CMD (run-par
```

194 lines (13.9 KB) - last update : Thu Nov 16 11:44:07 2006

그림 38.5. Log File Alert

Clicking on the messages log file displays the logs in the file with the new lines in bold as illustrated below.

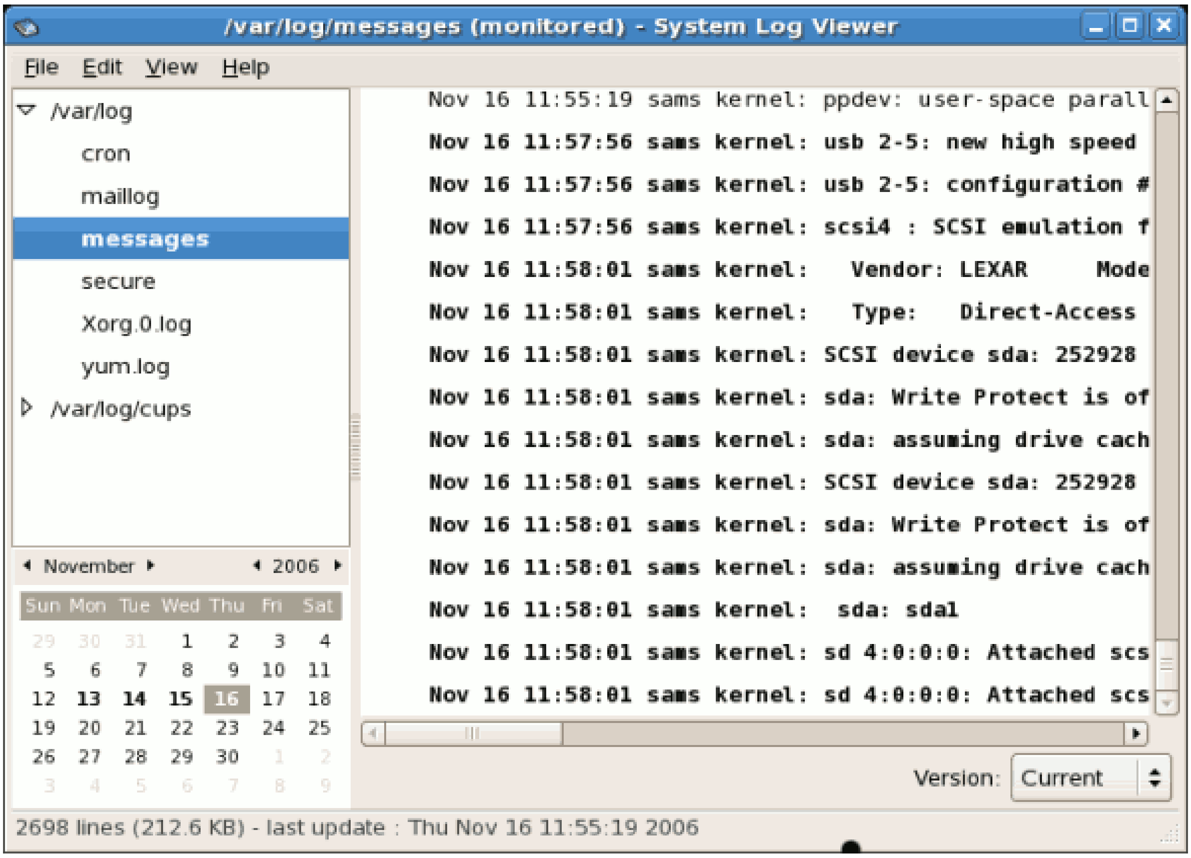

```
Nov 16 11:55:19 sams kernel: ppdev: user-space parall
Nov 16 11:57:56 sams kernel: usb 2-5: new high speed
Nov 16 11:57:56 sams kernel: usb 2-5: configuration #
Nov 16 11:57:56 sams kernel: scsi4 : SCSI emulation f
Nov 16 11:58:01 sams kernel:   Vendor: LEXAR     Mode
Nov 16 11:58:01 sams kernel:   Type:   Direct-Access
Nov 16 11:58:01 sams kernel: SCSI device sda: 252928
Nov 16 11:58:01 sams kernel: sda: Write Protect is of
Nov 16 11:58:01 sams kernel: sda: assuming drive cach
Nov 16 11:58:01 sams kernel: SCSI device sda: 252928
Nov 16 11:58:01 sams kernel: sda: Write Protect is of
Nov 16 11:58:01 sams kernel: sda: assuming drive cach
Nov 16 11:58:01 sams kernel:  sda: sda1
Nov 16 11:58:01 sams kernel: sd 4:0:0:0: Attached scs
Nov 16 11:58:01 sams kernel: sd 4:0:0:0: Attached scs
```

Version: Current

2698 lines (212.6 KB) - last update : Thu Nov 16 11:55:19 2006

그림 38.6. Log file contents

The new lines are displayed in bold for five seconds after which they are displayed in normal font.

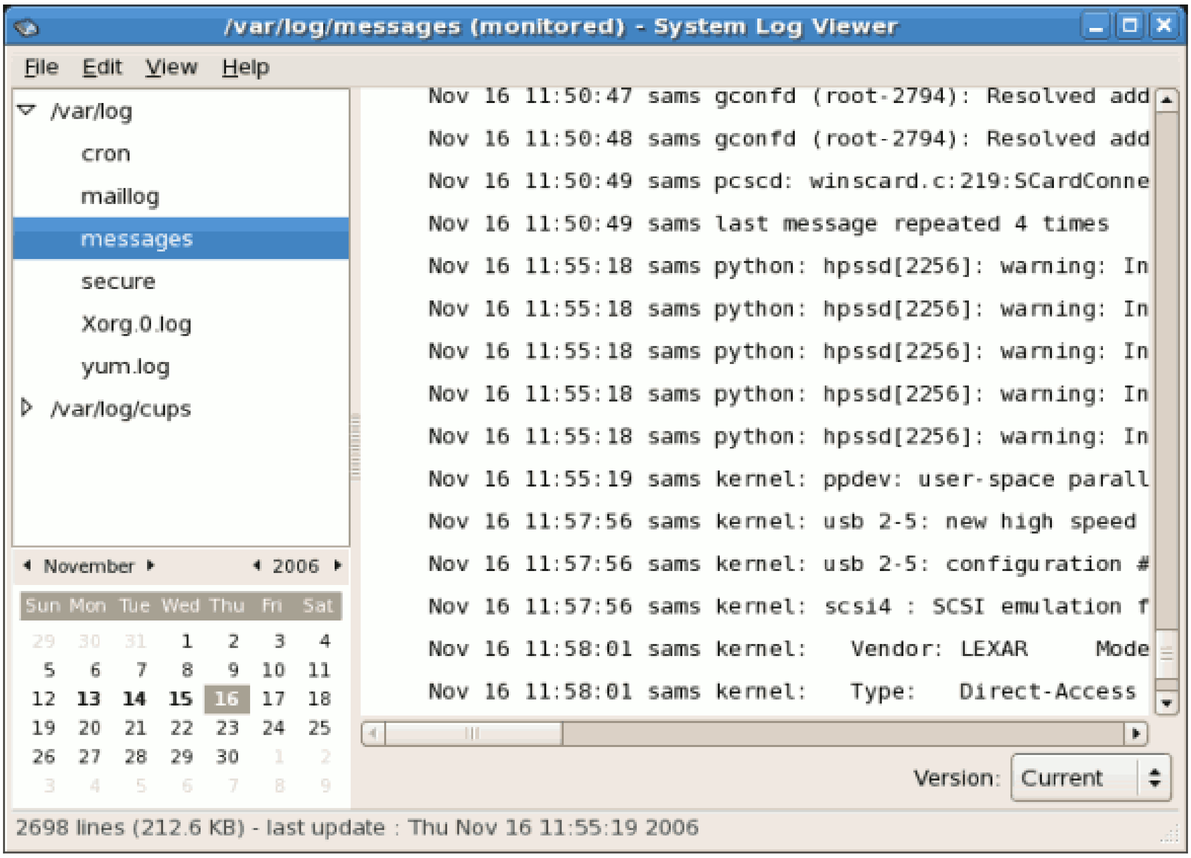

```
Nov 16 11:50:47 sams gconfd (root-2794): Resolved add
Nov 16 11:50:48 sams gconfd (root-2794): Resolved add
Nov 16 11:50:49 sams pcscd: winscard.c:219:SCardConne
Nov 16 11:50:49 sams last message repeated 4 times
Nov 16 11:55:18 sams python: hpssd[2256]: warning: In
Nov 16 11:55:18 sams python: hpssd[2256]: warning: In
Nov 16 11:55:18 sams python: hpssd[2256]: warning: In
Nov 16 11:55:18 sams python: hpssd[2256]: warning: In
Nov 16 11:55:18 sams python: hpssd[2256]: warning: In
Nov 16 11:55:19 sams kernel: ppdev: user-space parall
Nov 16 11:57:56 sams kernel: usb 2-5: new high speed
Nov 16 11:57:56 sams kernel: usb 2-5: configuration #
Nov 16 11:57:56 sams kernel: scsi4 : SCSI emulation f
Nov 16 11:58:01 sams kernel:   Vendor: LEXAR     Mode
Nov 16 11:58:01 sams kernel:   Type:   Direct-Access
```

Version: Current

2698 lines (212.6 KB) - last update : Thu Nov 16 11:55:19 2006

그림 38.7. Log file contents after five seconds

부 V. 시스템 감시

또한 시스템 관리자는 시스템 성능을 감시합니다. Red Hat Enterprise Linux에는 시스템 관리자의 업무를 보조하기 위한 도구가 포함되어 있습니다.

SystemTap

39.1. Introduction

SystemTap provides a simple command line interface and scripting language to simplify the gathering of information about the running Linux kernel so that it can be further analyzed. Data may be extracted, filtered, and summarized quickly and safely, to enable diagnoses of complex performance or functional problems.

SystemTap allows scripts to be written in the SystemTap scripting language, which are then compiled to C-code kernel modules and inserted into the kernel.

The essential idea behind a systemtap script is to name events, and to give them handlers. Whenever a specified event occurs, the Linux kernel runs the handler as if it were a quick subroutine, then resumes. There are several kind of events, such as entering or exiting a function, a timer expiring, or the entire systemtap session starting or stopping. A handler is a series of script language statements that specify the work to be done whenever the event occurs. This work normally includes extracting data from the event context, storing them into internal variables, or printing results.

39.2. Implementation

SystemTap takes a compiler-oriented approach to generating instrumentation. Refer to 그림 39.1. "Flow of Data in SystemTap" "Flow of data in SystemTap" for an overall diagram of SystemTap used in this discussion. In the upper right hand corner of the diagram is the probe.stp, the probe script the developer has written. This is parsed by the translator into parse trees. During this time the input is checked for syntax errors. The translator then performs elaboration, pulling in additional code from the script library and determining locations of probe points and variables from the debug information. After the elaboration is complete the translator can generate the probe.c, the kernel module in C.

The probe.c file is compiled into a regular kernel module, probe.ko, using the GCC compiler. The compilation may pull in support code from the runtime libraries. After GCC has generated the probe.ko, the SystemTap daemon is started to collect the output of the instrumentation module. The instrumentation module is loaded into the kernel, and data collection is started. Data from the instrumentation module is transferred to user-space via relayfs and displayed by the daemon. When the user hits Control-C the daemon unloads the module, which also shuts down the data collection process.

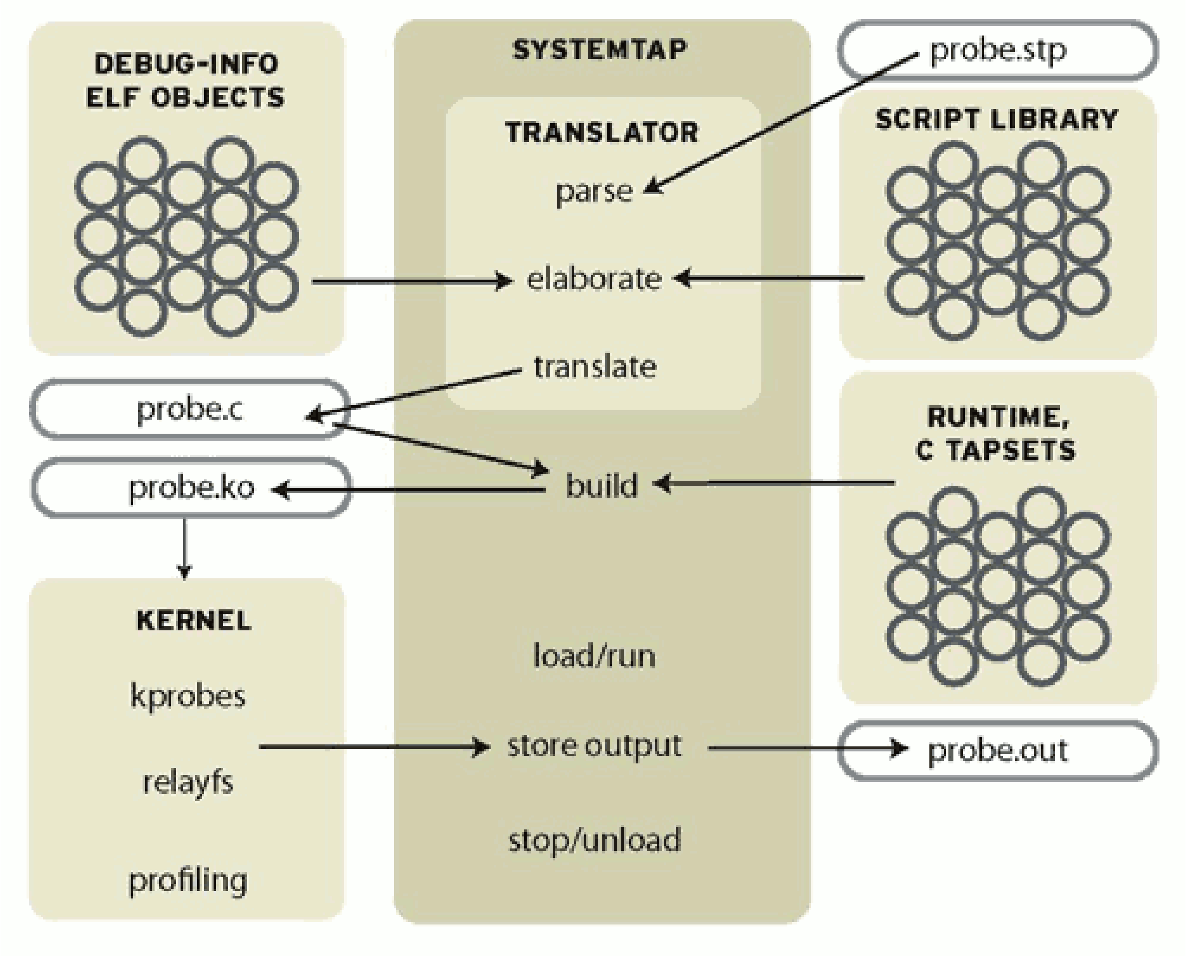

그림 39.1. Flow of Data in SystemTap

39.3. Using SystemTap

Systemtap works by translating a SystemTap script to C, running the system C compiler to create a kernel module from that. When the module is loaded, it activates all the probed events by hooking into the kernel. Then, as events occur on any processor, the compiled handlers run. Eventually, the session stops, the hooks are disconnected, and the module removed. This entire process is driven from a single command-line program, stap.

39.3.1. Tracing

The simplest kind of probe is simply to trace an event. This is the effect of inserting strategically located print statements into a program. This is often the first step of problem solving: explore by seeing a history of what has happened.

This style of instrumentation is the simplest. It just asks systemtap to print something at each event. To express this in the script language, you need to say where to probe and what to print there.

39.3.1.1. Where to Probe

Systemtap supports a number of built-in events. The library of scripts that comes with systemtap, each called a "tapset", may define additional ones defined in terms of the built-in family. See the stapprobes man page for details. All these events are named using a unified syntax that looks like dot-separated parameterized identifiers:

표 39.1. SystemTap Events

Event	Description
begin	The startup of the systemtap session.
end	The end of the systemtap session.
kernel.function("sys_open")	The entry to the function named sys_open in the kernel.
syscall.close.return	The return from the close system call..
module("ext3").statement(0xdeadbeef)	The addressed instruction in the ext3 filesystem driver.
timer.ms(200)	A timer that fires every 200 milliseconds.

We will use as a demonstration case that you would like to trace all function entries and exits in a source file, for example net/socket.c in the kernel. The kernel.function probe point lets you express that easily, since systemtap examines the kernel's debugging information to relate object code to source code. It works like a debugger: if you can name or place it, you can probe it. Use kernel.function("*@net/socket.c") for the function entries, and kernel.function("*@net/socket.c").return for the exits. Note the use of wildcards in the function name part, and the subsequent @FILENAME part. You can also put wildcards into the file name, and even add a colon (:) and a line number, if you want to restrict the search that precisely. Since systemtap will put a separate probe in every place that matches a probe point, a few wildcards can expand to hundreds or thousands of probes, so be careful what you ask for.

Once you identify the probe points, the skeleton of the systemtap script appears. The probe keyword introduces a probe point, or a comma-separated list of them. The following { and } braces enclose the handler for all listed probe points.

You can run this script as is, though with empty handlers there will be no output. Put the two lines into a new file. Run stap -v FILE. Terminate it any time with ^C. (The -v option tells systemtap to print more verbose messages during its processing. Try the -h option to see more options.)

39.3.1.2. What to Print

Since you are interested in each function that was entered and exited, a line should be printed for each, containing the function name. In order to make that list easy to read, systemtap should indent the lines so that functions called by other traced functions are nested deeper. To tell each single process apart from any others that may be running concurrently, systemtap should also print the process ID in the line.

시스템 정보 모으기

시스템을 설정하는 방법에 대해서 배우시기에 앞서 먼저 기본적인 시스템 정보를 모으는 방법부터 배우셔야 합니다. 예로 들면, 여유 메모리 용량과 사용 가능한 하드 드라이브 공간의 용량을 알아내는 방법, 하드 드라이브 파티션 하는 방법과 실행되고 있는 프로세스 알아내는 방법 등에 대한 정보를 알아 두셔야 합니다. 이 장에서는 간단한 명령어 및 몇개의 간단한 프로그램을 사용하여 Red Hat Enterprise Linux 시스템에서 이러한 유형의 정보를 검색하는 방법에 대해 설명해 보겠습니다.

40.1. 시스템 프로세스

The ps ax command displays a list of current system processes, including processes owned by other users. To display the owner alongside each process, use the ps aux command. This list is a static list; in other words, it is a snapshot of what was running when you invoked the command. If you want a constantly updated list of running processes, use top as described below.

The ps output can be long. To prevent it from scrolling off the screen, you can pipe it through less:

```
ps aux | less
```

You can use the ps command in combination with the grep command to see if a process is running. For example, to determine if Emacs is running, use the following command:

```
ps ax | grep emacs
```

The top command displays currently running processes and important information about them including their memory and CPU usage. The list is both real-time and interactive. An example of output from the top command is provided as follows:

```
top - 15:02:46 up 35 min,  4 users,  load average: 0.17, 0.65, 1.00
Tasks: 110 total,   1 running, 107 sleeping,   0 stopped,   2 zombie
Cpu(s): 41.1% us,  2.0% sy,  0.0% ni, 56.6% id,  0.0% wa,  0.3% hi,  0.0% si
Mem:    775024k total,   772028k used,     2996k free,    68468k buffers
Swap:  1048568k total,      176k used,  1048392k free,   441172k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
 4624 root      15   0 40192  18m 7228 S 28.4  2.4   1:23.21 X
 4926 mhideo    15   0 55564  33m 9784 S 13.5  4.4   0:25.96 gnome-terminal
 6475 mhideo    16   0  3612  968  760 R  0.7  0.1   0:00.11 top
 4920 mhideo    15   0 20872  10m 7808 S  0.3  1.4   0:01.61 wnck-applet
    1 root      16   0  1732  548  472 S  0.0  0.1   0:00.23 init
    2 root      34  19     0    0    0 S  0.0  0.0   0:00.00 ksoftirqd/0
    3 root       5 -10     0    0    0 S  0.0  0.0   0:00.03 events/0
    4 root       6 -10     0    0    0 S  0.0  0.0   0:00.02 khelper
    5 root       5 -10     0    0    0 S  0.0  0.0   0:00.00 kacpid
   29 root       5 -10     0    0    0 S  0.0  0.0   0:00.00 kblockd/0
   47 root      16   0     0    0    0 S  0.0  0.0   0:01.74 pdflush
   50 root      11 -10     0    0    0 S  0.0  0.0   0:00.00 aio/0
   30 root      15   0     0    0    0 S  0.0  0.0   0:00.05 khubd
   49 root      16   0     0    0    0 S  0.0  0.0   0:01.44 kswapd0
```

To exit top, press the q key.

표 40.1. "Interactive top commands" contains useful interactive commands that you can use with top. For more information, refer to the top(1) manual page.

표 40.1. Interactive top commands

명령어	설명
Space	즉시 화면을 재생합니다
h	도움말 화면을 보여줍니다.
k	프로세스를 강제 종료 (kill) 합니다. 종료할 프로세스 ID와 보낼 신호를 입력하셔야 합니다.
n	보여줄 프로세스의 수를 변경합니다. 보시려는 프로세스의 수를 입력하셔야 합니다.
u	사용자에 따라 목록 정렬
M	메모리 사용량에 따라 목록 정렬
P	CPU 사용량에 따라 CPU 목록 정렬

If you prefer a graphical interface for top, you can use the GNOME System Monitor. To start it from the desktop, select System > Administration > System Monitor or type gnome-system-monitor at a shell prompt (such as an XTerm). Select the Process Listing tab.

The GNOME System Monitor allows you to search for a process in the list of running processes. Using the Gnome System Monitor, you can also view all processes, your processes, or active processes.

The Edit menu item allows you to:

- 프로세스를 중지함.
- 프로세스를 계속 진행하거나 시작함.
- 프로세스를 종료함.
- 프로세스를 삭제함.
- 선택된 프로세스의 우선 순위를 변경함.
- 시스템 모니터 설정을 편집합니다. 이는 목록을 재생을 위한 초 간격 변경과 시스템 모니터 윈도우에서 보여주는 프로세스 영역의 선택을 포함합니다.

The View menu item allows you to:

- 활성화된 프로세스만 보기.
- 모든 프로세스 보기.
- 내 프로세스 보기
- 프로세스 의존성 보기
- 프로세스 숨기기
- 숨겨진 프로세스 보기.
- 메모리 맵 보기.
- 선택된 프로세스에 의해 열린 파일 보기.

To stop a process, select it and click End Process. Alternatively you can also stop a process by selecting it, clicking Edit on your menu and selecting Stop Process.

특정 행에 따라 정보를 정렬하시려면, 행 이름에 클릭하십시오. 이는 정보를 선택된 행에서 오름차순으로 정렬합니다. 행 이름을 다시 클릭하면 오름차순에서 내림차순으로 정렬하게 됩니다.

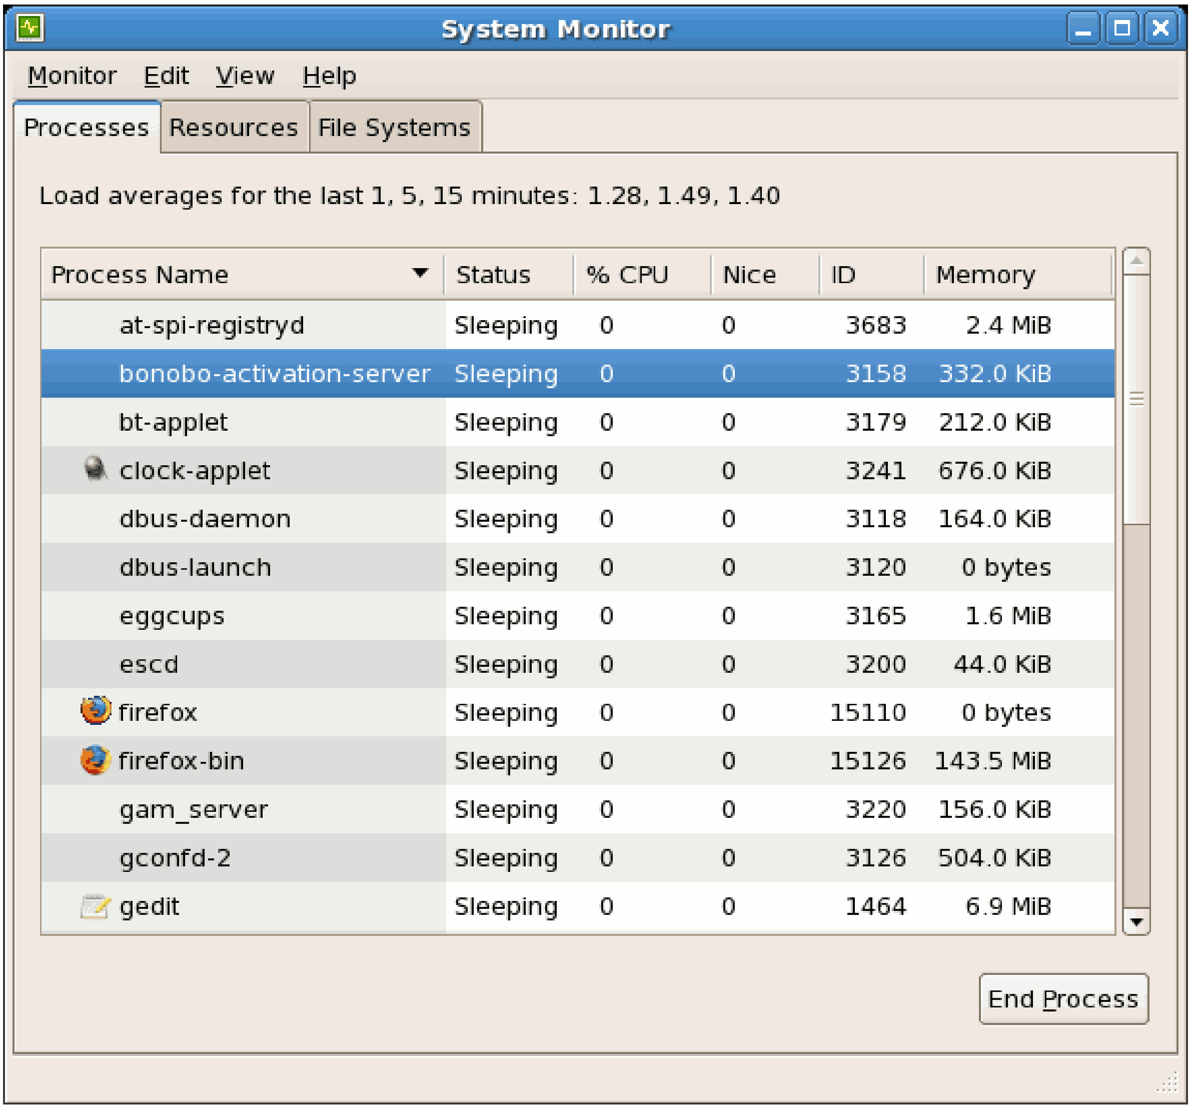

Load averages for the last 1, 5, 15 minutes: 1.28, 1.49, 1.40

Process Name	Status	% CPU	Nice	ID	Memory
at-spi-registryd	Sleeping	0	0	3683	2.4 MiB
bonobo-activation-server	Sleeping	0	0	3158	332.0 KiB
bt-applet	Sleeping	0	0	3179	212.0 KiB
clock-applet	Sleeping	0	0	3241	676.0 KiB
dbus-daemon	Sleeping	0	0	3118	164.0 KiB
dbus-launch	Sleeping	0	0	3120	0 bytes
eggcups	Sleeping	0	0	3165	1.6 MiB
escd	Sleeping	0	0	3200	44.0 KiB
firefox	Sleeping	0	0	15110	0 bytes
firefox-bin	Sleeping	0	0	15126	143.5 MiB
gam_server	Sleeping	0	0	3220	156.0 KiB
gconfd-2	Sleeping	0	0	3126	504.0 KiB
gedit	Sleeping	0	0	1464	6.9 MiB

그림 40.1. GNOME System Monitor

40.2. 메모리 사용량

The free command displays the total amount of physical memory and swap space for the system as well as the amount of memory that is used, free, shared, in kernel buffers, and cached.

```
              total       used       free     shared    buffers     cached
Mem:         645712     549720      95992          0     176248     224452
-/+ buffers/cache:      149020     496692
Swap:       1310712          0    1310712
```

The command free -m shows the same information in megabytes, which are easier to read.

```
              total       used       free     shared    buffers     cached
Mem:            630        536         93          0        172        219
-/+ buffers/cache:         145        485
Swap:          1279          0       1279
```

If you prefer a graphical interface for free, you can use the GNOME System Monitor. To start it from the desktop, go to System > Administration > System Monitor or type gnome-system-monitor at a shell prompt (such as an XTerm). Click on the Resources tab.

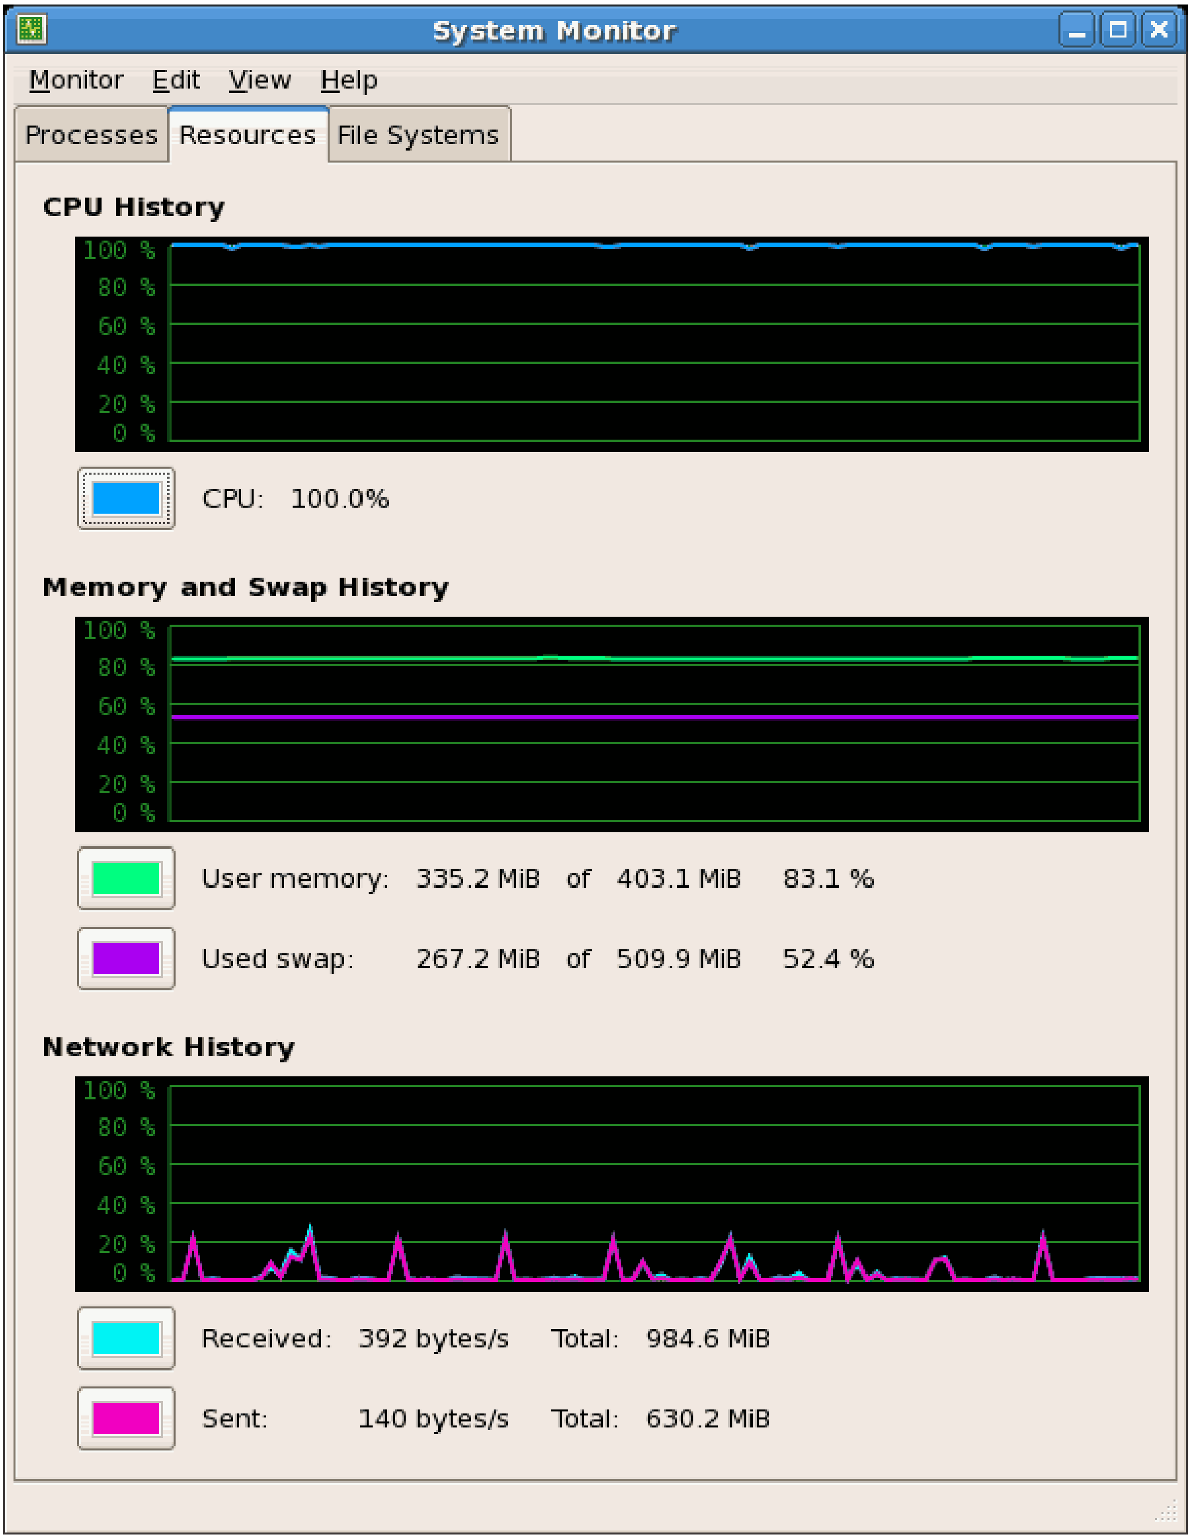

CPU: 100.0%

User memory: 335.2 MiB of 403.1 MiB 83.1 %

Used swap: 267.2 MiB of 509.9 MiB 52.4 %

Received: 392 bytes/s Total: 984.6 MiB

Sent: 140 bytes/s Total: 630.2 MiB

그림 40.2. GNOME System Monitor - Resources tab

40.3. 파일 시스템

The df command reports the system's disk space usage. If you type the command df at a shell prompt, the output looks similar to the following:

```
Filesystem              1K-blocks        Used  Available  Use%  Mounted  on
/dev/mapper/VolGroup00-LogVol00
                           11675568     6272120     4810348   57%  /  /dev/sda1
                        100691          9281        86211   10%  /boot
none                               322856             0      322856     0%  /dev/shm
```

By default, this utility shows the partition size in 1 kilobyte blocks and the amount of used and available disk space in kilobytes. To view the information in megabytes and gigabytes, use the command df -h. The -h argument stands for human-readable format. The output looks similar to the following:

```
Filesystem               Size   Used  Avail  Use%  Mounted  on
/dev/mapper/VolGroup00-LogVol00
                            12G    6.0G    4.6G    57%  /  /dev/sda1
    99M    9.1M     85M    10%  /boot
none       316M        0    316M     0%  /dev/shm
```

In the list of mounted partitions, there is an entry for /dev/shm. This entry represents the system's virtual memory file system.

The du command displays the estimated amount of space being used by files in a directory. If you type du at a shell prompt, the disk usage for each of the subdirectories is displayed in a list. The grand total for the current directory and subdirectories are also shown as the last line in the list. If you do not want to see the totals for all the subdirectories, use the command du -hs to see only the grand total for the directory in human-readable format. Use the du --help command to see more options.

To view the system's partitions and disk space usage in a graphical format, use the Gnome System Monitor by clicking on System > Administration > System Monitor or type gnome-system-monitor at a shell prompt (such as an XTerm). Select the File Systems tab to view the system's partitions. The figure below illustrates the File Systems tab.

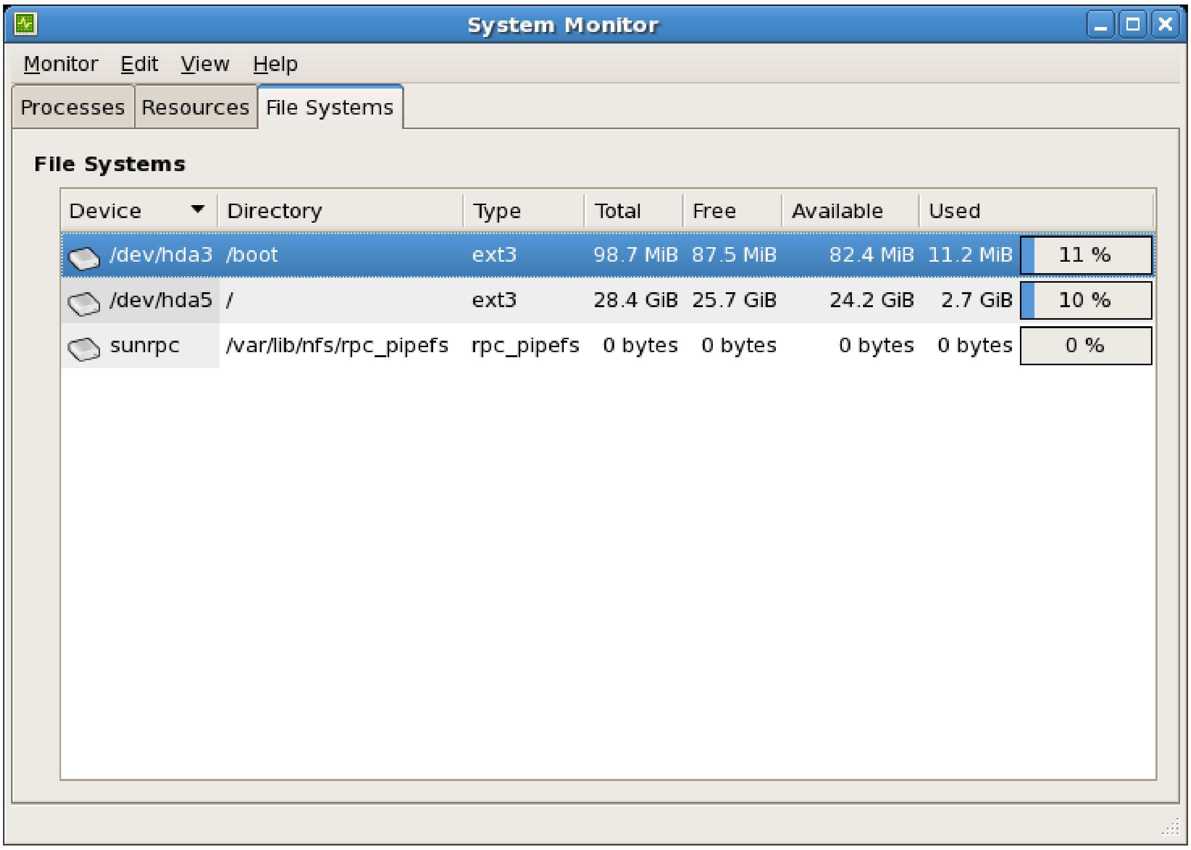

Device	Directory	Type	Total	Free	Available	Used	
/dev/hda3	/boot	ext3	98.7 MiB	87.5 MiB	82.4 MiB	11.2 MiB	11 %
/dev/hda5	/	ext3	28.4 GiB	25.7 GiB	24.2 GiB	2.7 GiB	10 %
sunrpc	/var/lib/nfs/rpc_pipefs	rpc_pipefs	0 bytes	0 bytes	0 bytes	0 bytes	0 %

그림 40.3. GNOME System Monitor - File Systems

40.4. 하드웨어

If you are having trouble configuring your hardware or just want to know what hardware is in your system, you can use the Hardware Browser application to display the hardware that can be probed. To start the program from the desktop, select System (the main menu on the panel) > Administration > Hardware or type hwbrowser at a shell prompt. As shown in 그림 40.4. "Hardware Browser" , it displays your CD-ROM devices, diskette drives, hard drives and their partitions, network devices, pointing devices, system devices, and video cards. Click on the category name in the left menu, and the information is displayed.

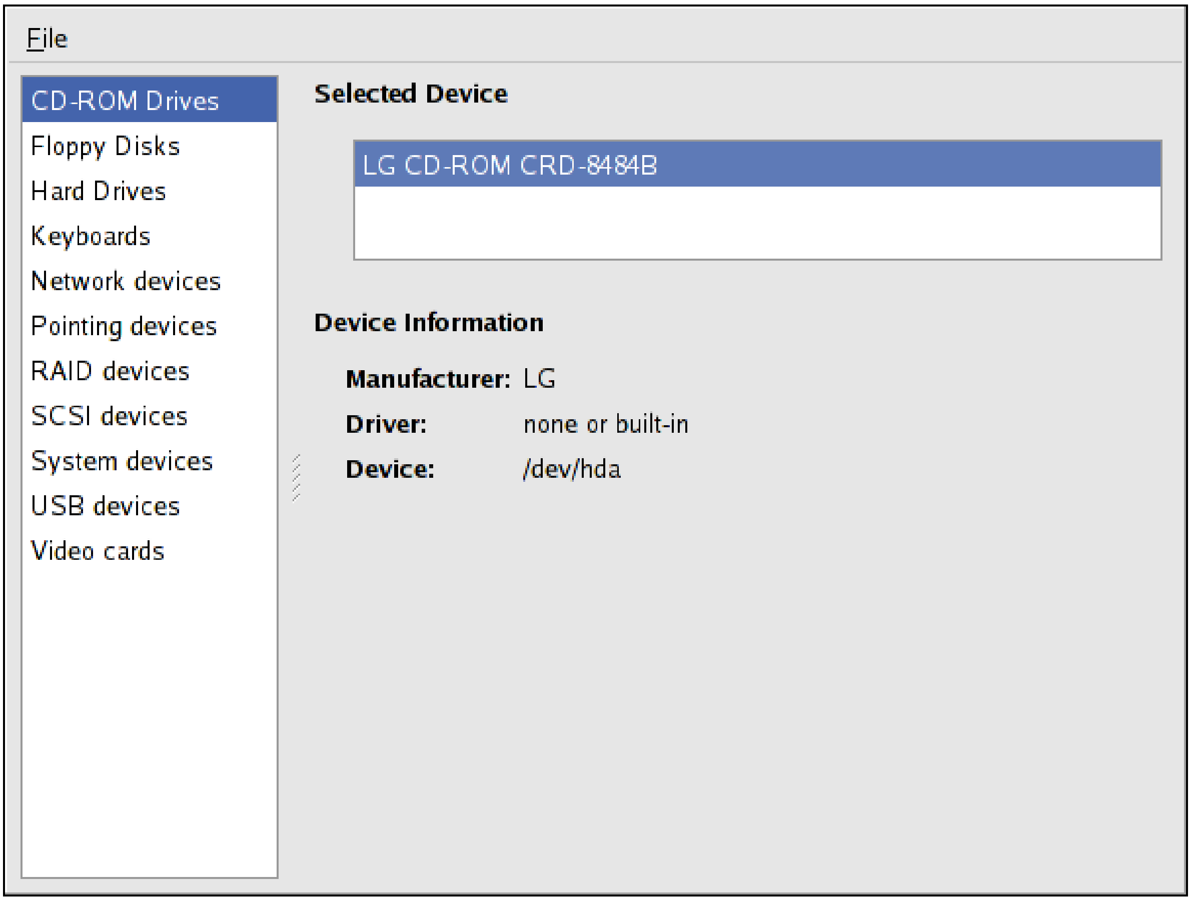

Selected Device

LG CD-ROM CRD-8484B

Device Information

Manufacturer: LG

Driver: none or built-in

Device: /dev/hda

그림 40.4. Hardware Browser

The Device Manager application can also be used to display your system hardware. This application can be started by selecting System (the main menu on the panel) > Administration > Hardware like the Hardware Browser. To start the application from a terminal, type hal-device-manager. Depending on your installation preferences, the graphical menu above may start this application or the Hardware Browser when clicked. The figure below illustrates the Device Manager window.

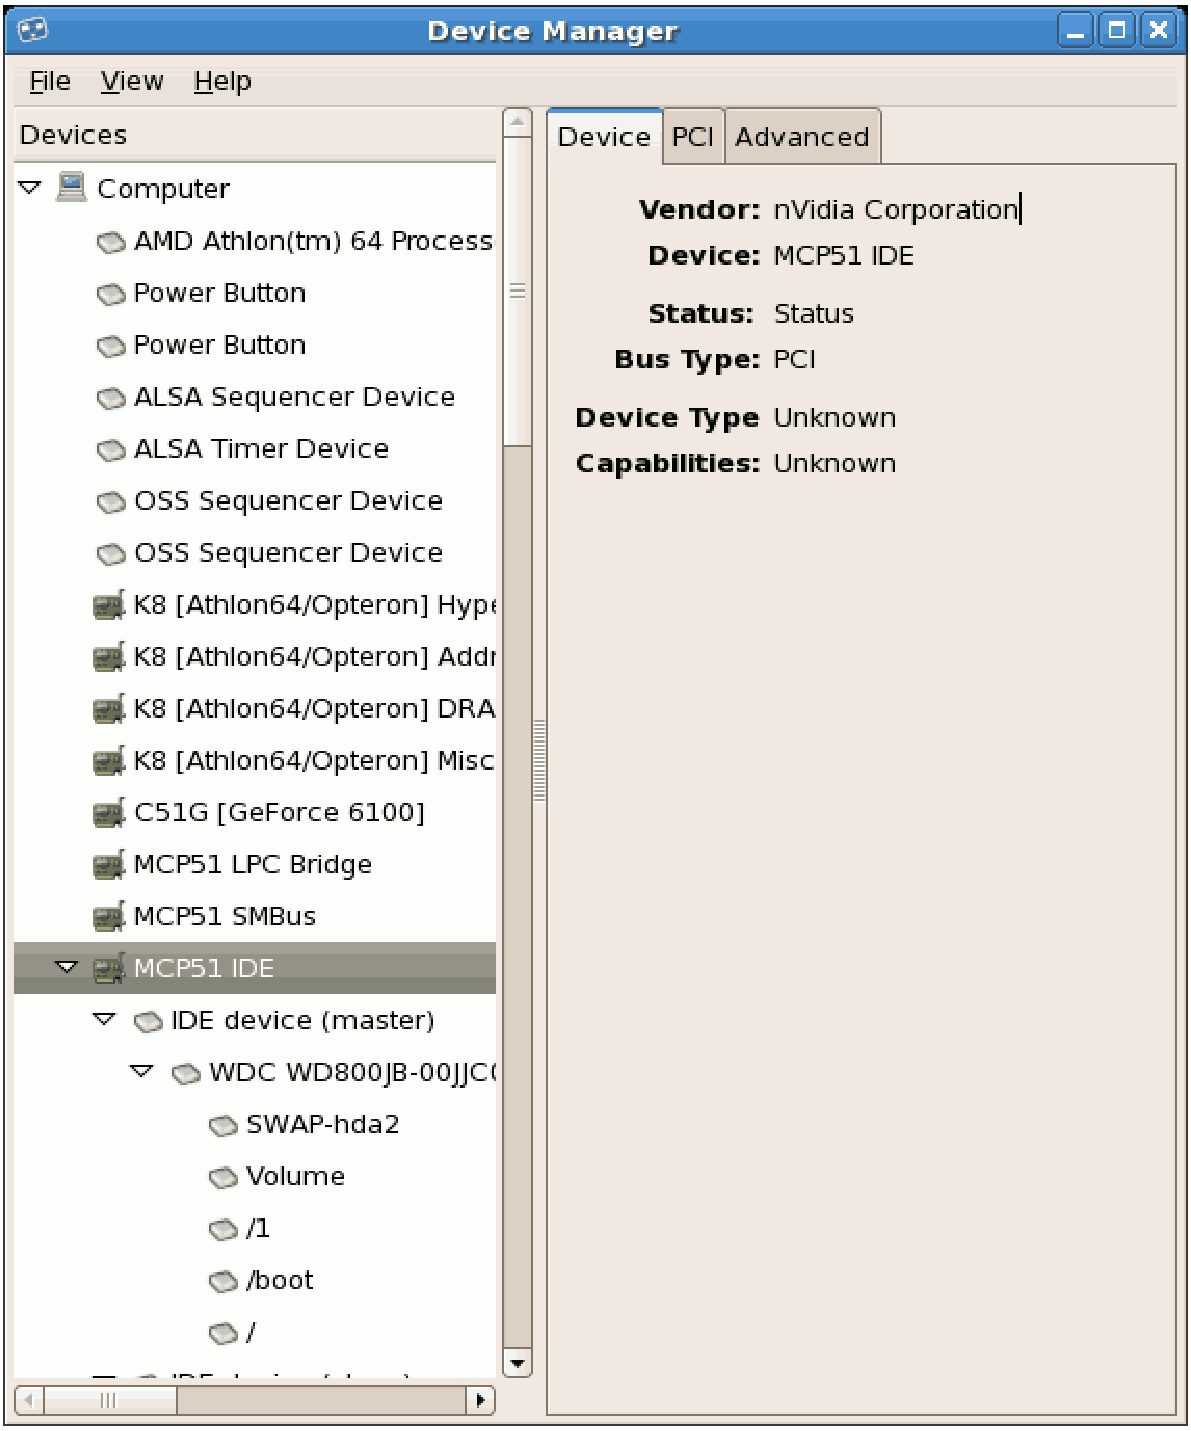

Vendor: nVidia Corporation
Device: MCP51 IDE
Status: Status
Bus Type: PCI
Device Type Unknown
Capabilities: Unknown

그림 40.5. Device Manager

You can also use the lspci command to list all PCI devices. Use the command lspci -v for more verbose information or lspci -vv for very verbose output.

For example, lspci can be used to determine the manufacturer, model, and memory size of a system's video card:

```
00:00.0 Host bridge: ServerWorks CNB20LE Host Bridge (rev 06)
00:00.1 Host bridge: ServerWorks CNB20LE Host Bridge (rev 06)
00:01.0 VGA compatible controller: S3 Inc. Savage 4 (rev 04)
00:02.0 Ethernet controller: Intel Corp. 82557/8/9 [Ethernet Pro 100] (rev 08)
00:0f.0 ISA bridge: ServerWorks OSB4 South Bridge (rev 50)
00:0f.1 IDE interface: ServerWorks OSB4 IDE Controller
```

```
00:0f.2 USB Controller: ServerWorks OSB4/CSB5 OHCI USB Controller (rev 04)
01:03.0 SCSI storage controller: Adaptec AIC-7892P U160/m (rev 02)
01:05.0 RAID bus controller: IBM ServeRAID Controller
```

The lspci is also useful to determine the network card in your system if you do not know the manufacturer or model number.

40.5. 추가 자료

시스템 정보 모으기에 대한 더 많은 정보를 원하신다면, 다음과 같은 자료를 참조하시기 바랍니다.

40.5.1. 설치된 문서 자료

- ps --help — Displays a list of options that can be used with ps.
- top manual page — Type man top to learn more about top and its many options.
- free manual page — type man free to learn more about free and its many options.
- df manual page — Type man df to learn more about the df command and its many options.
- du manual page — Type man du to learn more about the du command and its many options.
- lspci manual page — Type man lspci to learn more about the lspci command and its many options.
- /proc/ directory — The contents of the /proc/ directory can also be used to gather more detailed system information.

OProfile

OProfile is a low overhead, system-wide performance monitoring tool. It uses the performance monitoring hardware on the processor to retrieve information about the kernel and executables on the system, such as when memory is referenced, the number of L2 cache requests, and the number of hardware interrupts received. On a Red Hat Enterprise Linux system, the oprofile RPM package must be installed to use this tool.

Many processors include dedicated performance monitoring hardware. This hardware makes it possible to detect when certain events happen (such as the requested data not being in cache). The hardware normally takes the form of one or more counters that are incremented each time an event takes place. When the counter value, essentially rolls over, an interrupt is generated, making it possible to control the amount of detail (and therefore, overhead) produced by performance monitoring.

OProfile은 이 하드웨어 (또는 성능 감시 하드웨어가 없는 경우 타이머가 같은 대체 하드웨어)를 사용하여 카운터가 인터럽트를 발생할 때마다 성능과 관련된 데이터 샘플을 수집합니다. 이 샘플은 주기적으로 디스크에 기록되며; 이후 이 샘플에 포함된 데이터를 사용하여 시스템 수준 성능과 응용 프로그램 수준 성능에 대한 리포트를 생성할 수 있습니다.

OProfile은 유용한 도구이지만 사용하실때 몇가지 제한 사항을 알고 계셔야 합니다:

- 공유 라이브러리 사용 — --separate=library 옵션이 사용되지 않은 한 공유 라이브러리 코드의 샘플은 특정 응용 프로그램에 속하지 않습니다.

- 성능 감시 샘플은 정확하지 않습니다 — 성능 감시 기록기가 샘플을 수집시 인터럽트 처리는 0으로 나누는 것에 대해 예외사항(Exception)과 같이 정확하지 않습니다. 프로세서 명령이 순서대로 실행되지 않기 때문에 (out-of-order execution), 샘플은 근접하지만 정확하지 않게 기록될 수 있습니다.

- opreport does not associate samples for inline functions' properly — opreport uses a simple address range mechanism to determine which function an address is in. Inline function samples are not attributed to the inline function but rather to the function the inline function was inserted into.

- OProfile accumulates data from multiple runs — OProfile is a system-wide profiler and expects processes to start up and shut down multiple times. Thus, samples from multiple runs accumulate. Use the command opcontrol --reset to clear out the samples from previous runs.

- CPU 제한과 관련되지 않은 성능 문제들 — OProfile은 CPU 제한과 관련된 프로세스 문제점을 찾도록 고안되었습니다. 따라서 OProfile은 잠금이 해제되기를 기다리고 있거나 다른 작업이 발생하기를 (예, I/O 장치가 작업을 마치기를) 기다리면서 프로세스가 멈출 경우 그러한 프로세스를 찾아내지 못합니다.

41.1. 도구 개요

표 41.1. "OProfile 명령" provides a brief overview of the tools provided with the oprofile package.

표 41.1. OProfile 명령

명령	설명
ophelp	Displays available events for the system's processor along with a brief description of each.
opimport	샘플 데이터베이스 파일을 외부 이진 형식에서 시스템의 원시 형식으로 변환합니다. 다른 구조에서 수집한 샘플 데이터베이스를 분석할 경우에만 이 옵션을 사용하십시오.

명령	설명
opannotate	Creates annotated source for an executable if the application was compiled with debugging symbols. Refer to 41.5.4절. "Using opannotate" for details.
opcontrol	Configures what data is collected. Refer to 41.2절. "OProfile 설정" for details.
opreport	Retrieves profile data. Refer to 41.5.1절. "Using opreport" for details.
oprofiled	데몬으로 실행되어 샘플 데이터를 디스크에 주기적으로 기록합니다.

41.2. OProfile 설정

Before OProfile can be run, it must be configured. At a minimum, selecting to monitor the kernel (or selecting not to monitor the kernel) is required. The following sections describe how to use the opcontrol utility to configure OProfile. As the opcontrol commands are executed, the setup options are saved to the /root/.oprofile/daemonrc file.

41.2.1. 커널 지정

먼저 OProfile이 커널을 감시해야할지 여부를 설정해 주십시오. 이 옵션만 설정하시면 OProfile을 시작할 수 있습니다. 모든 다른 옵션은 선택 사항입니다.

커널을 감시하려면 루트로 로그인하신 후 다음 명령을 실행하십시오:

```
opcontrol --setup --vmlinux=/usr/lib/debug/lib/modules/`uname -r`/vmlinux
```

알림

The debuginfo package must be installed (which contains the uncompressed kernel) in order to monitor the kernel.

OProfile이 커널을 감시하지 않도록 설정하시려면 루트로 로그인하신 후 다음 명령을 실행하시기 바랍니다:

```
opcontrol --setup --no-vmlinux
```

This command also loads the oprofile kernel module, if it is not already loaded, and creates the /dev/oprofile/ directory, if it does not already exist. Refer to 41.6절. "Understanding /dev/oprofile/" for details about this directory.

알림

Even if OProfile is configured not to profile the kernel, the SMP kernel still must be running so that the oprofile module can be loaded from it.

Setting whether samples should be collected within the kernel only changes what data is collected, not how or where the collected data is stored. To generate different sample files for the kernel and application libraries, refer to 41.2.3절. "커널과 사용자 영역 프로파일 분리하기" .

41.2.2. 감시기가 기록할 사건 설정하기

Most processors contain counters, which are used by OProfile to monitor specific events. As shown in 표 41.2. "OProfile 프로세서와 카운터" , the number of counters available depends on the processor.

표 41.2. OProfile 프로세서와 카운터

프로세서	cpu_type	카운터 수
Pentium Pro	i386/ppro	2
Pentium II	i386/pii	2
Pentium III	i386/piii	2
Pentium 4 (non-hyper-threaded)	i386/p4	8
Pentium 4 (hyper-threaded)	i386/p4-ht	4
Athlon	i386/athlon	4
AMD64	x86-64/hammer	4
Itanium	ia64/itanium	4
Itanium 2	ia64/itanium2	4
TIMER_INT	timer	1
IBM eServer iSeries and pSeries	timer	1
	ppc64/power4	8
	ppc64/power5	6
	ppc64/970	8
IBM eServer S/390 and S/390x	timer	1
IBM eServer zSeries	timer	1

Use 표 41.2. "OProfile 프로세서와 카운터" to verify that the correct processor type was detected and to determine the number of events that can be monitored simultaneously. timer is used as the processor type if the processor does not have supported performance monitoring hardware.

If timer is used, events cannot be set for any processor because the hardware does not have support for hardware performance counters. Instead, the timer interrupt is used for profiling.

If timer is not used as the processor type, the events monitored can be changed, and counter 0 for the processor is set to a time-based event by default. If more than one counter exists on the processor, the counters other than counter 0 are not set to an event by default. The default events monitored are shown in 표 41.3. "기본 사건" .

표 41.3. 기본 사건

프로세서	Default Event for Counter	설명
Pentium Pro, Pentium II, Pentium III, Athlon, AMD64	CPU_CLK_UNHALTED	The processor's clock is not halted
Pentium 4 (HT and non-HT)	GLOBAL_POWER_EVENTS	프로세서가 멈추지 않은 시간
Itanium 2	CPU_CYCLES	CPU 사이클
TIMER_INT	(없음)	각 타이머 인터럽트의 샘플
ppc64/power4	CYCLES	Processor Cycles
ppc64/power5	CYCLES	Processor Cycles
ppc64/970	CYCLES	Processor Cycles

한번에 감시할 수 있는 사건의 수는 프로세서에서 사용되는 카운터의 수에 의해 결정됩니다. 그러나 한개 당 한개씩 상관 관계를 갖는 것이 아니라; 일부 프로세서에서 특정 사건은 특정 카운터에 대응해야 합니다. 사용 가능한 카운터의 수를 알아보시려면 다음 명령을 실행하시기 바랍니다:

```
ls -d /dev/oprofile/[0-9]*
```

The events available vary depending on the processor type. To determine the events available for profiling, execute the following command as root (the list is specific to the system's processor type):

```
ophelp
```

The events for each counter can be configured via the command line or with a graphical interface. For more information on the graphical interface, refer to 41.8절. "그래픽 인터페이스" . If the counter cannot be set to a specific event, an error message is displayed.

To set the event for each configurable counter via the command line, use opcontrol:

```
opcontrol --event=<event-name>:<sample-rate>
```

Replace <event-name> with the exact name of the event from ophelp, and replace <sample-rate> with the number of events between samples.

41.2.2.1. 샘플 수집 속도

By default, a time-based event set is selected. It creates a sample every 100,000 clock cycles per processor. If the timer interrupt is used, the timer is set to whatever the jiffy rate is and is not user-settable. If the cpu_type is not timer, each event can have a sampling rate set for it. The sampling rate is the number of events between each sample snapshot.

카운터에 사건을 설정할 때 샘플 속도도 지정 가능합니다:

```
opcontrol --event=<event-name>:<sample-rate>
```

Replace <sample-rate> with the number of events to wait before sampling again. The smaller the count, the more frequent the samples. For events that do not happen frequently, a lower count may be needed to capture the event instances.

주의

샘플링 속도를 설정하실때 매우 주의하셔야 합니다. 너무 자주 샘플링하게되면 시스템의 작업 부하가 높아져서 시스템이 정지한 것처럼 나타나거나 실제로 멈출 수도 있습니다.

41.2.2.2. 유닛 마스크 (Unit Masks)

Some user performance monitoring events may also require unit masks to further define the event.

Unit masks for each event are listed with the ophelp command. The values for each unit mask are listed in hexadecimal format. To specify more than one unit mask, the hexadecimal values must be combined using a bitwise or operation.

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>
```

41.2.3. 커널과 사용자 영역 프로파일 분리하기

By default, kernel mode and user mode information is gathered for each event. To configure OProfile to ignore events in kernel mode for a specific counter, execute the following command:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:0
```

해당 카운터에서 커널 모드를 다시 프로파일링하도록 설정하시려면 다음 명령을 실행하시기 바랍니다:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:1
```

To configure OProfile to ignore events in user mode for a specific counter, execute the following command:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:<kernel>:0
```

해당 카운터에서 사용자 모드를 다시 프로파일링하도록 설정하시려면 다음 명령을 실행하시기 바랍니다:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:<kernel>:1
```

OProfile 데몬이 프로파일 데이터를 샘플 파일에 기록시 커널 데이터와 라이브러리 프로파일 데이터를 별개의 샘플 파일에 분리하여 기록할 수 있습니다. 데몬이 샘플 파일을 기록하는 방식을 설정하시려면, 루트 사용자로 로그인하신 후 다음 명령을 실행하십시오:

```
opcontrol --separate=<choice>
```

<choice> can be one of the following:

- none — do not separate the profiles (default)
- library — generate per-application profiles for libraries

- kernel — generate per-application profiles for the kernel and kernel modules
- all — generate per-application profiles for libraries and per-application profiles for the kernel and kernel modules

만일 --separate=library 옵션이 사용된다면 라이브러리의 이름을 비롯하여 실행 파일의 이름을 입력하십시오.

> **알림**
>
> These configuration changes will take effect when oprofile is restarted.

41.3. OProfile 시작과 정지

OProfile을 사용하여 시스템 감시를 시작하시려면 루트로 로그인하신 후 다음 명령을 실행하십시오:

```
opcontrol --start
```

다음과 유사한 결과가 출력될 것입니다:

```
Using log file /var/lib/oprofile/oprofiled.log Daemon started. Profiler running.
```

The settings in /root/.oprofile/daemonrc are used.

The OProfile daemon, oprofiled, is started; it periodically writes the sample data to the /var/lib/oprofile/samples/ directory. The log file for the daemon is located at /var/lib/oprofile/oprofiled.log.

To stop the profiler, execute the following command as root:

```
opcontrol --shutdown
```

41.4. 데이터 저장

가끔은 특정 시기에 샘플을 저장하는 것이 유용합니다. 예를 들면 실행 프로그램을 감시할 경우 다른 입력 데이터 종류에 따라서 다른 샘플을 수집하는 것이 유용할 수 있습니다. 만일 감시할 사건의 수가 프로세서에서 사용 가능한 숫자를 초과한다면, OProfile을 여러개 실행하여 데이터를 수집하고 매번 샘플 데이터를 다른 파일에 저장할 수 있습니다.

To save the current set of sample files, execute the following command, replacing <name> with a unique descriptive name for the current session.

```
opcontrol --save=<name>
```

The directory /var/lib/oprofile/samples/name/ is created and the current sample files are copied to it.

41.5. 데이터 분석

Periodically, the OProfile daemon, oprofiled, collects the samples and writes them to the /var/lib/oprofile/samples/ directory. Before reading the data, make sure all data has been written to this directory by executing the following command as root:

```
opcontrol --dump
```

Each sample file name is based on the name of the executable. For example, the samples for the default event on a Pentium III processor for /bin/bash becomes:

```
\{root\}/bin/bash/\{dep\}/\{root\}/bin/bash/CPU_CLK_UNHALTED.100000
```

데이터가 수집되면 샘플 데이터를 프로파일하기 위하여 다음과 같은 도구를 사용하실 수 있습니다:

- opreport
- opannotate

프로파일된 이진(binary)와 함께 이 도구를 사용하여 보다 자세한 분석을 위해 리포트를 생성할 수 있습니다.

> **경고**
>
> The executable being profiled must be used with these tools to analyze the data. If it must change after the data is collected, backup the executable used to create the samples as well as the sample files. Please note that the sample file and the binary have to agree. Making a backup isn't going to work if they do not match. oparchive can be used to address this problem.

Samples for each executable are written to a single sample file. Samples from each dynamically linked library are also written to a single sample file. While OProfile is running, if the executable being monitored changes and a sample file for the executable exists, the existing sample file is automatically deleted. Thus, if the existing sample file is needed, it must be backed up, along with the executable used to create it before replacing the executable with a new version. The oprofile analysis tools use the executable file that created the samples during analysis. If the executable changes the analysis tools will be unable to analyze the associated samples. Refer to 41.4절. "데이터 저장" for details on how to backup the sample file.

41.5.1. Using opreport

The opreport tool provides an overview of all the executables being profiled.

The following is part of a sample output:

```
Profiling through timer interrupt
TIMER:0|
samples|      %|
------------------
25926 97.5212 no-vmlinux
```

```
359   1.3504  pi
65    0.2445  Xorg
62    0.2332  libvte.so.4.4.0
56    0.2106  libc-2.3.4.so
34    0.1279  libglib-2.0.so.0.400.7
19    0.0715  libXft.so.2.1.2
17    0.0639  bash
8     0.0301  ld-2.3.4.so
8     0.0301  libgdk-x11-2.0.so.0.400.13
6     0.0226  libgobject-2.0.so.0.400.7
5     0.0188  oprofiled
4     0.0150  libpthread-2.3.4.so
4     0.0150  libgtk-x11-2.0.so.0.400.13
3     0.0113  libXrender.so.1.2.2
3     0.0113  du
1     0.0038  libcrypto.so.0.9.7a
1     0.0038  libpam.so.0.77
1     0.0038  libtermcap.so.2.0.8
1     0.0038  libX11.so.6.2
1     0.0038  libgthread-2.0.so.0.400.7
1     0.0038  libwnck-1.so.4.9.0
```

Each executable is listed on its own line. The first column is the number of samples recorded for the executable. The second column is the percentage of samples relative to the total number of samples. The third column is the name of the executable.

Refer to the opreport man page for a list of available command line options, such as the -r option used to sort the output from the executable with the smallest number of samples to the one with the largest number of samples.

41.5.2. Using opreport on a Single Executable

To retrieve more detailed profiled information about a specific executable, use opreport:

```
opreport <mode> <executable>
```

<executable> must be the full path to the executable to be analyzed. <mode> must be one of the following:

-l

List sample data by symbols. For example, the following is part of the output from running the command opreport -l /lib/tls/libc-<version>.so:

```
samples  %         symbol name
12       21.4286   __gconv_transform_utf8_internal
5         8.9286   _int_malloc
4         7.1429   malloc
3         5.3571   __i686.get_pc_thunk.bx
3         5.3571   _dl_mcount_wrapper_check
3         5.3571   mbrtowc
3         5.3571   memcpy
2         3.5714   _int_realloc
2         3.5714   _nl_intern_locale_data
2         3.5714   free
2         3.5714   strcmp
1         1.7857   __ctype_get_mb_cur_max
1         1.7857   __unregister_atfork
1         1.7857   __write_nocancel
1         1.7857   _dl_addr
1         1.7857   _int_free
```

```
1           1.7857   _itoa_word
1           1.7857   calc_eclosure_iter
1           1.7857   fopen@@GLIBC_2.1
1           1.7857   getpid
1           1.7857   memmove
1           1.7857   msort_with_tmp
1           1.7857   strcpy
1           1.7857   strlen
1           1.7857   vfprintf
1           1.7857   write
```

The first column is the number of samples for the symbol, the second column is the percentage of samples for this symbol relative to the overall samples for the executable, and the third column is the symbol name.

샘플을 가장 큰 숫자에서 가장 작은 숫자 (반대 순서)로 정렬하여 출력하시려면 -l 옵션에 -r 옵션을 함께 사용하시면 됩니다.

-i <symbol-name>

List sample data specific to a symbol name. For example, the following output is from the command opreport -l -i __gconv_transform_utf8_internal /lib/tls/libc-<version>.so:

```
samples   %          symbol name
12         100.000   __gconv_transform_utf8_internal
```

첫번째 줄은 심볼/실행가능 프로그램 조합에 대한 요약 정보입니다.

The first column is the number of samples for the memory symbol. The second column is the percentage of samples for the memory address relative to the total number of samples for the symbol. The third column is the symbol name.

-d

List sample data by symbols with more detail than -l. For example, the following output is from the command opreport -l -d __gconv_transform_utf8_internal /lib/tls/libc-<version>.so:

```
vma        samples  %          symbol name
00a98640  12         100.000   __gconv_transform_utf8_internal
00a98640  1            8.3333
00a9868c  2           16.6667
00a9869a  1            8.3333
00a986c1  1            8.3333
00a98720  1            8.3333
00a98749  1            8.3333
00a98753  1            8.3333
00a98789  1            8.3333
00a98864  1            8.3333
00a98869  1            8.3333
00a98b08  1            8.3333
```

각 심볼마다 -l 옵션을 사용할 때와 동일한 데이터가 나타나지만, 사용된 가상 메모리 주소가 나타나는 차이점이 있습니다. 각 가상 메모리 주소에 대하여 샘플의 숫자와 심볼의 샘플 수에 비교한 샘플의 비율을 보여줍니다.

-x<symbol-name>

출력 결과에서 콤마로 구분된 심볼 목록을 제외시킵니다.

session:<name>

Specify the full path to the session or a directory relative to the /var/lib/oprofile/samples/ directory.

41.5.3. Getting more detailed output on the modules

OProfile collects data on a system-wide basis for kernel- and user-space code running on the machine. However, once a module is loaded into the kernel, the information about the origin of the kernel module is lost. The module could have come from the initrd file on boot up, the directory with the various kernel modules, or a locally created kernel module. As a result when OProfile records sample for a module, it just lists the samples for the modules for an executable in the root directory, but this is unlikely to be the place with the actual code for the module. You will need to take some steps to make sure that analysis tools get the executable.

For example on an AMD64 machine the sampling is set up to record "Data cache accesses" and "Data cache misses" and assuming you would like to see the data for the ext3 module:

```
~]$ opreport /ext3
CPU: AMD64 processors, speed 797.948 MHz (estimated)
Counted DATA_CACHE_ACCESSES events (Data cache accesses) with a unit mask of 0x00 (No unit mask) count 500000
Counted DATA_CACHE_MISSES events (Data cache misses) with a unit mask of 0x00 (No unit mask) count 500000
DATA_CACHE_ACC...|DATA_CACHE_MIS...|
samples|        %|  samples|        %|
------------------------------------
148721 100.000        1493 100.000 ext3
```

To get a more detailed view of the actions of the module, you will need to either have the module unstripped (e.g. installed from a custom build) or have the debuginfo RPM installed for the kernel.

Find out which kernel is running, "uname -a", get the appropriate debuginfo rpm, and install on the machine.

Then make a symbolic link so oprofile finds the code for the module in the correct place:

```
~]# ln -s /lib/modules/`uname -r`/kernel/fs/ext3/ext3.ko /ext3
```

Then the detailed information can be obtained with:

```
~]# opreport image:/ext3 -l|more
warning: could not check that the binary file /ext3 has not been modified since the profile was taken. Results may be
 inaccurate.
CPU: AMD64 processors, speed 797.948 MHz (estimated)
Counted DATA_CACHE_ACCESSES events (Data cache accesses) with a unit mask of 0x00 (No unit mask) count 500000
Counted DATA_CACHE_MISSES events (Data cache misses) with a unit mask of 0x00 (No unit mask) count 500000
samples  %        samples  %        symbol name
16728    11.2479  7         0.4689  ext3_group_sparse
16454    11.0637  4         0.2679  ext3_count_free_blocks
14583     9.8056  51        3.4159  ext3_fill_super
8281      5.5681  129       8.6403  ext3_ioctl
7810      5.2514  62        4.1527  ext3_write_info
7286      4.8991  67        4.4876  ext3_ordered_writepage
6509      4.3767  130       8.7073  ext3_new_inode
6378      4.2886  156      10.4488  ext3_new_block
5932      3.9887  87        5.8272  ext3_xattr_block_list
...
```

41.5.4. Using opannotate

The opannotate tool tries to match the samples for particular instructions to the corresponding lines in the source code. The resulting files generated should have the samples for the lines at the left. It also puts in a comment at the beginning of each function listing the total samples for the function.

For this utility to work, the executable must be compiled with GCC's -g option. By default, Red Hat Enterprise Linux packages are not compiled with this option.

The general syntax for opannotate is as follows:

```
opannotate --search-dirs <src-dir> --source <executable>
```

The directory containing the source code and the executable to be analyzed must be specified. Refer to the opannotate man page for a list of additional command line options.

41.6. Understanding /dev/oprofile/

The /dev/oprofile/ directory contains the file system for OProfile. Use the cat command to display the values of the virtual files in this file system. For example, the following command displays the type of processor OProfile detected:

```
cat /dev/oprofile/cpu_type
```

A directory exists in /dev/oprofile/ for each counter. For example, if there are 2 counters, the directories /dev/oprofile/0/ and dev/oprofile/1/ exist.

카운터에 사용되는 각 디렉토리는 다음 파일들을 포함합니다:

- count — The interval between samples.
- enabled — If 0, the counter is off and no samples are collected for it; if 1, the counter is on and samples are being collected for it.
- event — The event to monitor.
- kernel — If 0, samples are not collected for this counter event when the processor is in kernel-space; if 1, samples are collected even if the processor is in kernel-space.
- unit_mask — Defines which unit masks are enabled for the counter.
- user — If 0, samples are not collected for the counter event when the processor is in user-space; if 1, samples are collected even if the processor is in user-space.

The values of these files can be retrieved with the cat command. For example:

```
cat /dev/oprofile/0/count
```

41.7. 사용법 예제

OProfile는 개발자가 응용 프로그램의 성능을 분석하는데 사용될 뿐만 아니라 다음과 같이 시스템 관리자가 시스템 분석을 수행하는데도 사용될 수 있습니다:

- Determine which applications and services are used the most on a system — opreport can be used to determine how much processor time an application or service uses. If the system is used for multiple services but is under performing, the services consuming the most processor time can be moved to dedicated systems.
- Determine processor usage — The CPU_CLK_UNHALTED event can be monitored to determine the processor load over a given period of time. This data can then be used to determine if additional processors or a faster processor might improve system performance.

41.8. 그래픽 인터페이스

Some OProfile preferences can be set with a graphical interface. To start it, execute the oprof_start command as root at a shell prompt. To use the graphical interface, you will need to have the oprofile-gui package installed.

After changing any of the options, save them by clicking the Save and quit button. The preferences are written to /root/.oprofile/daemonrc, and the application exits. Exiting the application does not stop OProfile from sampling.

On the Setup tab, to set events for the processor counters as discussed in 41.2.2절. "감시기가 기록할 사건 설정하기", select the counter from the pulldown menu and select the event from the list. A brief description of the event appears in the text box below the list. Only events available for the specific counter and the specific architecture are displayed. The interface also displays whether the profiler is running and some brief statistics about it.

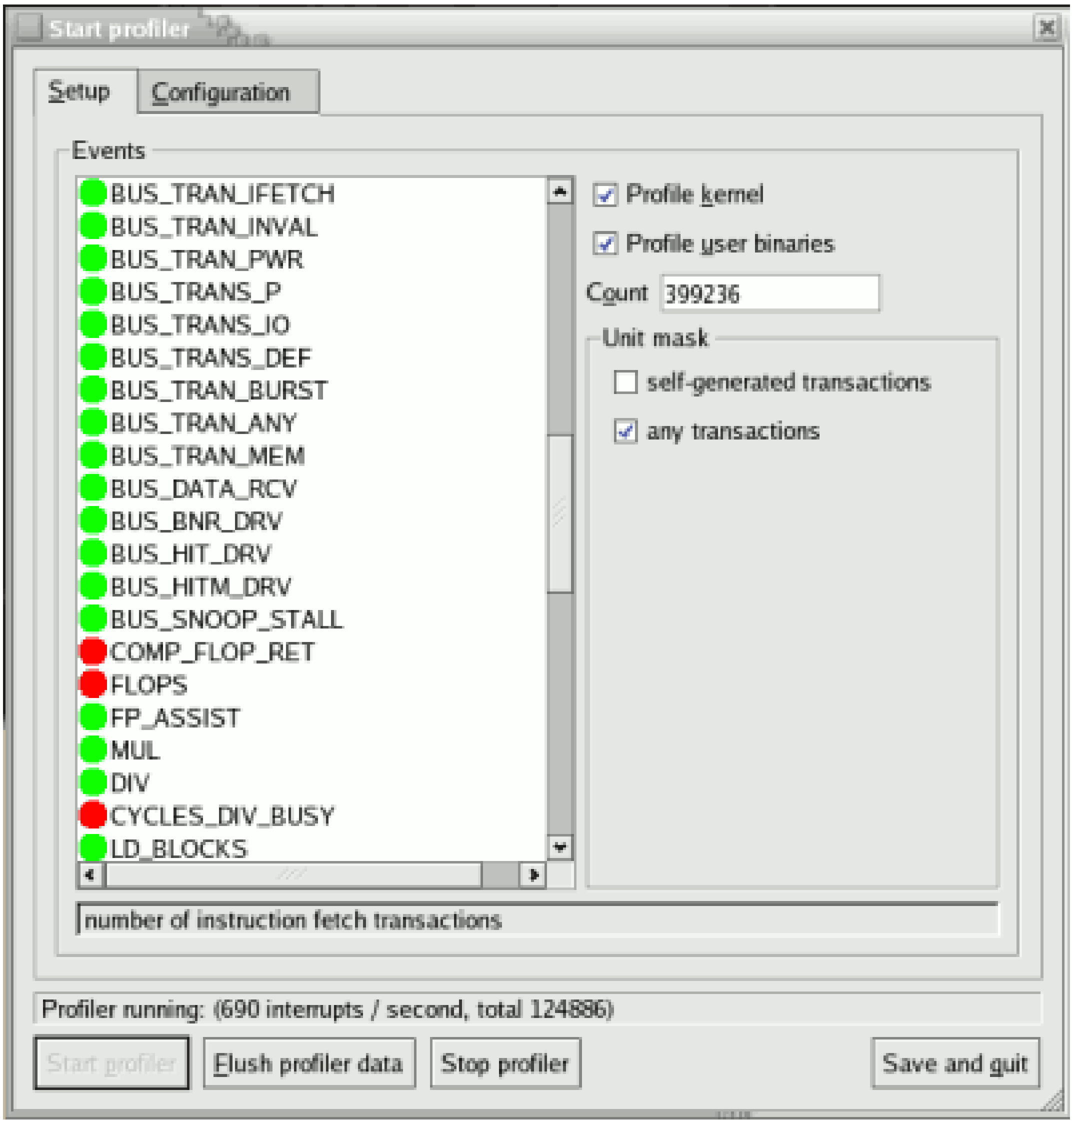

🟢 BUS_TRAN_IFETCH
🟢 BUS_TRAN_INVAL
🟢 BUS_TRAN_PWR
🟢 BUS_TRANS_P
🟢 BUS_TRANS_IO
🟢 BUS_TRANS_DEF
🟢 BUS_TRAN_BURST
🟢 BUS_TRAN_ANY
🟢 BUS_TRAN_MEM
🟢 BUS_DATA_RCV
🟢 BUS_BNR_DRV
🟢 BUS_HIT_DRV
🟢 BUS_HITM_DRV
🟢 BUS_SNOOP_STALL
🔴 COMP_FLOP_RET
🔴 FLOPS
🟢 FP_ASSIST
🟢 MUL
🟢 DIV
🔴 CYCLES_DIV_BUSY
🟢 LD_BLOCKS

☑ Profile kernel

☑ Profile user binaries

Count 399236

☐ self-generated transactions

☑ any transactions

number of instruction fetch transactions

Profiler running: (690 interrupts / second, total 124886)

그림 41.1. OProfile 설정

On the right side of the tab, select the Profile kernel option to count events in kernel mode for the currently selected event, as discussed in 41.2.3절. "커널과 사용자 영역 프로파일 분리하기". If this option is unselected, no samples are collected for the kernel.

Select the Profile user binaries option to count events in user mode for the currently selected event, as discussed in 41.2.3절. “커널과 사용자 영역 프로파일 분리하기” . If this option is unselected, no samples are collected for user applications.

Use the Count text field to set the sampling rate for the currently selected event as discussed in 41.2.2.1절. “샘플 수집 속도” .

If any unit masks are available for the currently selected event, as discussed in 41.2.2.2절. “유닛 마스크 (Unit Masks)” , they are displayed in the Unit Masks area on the right side of the Setup tab. Select the checkbox beside the unit mask to enable it for the event.

On the Configuration tab, to profile the kernel, enter the name and location of the vmlinux file for the kernel to monitor in the Kernel image file text field. To configure OProfile not to monitor the kernel, select No kernel image.

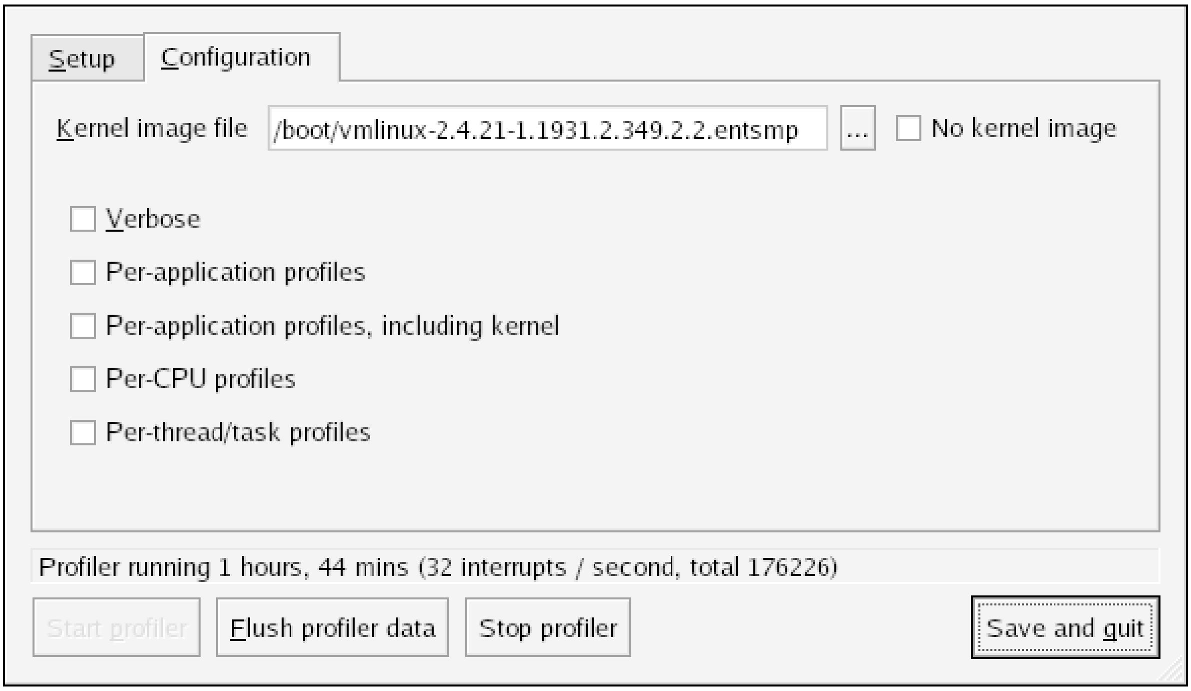

Kernel image file /boot/vmlinux-2.4.21-1.1931.2.349.2.2.entsmp ... □ No kernel image

□ Verbose

□ Per-application profiles

□ Per-application profiles, including kernel

□ Per-CPU profiles

□ Per-thread/task profiles

Profiler running 1 hours, 44 mins (32 interrupts / second, total 176226)

그림 41.2. OProfile Configuration

If the Verbose option is selected, the oprofiled daemon log includes more information.

If Per-application kernel samples files is selected, OProfile generates per-application profiles for the kernel and kernel modules as discussed in 41.2.3절. “커널과 사용자 영역 프로파일 분리하기” . This is equivalent to the opcontrol --separate=kernel command. If Per-application shared libs samples files is selected, OProfile generates per-application profiles for libraries. This is equivalent to the opcontrol --separate=library command.

To force data to be written to samples files as discussed in 41.5절. “데이터 분석” , click the Flush profiler data button. This is equivalent to the opcontrol --dump command.

To start OProfile from the graphical interface, click Start profiler. To stop the profiler, click Stop profiler. Exiting the application does not stop OProfile from sampling.

41.9. 추가 자료

이 장에서는 OProfile의 설정 방법과 사용법을 중점적으로 설명하고 있습니다. 보다 많은 정보를 원하신다면, 다음과 같은 자료를 참조하시기 바랍니다.

41.9.1. 설치된 문서 자료

- /usr/share/doc/oprofile-<version>/oprofile.html — OProfile Manual

- oprofile man page — Discusses opcontrol, opreport, opannotate, and ophelp

41.9.2. 유용한 웹사이트

- http://oprofile.sourceforge.net/ — Contains the latest documentation, mailing lists, IRC channels, and more.

부 VI. 커널 및 드라이버 설정

시스템 관리자는 커널에 대해 알고 이를 사용자 정의할 수 있습니다. Red Hat Enterprise Linux에는 시스템 관리자의 사용자 설정을 보조하기 위한 도구가 포함되어 있습니다.

Manually Upgrading the Kernel

The Red Hat Enterprise Linux kernel is custom built by the Red Hat Enterprise Linux kernel team to ensure its integrity and compatibility with supported hardware. Before Red Hat releases a kernel, it must first pass a rigorous set of quality assurance tests.

Red Hat Enterprise Linux kernels are packaged in RPM format so that they are easy to upgrade and verify using the Package Management Tool, or the yum command. The Package Management Tool automatically queries the Red Hat Enterprise Linux servers and determines which packages need to be updated on your machine, including the kernel. This chapter is only useful for those individuals that require manual updating of kernel packages, without using the yum command.

> **경고**
>
> Building a custom kernel is not supported by the Red Hat Global Services Support team, and therefore is not explored in this manual.

> **Tip**
>
> The use of yum is highly recommended by Red Hat for installing upgraded kernels.

For more information on Red Hat Network, the Package Management Tool, and yum, refer to 14장. Product Subscriptions and Entitlements.

42.1. 커널 패키지 개요

Red Hat Enterprise Linux contains the following kernel packages (some may not apply to your architecture):

- kernel — Contains the kernel for multi-processor systems. For x86 system, only the first 4GB of RAM is used. As such, x86 systems with over 4GB of RAM should use the kernel-PAE.
- kernel-devel — Contains the kernel headers and makefiles sufficient to build modules against the kernel package.
- kernel-PAE (only for i686 systems) — This package offers the following key configuration option (in addition to the options already enabled for the kernel package):
  - PAE (Physical Address Extension) support for systems with more than 4GB of RAM, and reliably up to 16GB.

중요

Physical Address Extension allows x86 processors to address up to 64GB of physical RAM, but due to differences between the Red Hat Enterprise Linux 4 and 5 kernels, only Red Hat Enterprise Linux 4 (with the kernel-hugemem package) is able to reliably address all 64GB of memory. Additionally, the Red Hat Enterprise Linux 5 PAE variant does not allow 4GB of addressable memory per-process like the Red Hat Enterprise Linux 4 kernel-hugemem variant does. However, the x86_64 kernel does not suffer from any of these limitations, and is the suggested Red Hat Enterprise Linux 5 architecture to use with large-memory systems.

- kernel-PAE-devel — Contains the kernel headers and makefiles required to build modules against the kernel-PAE package.
- kernel-doc — Contains documentation files from the kernel source. Various portions of the Linux kernel and the device drivers shipped with it are documented in these files. Installation of this package provides a reference to the options that can be passed to Linux kernel modules at load time.

  By default, these files are placed in the /usr/share/doc/kernel-doc-<version>/ directory.
- kernel-headers — Includes the C header files that specify the interface between the Linux kernel and userspace libraries and programs. The header files define structures and constants that are needed for building most standard programs.
- kernel-xen — Includes a version of the Linux kernel which is needed to run Virtualization.
- kernel-xen-devel — Contains the kernel headers and makefiles required to build modules against the kernel-xen package

Note

The kernel-source package has been removed and replaced with an RPM that can only be retrieved from Red Hat Network. This *.src.rpm package must then be rebuilt locally using the rpmbuild command. For more information on obtaining and installing the kernel source package, refer to the latest updated Release Notes (including all updates) at http://www.redhat.com/docs/manuals/enterprise/

42.2. 업그레이드 준비

Before upgrading the kernel, it is recommended that you take some precautionary steps. The first step is to make sure working boot media exists for the system in case a problem occurs. If the boot loader is not configured properly to boot the new kernel, the system cannot be booted into Red Hat Enterprise Linux without working boot media.

To create a boot diskette, login as root, and run the command /sbin/mkbootdisk `uname -r` at a shell prompt.

Tip

Refer to the mkbootdisk man page for more options. You can create bootable media via CD-Rs, CD-RWs, and USB flash drives, provided that your system BIOS also supports it.

Reboot the machine with the boot media and verify that it works before continuing.

To determine which kernel packages are installed, execute the command rpm -qa | grep kernel at a shell prompt:

The output contains some or all of the following packages, depending on the system's architecture (the version numbers and packages may differ):

```
kernel-2.6.9-5.EL
kernel-devel-2.6.9-5.EL
kernel-utils-2.6.9-5.EL
kernel-doc-2.6.9-5.EL
kernel-smp-2.6.9-5.EL
kernel-smp-devel-2.6.9-5.EL
kernel-hugemem-devel-2.6.9-5.EL
```

From the output, determine which packages need to be download for the kernel upgrade. For a single processor system, the only required package is the kernel package. Refer to 42.1절. "커널 패키지 개요" for descriptions of the different packages.

In the file name, each kernel package contains the architecture for which the package was built. The format is kernel-<variant>-<version>.<arch>.rpm, where <variant> is one of either PAE, xen, and so forth. The <arch> is one of the following:

- x86_64 for the AMD64 and Intel EM64T architectures
- ia64 for the Intel® Itanium™ architecture
- ppc64 for the IBM® eServer™ pSeries™ architecture
- s390 for the IBM® S/390® architecture
- s390x for the IBM® eServer™ System z® architecture
- i686 for Intel® Pentium® II, Intel® Pentium® III, Intel® Pentium® 4, AMD Athlon®, and AMD Duron® systems

42.3. 업그레이드된 커널 다운로드 받기

There are several ways to determine if an updated kernel is available for the system.

- Security Errata — Refer to http://www.redhat.com/security/updates/ for information on security errata, including kernel upgrades that fix security issues.
- Via Red Hat Network — Download and install the kernel RPM packages. Red Hat Network can download the latest kernel, upgrade the kernel on the system, create an initial RAM disk image

if needed, and configure the boot loader to boot the new kernel. For more information, refer to http://www.redhat.com/docs/manuals/RHNetwork/[1].

If Red Hat Network was used to download and install the updated kernel, follow the instructions in 42.5절. "초기 RAM 디스크 이미지 확인하기" and 42.6절. "부트로더 확인하기", only do not change the kernel to boot by default. Red Hat Network automatically changes the default kernel to the latest version. To install the kernel manually, continue to 42.4절. "업그레이드 수행하기".

42.4. 업그레이드 수행하기

After retrieving all of the necessary packages, it is time to upgrade the existing kernel.

중요

It is strongly recommended that you keep the old kernel in case there are problems with the new kernel.

At a shell prompt, change to the directory that contains the kernel RPM packages. Use -i argument with the rpm command to keep the old kernel. Do not use the -U option, since it overwrites the currently installed kernel, which creates boot loader problems. For example:

```
rpm -ivh kernel-<kernel version>.<arch>.rpm
```

The next step is to verify that the initial RAM disk image has been created. Refer to 42.5절. "초기 RAM 디스크 이미지 확인하기" for details.

42.5. 초기 RAM 디스크 이미지 확인하기

If the system uses the ext3 file system, a SCSI controller, or labels to reference partitions in /etc/fstab, an initial RAM disk is needed. The initial RAM disk allows a modular kernel to have access to modules that it might need to boot from before the kernel has access to the device where the modules normally reside.

On architectures other than IBM eServer iSeries, the initial RAM disk can be created with the mkinitrd command. However, this step is performed automatically if the kernel and its associated packages are installed or upgraded from the RPM packages distributed by Red Hat; in such cases, you do not need to create the initial RAM disk manually. To verify that an initial RAM disk already exists, use the command ls -l /boot to make sure the initrd-<version>.img file was created (the version should match the version of the kernel just installed).

On iSeries systems, the initial RAM disk file and vmlinux file are combined into one file, which is created with the addRamDisk command. This step is performed automatically if the kernel and its associated packages are installed or upgraded from the RPM packages distributed by Red Hat, Inc.; thus, it does not need to be executed manually. To verify that it was created, use the command ls -l /boot to make sure the /boot/vmlinitrd-<kernel-version> file already exists (the <kernel-version> should match the version of the kernel just installed).

The next step is to verify that the boot loader has been configured to boot the new kernel. Refer to 42.6절. "부트로더 확인하기" for details.

[1] http://www.redhat.com/docs/manuals/RHNetwork/

42.6. 부트로더 확인하기

The kernel RPM package configures the boot loader to boot the newly installed kernel (except for IBM eServer iSeries systems). However, it does not configure the boot loader to boot the new kernel by default.

It is always a good idea to confirm that the boot loader has been configured correctly. This is a crucial step. If the boot loader is configured incorrectly, the system will not boot into Red Hat Enterprise Linux properly. If this happens, boot the system with the boot media created earlier and try configuring the boot loader again.

42.6.1. x86 시스템

All x86 systems (including all AMD64 systems) use GRUB as the boot loader.

42.6.1.1. GRUB

Confirm that the file /boot/grub/grub.conf contains a title section with the same version as the kernel package just installed

```
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/hda2
#          initrd /initrd-version.img
#boot=/dev/hda
default=1 timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Enterprise Linux (2.6.9-5.EL)
         root (hd0,0)
  kernel /vmlinuz-2.6.9-5.EL ro root=LABEL=/
  initrd /initrd-2.6.9-5.EL.img
title Red Hat Enterprise Linux (2.6.9-1.906_EL)
         root (hd0,0)
  kernel /vmlinuz-2.6.9-1.906_EL ro root=LABEL=/
  initrd /initrd-2.6.9-1.906_EL.img
```

If a separate /boot/ partition was created, the paths to the kernel and initrd image are relative to /boot/.

Notice that the default is not set to the new kernel. To configure GRUB to boot the new kernel by default, change the value of the default variable to the title section number for the title section that contains the new kernel. The count starts with 0. For example, if the new kernel is the first title section, set default to 0.

새 커널을 테스트하기 위하여 컴퓨터를 재부팅한 후 하드웨어가 제대로 검색되는지 확인하기 위하여 메시지들을 살펴보십시오.

42.6.2. Itanium 시스템

Itanium systems use ELILO as the boot loader, which uses /boot/efi/EFI/redhat/elilo.conf as the configuration file. Confirm that this file contains an image section with the same version as the kernel package just installed:

```
prompt timeout=50 default=old  image=vmlinuz-2.6.9-5.EL
         label=linux
  initrd=initrd-2.6.9-5.EL.img          read-only
```

```
    append="root=LABEL=/" image=vmlinuz-2.6.9-1.906_EL
    label=old
    initrd=initrd-2.6.9-1.906.img              read-only
    append="root=LABEL=/"
```

Notice that the default is not set to the new kernel. To configure ELILO to boot the new kernel, change the value of the default variable to the value of the label for the image section that contains the new kernel.

새 커널을 테스트하기 위하여 컴퓨터를 재부팅한 후 하드웨어가 제대로 검색되는지 확인하기 위하여 메시지들을 살펴보십시오.

42.6.3. IBM S/390 and IBM System z Systems

The IBM S/390 and IBM System z systems use z/IPL as the boot loader, which uses /etc/zipl.conf as the configuration file. Confirm that the file contains a section with the same version as the kernel package just installed:

```
[defaultboot] default=old target=/boot/
[linux]
          image=/boot/vmlinuz-2.6.9-5.EL
  ramdisk=/boot/initrd-2.6.9-5.EL.img
  parameters="root=LABEL=/"
[old]
          image=/boot/vmlinuz-2.6.9-1.906_EL
  ramdisk=/boot/initrd-2.6.9-1.906_EL.img
  parameters="root=LABEL=/"
```

Notice that the default is not set to the new kernel. To configure z/IPL to boot the new kernel by default, change the value of the default variable to the name of the section that contains the new kernel. The first line of each section contains the name in brackets.

After modifying the configuration file, run /sbin/zipl as root to enable the changes.

새 커널을 테스트하기 위하여 컴퓨터를 재부팅한 후 하드웨어가 제대로 검색되는지 확인하기 위하여 메시지들을 살펴보십시오.

42.6.4. IBM eServer iSeries 시스템

The /boot/vmlinitrd-<kernel-version> file is installed when you upgrade the kernel. However, you must use the dd command to configure the system to boot the new kernel:

1. As root, issue the command cat /proc/iSeries/mf/side to determine the default side (either A, B, or C).

2. As root, issue the following command, where <kernel-version> is the version of the new kernel and <side> is the side from the previous command:

   ```
   dd if=/boot/vmlinitrd-<kernel-version> of=/proc/iSeries/mf/<side>/vmlinux bs=8k
   ```

새 커널을 테스트하기 위하여 컴퓨터를 재부팅한 후 하드웨어가 제대로 검색되는지 확인하기 위하여 메시지들을 살펴보십시오.

42.6.5. IBM eServer pSeries 시스템

IBM eServer pSeries systems use YABOOT as the boot loader, which uses /etc/aboot.conf as the configuration file. Confirm that the file contains an image section with the same version as the kernel package just installed:

```
boot=/dev/sda1 init-message=Welcome to Red Hat Enterprise Linux! Hit <TAB> for boot options
partition=2 timeout=30 install=/usr/lib/yaboot/yaboot delay=10 nonvram
image=/vmlinux--2.6.9-5.EL
          label=old
  read-only
  initrd=/initrd--2.6.9-5.EL.img
  append="root=LABEL=/"
image=/vmlinux-2.6.9-5.EL
  label=linux
  read-only
  initrd=/initrd-2.6.9-5.EL.img
  append="root=LABEL=/"
```

Notice that the default is not set to the new kernel. The kernel in the first image is booted by default. To change the default kernel to boot either move its image stanza so that it is the first one listed or add the directive default and set it to the label of the image stanza that contains the new kernel.

새 커널을 테스트하기 위하여 컴퓨터를 재부팅한 후 하드웨어가 제대로 검색되는지 확인하기 위하여 메시지들을 살펴보십시오.

General Parameters and Modules

This chapter is provided to illustrate some of the possible parameters available for common hardware device drivers [1], which under Red Hat Enterprise Linux are called kernel modules. In most cases, the default parameters do work. However, there may be times when extra module parameters are necessary for a device to function properly or to override the module's default parameters for the device.

During installation, Red Hat Enterprise Linux uses a limited subset of device drivers to create a stable installation environment. Although the installation program supports installation on many different types of hardware, some drivers (including those for SCSI adapters and network adapters) are not included in the installation kernel. Rather, they must be loaded as modules by the user at boot time.

Once installation is completed, support exists for a large number of devices through kernel modules.

Important

Red Hat provides a large number of unsupported device drivers in groups of packages called kernel-smp-unsupported-<kernel-version> and kernel-hugemem-unsupported-<kernel-version> . Replace <kernel-version> with the version of the kernel installed on the system. These packages are not installed by the Red Hat Enterprise Linux installation program, and the modules provided are not supported by Red Hat, Inc.

43.1. Kernel Module Utilities

A group of commands for managing kernel modules is available if the module-init-tools package is installed. Use these commands to determine if a module has been loaded successfully or when trying different modules for a piece of new hardware.

The command /sbin/lsmod displays a list of currently loaded modules. For example:

```
Module                  Size  Used by
tun                    11585  1
autofs4                21573  1
hidp                   16193  2
rfcomm                 37849  0
l2cap                  23873  10 hidp,rfcomm
bluetooth              50085  5 hidp,rfcomm,l2cap
sunrpc                153725  1
dm_mirror              29073  0
dm_mod                 57433  1 dm_mirror
video                  17221  0
sbs                    16257  0
i2c_ec                  5569  1 sbs
container               4801  0
button                  7249  0
battery                10565  0
asus_acpi              16857  0
ac                      5701  0
ipv6                  246113  12
```

[1] A driver is software which enables Linux to use a particular hardware device. Without a driver, the kernel cannot communicate with attached devices.

```
lp                    13065  0
parport_pc            27493  1
parport               37001  2 lp,parport_pc
uhci_hcd              23885  0
floppy                57317  1
sg                    34653  0
snd_ens1371           26721  1
gameport              16073  1 snd_ens1371
snd_rawmidi           24897  1 snd_ens1371
snd_ac97_codec        91360  1 snd_ens1371
snd_ac97_bus           2753  1 snd_ac97_codec
snd_seq_dummy          4293  0
snd_seq_oss           32705  0
serio_raw              7493  0
snd_seq_midi_event     8001  1 snd_seq_oss
snd_seq               51633  5 snd_seq_dummy,snd_seq_oss,snd_seq_midi_event
snd_seq_device         8781  4 snd_rawmidi,snd_seq_dummy,snd_seq_oss,snd_seq
snd_pcm_oss           42849  0
snd_mixer_oss         16833  1 snd_pcm_oss
snd_pcm               76485  3 snd_ens1371,snd_ac97_codec,snd_pcm_oss
snd_timer             23237  2 snd_seq,snd_pcm
snd                   52933  12
 snd_ens1371,snd_rawmidi,snd_ac97_codec,snd_seq_oss,snd_seq,snd_seq_device,snd_pcm_oss,snd_mixer_oss,snd_pcm,snd_timer
soundcore             10145  1 snd
i2c_piix4              8909  0
ide_cd                38625  3
snd_page_alloc        10569  1 snd_pcm
i2c_core              21697  2 i2c_ec,i2c_piix4
pcnet32               34117  0
cdrom                 34913  1 ide_cd
mii                    5825  1 pcnet32
pcspkr                 3521  0
ext3                 129737  2
jbd                   58473  1 ext3
mptspi                17353  3
scsi_transport_spi    25025  1 mptspi
mptscsih              23361  1 mptspi
sd_mod                20929  16
scsi_mod             134121  5 sg,mptspi,scsi_transport_spi,mptscsih,sd_mod
mptbase               52193  2 mptspi,mptscsih
```

For each line, the first column is the name of the module, the second column is the size of the module, and the third column is the use count.

The /sbin/lsmod output is less verbose and easier to read than the output from viewing /proc/modules.

To load a kernel module, use the /sbin/modprobe command followed by the kernel module name. By default, modprobe attempts to load the module from the /lib/modules/<kernel-version>/kernel/drivers/ subdirectories. There is a subdirectory for each type of module, such as the net/ subdirectory for network interface drivers. Some kernel modules have module dependencies, meaning that other modules must be loaded first for it to load. The /sbin/modprobe command checks for these dependencies and loads the module dependencies before loading the specified module.

For example, the command

```
modprobe e100
```

loads any module dependencies and then the e100 module.

To print to the screen all commands as /sbin/modprobe executes them, use the -v option. For example:

```
modprobe -v e100
```

Output similar to the following is displayed:

```
insmod /lib/modules/2.6.9-5.EL/kernel/drivers/net/e100.ko
Using /lib/modules/2.6.9-5.EL/kernel/drivers/net/e100.ko
Symbol version prefix 'smp_'
```

The /sbin/insmod command also exists to load kernel modules; however, it does not resolve dependencies. Thus, it is recommended that the /sbin/modprobe command be used.

To unload kernel modules, use the /sbin/rmmod command followed by the module name. The rmmod utility only unloads modules that are not in use and that are not a dependency of other modules in use.

For example, the command

```
rmmod e100
```

unloads the e100 kernel module.

Another useful kernel module utility is modinfo. Use the command /sbin/modinfo to display information about a kernel module. The general syntax is:

```
modinfo [options] <module>
```

Options include -d, which displays a brief description of the module, and -p, which lists the parameters the module supports. For a complete list of options, refer to the modinfo man page (man modinfo).

43.2. Persistent Module Loading

Kernel modules are usually loaded directly by the facility that requires them, which is given correct settings in the /etc/modprobe.conf file. However, it is sometimes necessary to explicitly force the loading of a module at boot time.

Red Hat Enterprise Linux checks for the existence of the /etc/rc.modules file at boot time, which contains various commands to load modules. The rc.modules should be used, and not rc.local because rc.modules is executed earlier in the boot process.

For example, the following commands configure loading of the foo module at boot time (as root):

```
echo modprobe foo >> /etc/rc.modules
chmod +x /etc/rc.modules
```

Tip

This approach is not necessary for network and SCSI interfaces because they have their own specific mechanisms.

43.3. Specifying Module Parameters

In some situations, it may be necessary to supply parameters to a module as it is loaded for it to function properly.

For instance, to enable full duplex at 100Mbps connection speed for an Intel Ether Express/100 card, load the e100 driver with the e100_speed_duplex=4 option.

Caution

When a parameter has commas, be sure not to put a space after a comma.

Tip

The modinfo command is also useful for listing various information about a kernel module, such as version, dependencies, parameter options, and aliases.

43.4. Storage parameters

표 43.1. Storage Module Parameters

Hardware	Module	Parameters
3ware Storage Controller and 9000 series	3w-xxxx.ko, 3w-9xxx.ko	
Adaptec Advanced Raid Products, Dell PERC2, 2/Si, 3/Si, 3/Di, HP NetRAID-4M, IBM ServeRAID, and ICP SCSI driver	aacraid.ko	nondasd — Control scanning of hba for nondasd devices. 0=off, 1=on dacmode — Control whether dma addressing is using 64 bit DAC. 0=off, 1=on commit — Control whether a COMMIT_CONFIG is issued to the adapter for foreign arrays. This is typically needed in

Hardware	Module	Parameters
		systems that do not have a BIOS. 0=off, 1=on startup_timeout — The duration of time in seconds to wait for adapter to have it's kernel up and running. This is typically adjusted for large systems that do not have a BIOS aif_timeout — The duration of time in seconds to wait for applications to pick up AIFs before deregistering them. This is typically adjusted for heavily burdened systems. numacb — Request a limit to the number of adapter control blocks (FIB) allocated. Valid values are 512 and down. Default is to use suggestion from Firmware. acbsize — Request a specific adapter control block (FIB) size. Valid values are 512, 2048, 4096 and 8192. Default is to use suggestion from Firmware.
Adaptec 28xx, R9xx, 39xx AHA-284x, AHA-29xx, AHA-394x, AHA-398x, AHA-274x, AHA-274xT, AHA-2842, AHA-2910B, AHA-2920C, AHA-2930/U/ U2, AHA-2940/W/U/UW/AU/, U2W/U2/U2B/, U2BOEM, AHA-2944D/WD/UD/UWD, AHA-2950U2/W/B, AHA-3940/ U/W/UW/, AUW/U2W/U2B, AHA-3950U2D, AHA-3985/U/ W/UW, AIC-777x, AIC-785x, AIC-786x, AIC-787x, AIC-788x , AIC-789x, AIC-3860	aic7xxx.ko	verbose — Enable verbose/ diagnostic logging allow_memio — Allow device registers to be memory mapped debug — Bitmask of debug values to enable no_probe — Toggle EISA/VLB controller probing probe_eisa_vl — Toggle EISA/ VLB controller probing no_reset — Supress initial bus resets extended — Enable extended geometry on all controllers periodic_otag — Send an ordered tagged transaction periodically to prevent tag

Hardware	Module	Parameters
		starvation. This may be required by some older disk drives or RAID arrays. tag_info:<tag_str> — Set per-target tag depth global_tag_depth:<int> — Global tag depth for every target on every bus seltime:<int> — Selection Timeout (0/256ms,1/128ms,2/64ms,3/32ms)
IBM ServeRAID	ips.ko	
LSI Logic MegaRAID Mailbox Driver	megaraid_mbox.ko	unconf_disks — Set to expose unconfigured disks to kernel (default=0) busy_wait — Max wait for mailbox in microseconds if busy (default=10) max_sectors — Maximum number of sectors per IO command (default=128) cmd_per_lun — Maximum number of commands per logical unit (default=64) fast_load — Faster loading of the driver, skips physical devices! (default=0) debug_level — Debug level for driver (default=0)
Emulex LightPulse Fibre Channel SCSI driver	lpfc.ko	lpfc_poll — FCP ring polling mode control: 0 - none, 1 - poll with interrupts enabled 3 - poll and disable FCP ring interrupts lpfc_log_verbose — Verbose logging bit-mask lpfc_lun_queue_depth — Max number of FCP commands we can queue to a specific LUN lpfc_hba_queue_depth — Max number of FCP commands we can queue to a lpfc HBA

Hardware	Module	Parameters
		lpfc_scan_down — Start scanning for devices from highest ALPA to lowest lpfc_nodev_tmo — Seconds driver will hold I/O waiting for a device to come back lpfc_topology — Select Fibre Channel topology lpfc_link_speed — Select link speed lpfc_fcp_class — Select Fibre Channel class of service for FCP sequences lpfc_use_adisc — Use ADISC on rediscovery to authenticate FCP devices lpfc_ack0 — Enable ACK0 support lpfc_cr_delay — A count of milliseconds after which an interrupt response is generated lpfc_cr_count — A count of I/O completions after which an interrupt response is generated lpfc_multi_ring_support — Determines number of primary SLI rings to spread IOCB entries across lpfc_fdmi_on — Enable FDMI support lpfc_discovery_threads — Maximum number of ELS commands during discovery lpfc_max_luns — Maximum allowed LUN lpfc_poll_tmo — Milliseconds driver will wait between polling FCP ring
HP Smart Array	cciss.ko	

Hardware	Module	Parameters
LSI Logic MPT Fusion	mptbase.ko mptctl.ko mptfc.ko mptlan.ko mptsas.ko mptscsih.ko mptspi.ko	mpt_msi_enable — MSI Support Enable mptfc_dev_loss_tmo — Initial time the driver programs the transport to wait for an rport to return following a device loss event. mpt_pt_clear — Clear persistency table mpt_saf_te — Force enabling SEP Processor
QLogic Fibre Channel Driver	qla2xxx.ko	ql2xlogintimeout — Login timeout value in seconds. qlport_down_retry — Maximum number of command retries to a port that returns a PORT-DOWN status ql2xplogiabsentdevice — Option to enable PLOGI to devices that are not present after a Fabric scan. ql2xloginretrycount — Specify an alternate value for the NVRAM login retry count. ql2xallocfwdump — Option to enable allocation of memory for a firmware dump during HBA initialization. Default is 1 - allocate memory. extended_error_logging — Option to enable extended error logging. ql2xfdmienable — Enables FDMI registrations.
NCR, Symbios and LSI 8xx and 1010	sym53c8xx	cmd_per_lun — The maximum number of tags to use by default tag_ctrl — More detailed control over tags per LUN burst — Maximum burst. 0 to disable, 255 to read from registers

Hardware	Module	Parameters
		led — Set to 1 to enable LED support diff — 0 for no differential mode, 1 for BIOS, 2 for always, 3 for not GPIO3 irqm — 0 for open drain, 1 to leave alone, 2 for totem pole buschk — 0 to not check, 1 for detach on error, 2 for warn on error hostid — The SCSI ID to use for the host adapters verb — 0 for minimal verbosity, 1 for normal, 2 for excessive debug — Set bits to enable debugging settle — Settle delay in seconds. Default 3 nvram — Option currently not used excl — List ioport addresses here to prevent controllers from being attached safe — Set other settings to a "safe mode"

43.5. Ethernet Parameters

Important

Most modern Ethernet-based network interface cards (NICs), do not require module parameters to alter settings. Instead, they can be configured using ethtool or mii-tool. Only after these tools fail to work should module parameters be adjusted. Module parameters can be viewed using the modinfo command.

Note

For information about using these tools, consult the man pages for ethtool, mii-tool, and modinfo.

표 43.2. Ethernet Module Parameters

Hardware	Module	Parameters
3Com EtherLink PCI III/XL Vortex (3c590, 3c592, 3c595, 3c597) Boomerang (3c900, 3c905, 3c595)	3c59x.ko	debug — 3c59x debug level (0-6) options — 3c59x: Bits 0-3: media type, bit 4: bus mastering, bit 9: full duplex global_options — 3c59x: same as options, but applies to all NICs if options is unset full_duplex — 3c59x full duplex setting(s) (1) global_full_duplex — 3c59x: same as full_duplex, but applies to all NICs if full_duplex is unset hw_checksums — 3c59x Hardware checksum checking by adapter(s) (0-1) flow_ctrl — 3c59x 802.3x flow control usage (PAUSE only) (0-1) enable_wol — 3c59x: Turn on Wake-on-LAN for adapter(s) (0-1) global_enable_wol — 3c59x: same as enable_wol, but applies to all NICs if enable_wol is unset rx_copybreak — 3c59x copy breakpoint for copy-only-tiny-frames max_interrupt_work — 3c59x maximum events handled per interrupt

Hardware	Module	Parameters
		compaq_ioaddr — 3c59x PCI I/O base address (Compaq BIOS problem workaround) compaq_irq — 3c59x PCI IRQ number (Compaq BIOS problem workaround) compaq_device_id — 3c59x PCI device ID (Compaq BIOS problem workaround) watchdog — 3c59x transmit timeout in milliseconds global_use_mmio — 3c59x: same as use_mmio, but applies to all NICs if options is unset use_mmio — 3c59x: use memory-mapped PCI I/O resource (0-1)
RTL8139, SMC EZ Card Fast Ethernet, RealTek cards using RTL8129, or RTL8139 Fast Ethernet chipsets	8139too.ko	
Broadcom 4400 10/100 PCI ethernet driver	b44.ko	b44_debug — B44 bitmapped debugging message enable value
Broadcom NetXtreme II BCM5706/5708 Driver	bnx2.ko	disable_msi — Disable Message Signaled Interrupt (MSI)
Intel Ether Express/100 driver	e100.ko	debug — Debug level (0=none,...,16=all) eeprom_bad_csum_allow — Allow bad eeprom checksums
Intel EtherExpress/1000 Gigabit	e1000.ko	TxDescriptors — Number of transmit descriptors RxDescriptors — Number of receive descriptors Speed — Speed setting Duplex — Duplex setting AutoNeg — Advertised auto-negotiation setting FlowControl — Flow Control setting XsumRX — Disable or enable Receive Checksum offload

Hardware	Module	Parameters
		TxIntDelay — Transmit Interrupt Delay TxAbsIntDelay — Transmit Absolute Interrupt Delay RxIntDelay — Receive Interrupt Delay RxAbsIntDelay — Receive Absolute Interrupt Delay InterruptThrottleRate — Interrupt Throttling Rate SmartPowerDownEnable — Enable PHY smart power down KumeranLockLoss — Enable Kumeran lock loss workaround
Myricom 10G driver (10GbE)	myri10ge.ko	myri10ge_fw_name — Firmware image name myri10ge_ecrc_enable — Enable Extended CRC on PCI-E myri10ge_max_intr_slots — Interrupt queue slots myri10ge_small_bytes — Threshold of small packets myri10ge_msi — Enable Message Signalled Interrupts myri10ge_intr_coal_delay — Interrupt coalescing delay myri10ge_flow_control — Pause parameter myri10ge_deassert_wait — Wait when deasserting legacy interrupts myri10ge_force_firmware — Force firmware to assume aligned completions myri10ge_skb_cross_4k — Can a small skb cross a 4KB boundary? myri10ge_initial_mtu — Initial MTU

Hardware	Module	Parameters
		myri10ge_napi_weight — Set NAPI weight myri10ge_watchdog_timeout — Set watchdog timeout myri10ge_max_irq_loops — Set stuck legacy IRQ detection threshold
NatSemi DP83815 Fast Ethernet	natsemi.ko	mtu — DP8381x MTU (all boards) debug — DP8381x default debug level rx_copybreak — DP8381x copy breakpoint for copy-only-tiny-frames options — DP8381x: Bits 0-3: media type, bit 17: full duplex full_duplex — DP8381x full duplex setting(s) (1)
AMD PCnet32 and AMD PCnetPCI	pcnet32.ko	
PCnet32 and PCnetPCI	pcnet32.ko	debug — pcnet32 debug level max_interrupt_work — pcnet32 maximum events handled per interrupt rx_copybreak — pcnet32 copy breakpoint for copy-only-tiny-frames tx_start_pt — pcnet32 transmit start point (0-3) pcnet32vlb — pcnet32 Vesa local bus (VLB) support (0/1) options — pcnet32 initial option setting(s) (0-15) full_duplex — pcnet32 full duplex setting(s) (1) homepna — pcnet32 mode for 79C978 cards (1 for HomePNA, 0 for Ethernet, default Ethernet
RealTek RTL-8169 Gigabit Ethernet driver	r8169.ko	media — force phy operation. Deprecated by ethtool (8).

Hardware	Module	Parameters
		rx_copybreak — Copy breakpoint for copy-only-tiny-frames use_dac — Enable PCI DAC. Unsafe on 32 bit PCI slot. debug — Debug verbosity level (0=none, ..., 16=all)
Neterion Xframe 10GbE Server Adapter	s2io.ko	
SIS 900/701G PCI Fast Ethernet	sis900.ko	multicast_filter_limit — SiS 900/7016 maximum number of filtered multicast addresses max_interrupt_work — SiS 900/7016 maximum events handled per interrupt sis900_debug — SiS 900/7016 bitmapped debugging message level
Adaptec Starfire Ethernet driver	starfire.ko	max_interrupt_work — Maximum events handled per interrupt mtu — MTU (all boards) debug — Debug level (0-6) rx_copybreak — Copy breakpoint for copy-only-tiny-frames intr_latency — Maximum interrupt latency, in microseconds small_frames — Maximum size of receive frames that bypass interrupt latency (0,64,128,256,512) options — Deprecated: Bits 0-3: media type, bit 17: full duplex full_duplex — Deprecated: Forced full-duplex setting (0/1) enable_hw_cksum — Enable/disable hardware cksum support (0/1)

Hardware	Module	Parameters
Broadcom Tigon3	tg3.ko	tg3_debug — Tigon3 bitmapped debugging message enable value
ThunderLAN PCI	tlan.ko	aui — ThunderLAN use AUI port(s) (0-1) duplex — ThunderLAN duplex setting(s) (0-default, 1-half, 2-full) speed — ThunderLAN port speen setting(s) (0,10,100) debug — ThunderLAN debug mask bbuf — ThunderLAN use big buffer (0-1)
Digital 21x4x Tulip PCI Ethernet cards SMC EtherPower 10 PCI(8432T/8432BT) SMC EtherPower 10/100 PCI(9332DST) DEC EtherWorks 100/10 PCI(DE500-XA) DEC EtherWorks 10 PCI(DE450) DEC QSILVER's, Znyx 312 etherarray Allied Telesis LA100PCI-T Danpex EN-9400, Cogent EM110	tulip.ko	io io_port
VIA Rhine PCI Fast Ethernet cards with either the VIA VT86c100A Rhine-II PCI or 3043 Rhine-I D-Link DFE-930-TX PCI 10/100	via-rhine.ko	max_interrupt_work — VIA Rhine maximum events handled per interrupt debug — VIA Rhine debug level (0-7) rx_copybreak — VIA Rhine copy breakpoint for copy-only-tiny-frames avoid_D3 — Avoid power state D3 (work-around for broken BIOSes)

43.5.1. Using Multiple Ethernet Cards

It is possible to use multiple Ethernet cards on a single machine. For each card there must be an alias and, possibly, options lines for each card in /etc/modprobe.conf.

For additional information about using multiple Ethernet cards, refer to the Linux Ethernet-HOWTO online at http://www.redhat.com/mirrors/LDP/HOWTO/Ethernet-HOWTO.html.

43.5.2. The Channel Bonding Module

Red Hat Enterprise Linux allows administrators to bind NICs together into a single channel using the bonding kernel module and a special network interface, called a channel bonding interface. Channel bonding enables two or more network interfaces to act as one, simultaneously increasing the bandwidth and providing redundancy.

To channel bond multiple network interfaces, the administrator must perform the following steps:

1. Add the following line to /etc/modprobe.conf:

   ```
   alias bond<N> bonding
   ```

   Replace <N> with the interface number, such as 0. For each configured channel bonding interface, there must be a corresponding entry in /etc/modprobe.conf.

2. Configure a channel bonding interface as outlined in 15.2.3절. “채널 결합 인터페이스” .

3. To enhance performance, adjust available module options to ascertain what combination works best. Pay particular attention to the miimon or arp_interval and the arp_ip_target parameters. Refer to 43.5.2.1절. “bonding Module Directives” for a list of available options and how to quickly determine the best ones for your bonded interface.

43.5.2.1. bonding Module Directives

It is a good idea to test which channel bonding module parameters work best for your bonded interfaces before adding them to the BONDING_OPTS="<bonding parameters>" directive in your bonding interface configuration file (ifcfg-bond0 for example). Parameters to bonded interfaces can be configured without unloading (and reloading) the bonding module by manipulating files in the sysfs file system.

sysfs is a virtual file system that represents kernel objects as directories, files and symbolic links. sysfs can be used to query for information about kernel objects, and can also manipulate those objects through the use of normal file system commands. The sysfs virtual file system has a line in /etc/fstab, and is mounted under /sys. All bonded interfaces can be configured dynamically by interacting with and manipulating files under the /sys/class/net/ directory.

After you have created a channel bonding interface file such as ifcfg-bond0 and inserted SLAVE=yes and MASTER=bond0 directives in the bonded interfaces following the instructions in 15.2.3절. “채널 결합 인터페이스” , you can proceed to testing and determining the best parameters for your bonded interface.

First, bring up the bond you created by running ifconfig bond<N> up as root:

```
ifconfig bond0 up
```

If you have correctly created the ifcfg-bond0 bonding interface file, you will be able to see bond0 listed in the output of running ifconfig (without any options):

```
~]# ifconfig
bond0     Link encap:Ethernet  HWaddr 00:00:00:00:00:00
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
```

```
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
eth0      Link encap:Ethernet  HWaddr 52:54:00:26:9E:F1
          inet addr:192.168.122.251  Bcast:192.168.122.255  Mask:255.255.255.0
          inet6 addr: fe80::5054:ff:fe26:9ef1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:207 errors:0 dropped:0 overruns:0 frame:0
          TX packets:205 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:70374 (68.7 KiB)  TX bytes:25298 (24.7 KiB)
[output truncated]
```

To view all existing bonds, even if they are not up, run:

```
~]# cat /sys/class/net/bonding_masters
bond0
```

You can configure each bond individually by manipulating the files located in the /sys/class/net/bond<N>/bonding/ directory. First, the bond you are configuring must be taken down:

```
ifconfig bond0 down
```

As an example, to enable MII monitoring on bond0 with a 1 second interval, you could run (as root):

```
echo 1000 > /sys/class/net/bond0/bonding/miimon
```

To configure bond0 for balance-alb mode, you could run either:

```
echo 6 > /sys/class/net/bond0/bonding/mode
```

...or, using the name of the mode:

```
echo balance-alb > /sys/class/net/bond0/bonding/mode
```

After configuring some options for the bond in question, you can bring it up and test it by running ifconfig bond<N> up . If you decide to change the options, take the interface down, modify its parameters using sysfs, bring it back up, and re-test.

Once you have determined the best set of parameters for your bond, add those parameters as a space-separated list to the BONDING_OPTS= directive of the /etc/sysconfig/network-scripts/ifcfg-bond<N> file for the bonded interface you are configuring. Whenever that bond is brought up (for example, by the system during the boot sequence if the ONBOOT=yes directive is set), the bonding options specified in the BONDING_OPTS will take effect for that bond. For more information on configuring bonded interfaces (and BONDING_OPTS), refer to 15.2.3절. “채널 결합 인터페이스” .

The following is a list of available channel bonding module parameters for the bonding module. For more in-depth information on configuring channel bonding and the exhaustive list of bonding module parameters, install the kernel-doc package and then locating and opening the included bonding.txt file:

```
yum -y install kernel-doc
nano -w $(rpm -ql kernel-doc | grep bonding.txt)
```

Bonding Interface Parameters

arp_interval=<time_in_milliseconds>

Specifies (in milliseconds) how often ARP monitoring occurs.

> **Important**
>
> It is essential that both arp_interval and arp_ip_target parameters are specified, or, alternatively, the miimon parameter is specified. Failure to do so can cause degradation of network performance in the event that a link fails.

If using this setting while in mode=0 or mode=1 (the two load-balancing modes), the network switch must be configured to distribute packets evenly across the NICs. For more information on how to accomplish this, refer to /usr/share/doc/kernel-doc-<kernel_version>/Documentation/networking/bonding.txt

The value is set to 0 by default, which disables it.

arp_ip_target=<ip_address> [,<ip_address_2>,...<ip_address_16>]

Specifies the target IP address of ARP requests when the arp_interval parameter is enabled. Up to 16 IP addresses can be specified in a comma separated list.

arp_validate=<value>

Validate source/distribution of ARP probes; default is none. Other valid values are active, backup, and all.

debug=<number>

Enables debug messages. Possible values are:

- 0 — Debug messages are disabled. This is the default.
- 1 — Debug messages are enabled.

downdelay=<time_in_milliseconds>

Specifies (in milliseconds) how long to wait after link failure before disabling the link. The value must be a multiple of the value specified in the miimon parameter. The value is set to 0 by default, which disables it.

lacp_rate=<value>

Specifies the rate at which link partners should transmit LACPDU packets in 802.3ad mode. Possible values are:

- slow or 0 — Default setting. This specifies that partners should transmit LACPDUs every 30 seconds.
- fast or 1 — Specifies that partners should transmit LACPDUs every 1 second.

miimon=<time_in_milliseconds>

Specifies (in milliseconds) how often MII link monitoring occurs. This is useful if high availability is required because MII is used to verify that the NIC is active. To verify that the driver for a particular NIC supports the MII tool, type the following command as root:

```
ethtool <interface_name> | grep "Link detected:"
```

In this command, replace <interface_name> with the name of the device interface, such as eth0, not the bond interface. If MII is supported, the command returns:

```
Link detected: yes
```

If using a bonded interface for high availability, the module for each NIC must support MII. Setting the value to 0 (the default), turns this feature off. When configuring this setting, a good starting point for this parameter is 100.

Important

It is essential that both arp_interval and arp_ip_target parameters are specified, or, alternatively, the miimon parameter is specified. Failure to do so can cause degradation of network performance in the event that a link fails.

mode=<value>

...where <value> is one of:

- balance-rr or 0 — Sets a round-robin policy for fault tolerance and load balancing. Transmissions are received and sent out sequentially on each bonded slave interface beginning with the first one available.

- active-backup or 1 — Sets an active-backup policy for fault tolerance. Transmissions are received and sent out via the first available bonded slave interface. Another bonded slave interface is only used if the active bonded slave interface fails.

- balance-xor or 2 — Sets an XOR (exclusive-or) policy for fault tolerance and load balancing. Using this method, the interface matches up the incoming request's MAC address with the MAC address for one of the slave NICs. Once this link is established, transmissions are sent out sequentially beginning with the first available interface.

- broadcast or 3 — Sets a broadcast policy for fault tolerance. All transmissions are sent on all slave interfaces.

- 802.3ad or 4 — Sets an IEEE 802.3ad dynamic link aggregation policy. Creates aggregation groups that share the same speed and duplex settings. Transmits and receives on all slaves in the active aggregator. Requires a switch that is 802.3ad compliant.

- balance-tlb or 5 — Sets a Transmit Load Balancing (TLB) policy for fault tolerance and load balancing. The outgoing traffic is distributed according to the current load on each slave interface. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed slave.

- balance-alb or 6 — Sets an Active Load Balancing (ALB) policy for fault tolerance and load balancing. Includes transmit and receive load balancing for IPV4 traffic. Receive load balancing is achieved through ARP negotiation.

num_unsol_na=<number>

Specifies the number of unsolicited IPv6 Neighbor Advertisements to be issued after a failover event. One unsolicited NA is issued immediately after the failover.

The valid range is 0 - 255; the default value is 1. This option affects only the active-backup mode.

primary=<interface_name>

Specifies the interface name, such as eth0, of the primary device. The primary device is the first of the bonding interfaces to be used and is not abandoned unless it fails. This setting is particularly useful when one NIC in the bonding interface is faster and, therefore, able to handle a bigger load.

This setting is only valid when the bonding interface is in active-backup mode. Refer to /usr/share/doc/kernel-doc-<kernel-version>/Documentation/networking/bonding.txt for more information.

primary_reselect=<value>

Specifies the reselection policy for the primary slave. This affects how the primary slave is chosen to become the active slave when failure of the active slave or recovery of the primary slave occurs. This option is designed to prevent flip-flopping between the primary slave and other slaves. Possible values are:

- always or 0 (default) — The primary slave becomes the active slave whenever it comes back up.

- better or 1 — The primary slave becomes the active slave when it comes back up, if the speed and duplex of the primary slave is better than the speed and duplex of the current active slave.

- failure or 2 — The primary slave becomes the active slave only if the current active slave fails and the primary slave is up.

The primary_reselect setting is ignored in two cases:

- If no slaves are active, the first slave to recover is made the active slave.

- When initially enslaved, the primary slave is always made the active slave.

Changing the primary_reselect policy via sysfs will cause an immediate selection of the best active slave according to the new policy. This may or may not result in a change of the active slave, depending upon the circumstances

updelay=<time_in_milliseconds>

Specifies (in milliseconds) how long to wait before enabling a link. The value must be a multiple of the value specified in the miimon parameter. The value is set to 0 by default, which disables it.

use_carrier=<number>

Specifies whether or not miimon should use MII/ETHTOOL ioctls or netif_carrier_ok() to determine the link state. The netif_carrier_ok() function relies on the device driver to maintains its state with netif_carrier_on/off ; most device drivers support this function.

The MII/ETHROOL ioctls tools utilize a deprecated calling sequence within the kernel. However, this is still configurable in case your device driver does not support netif_carrier_on/off .

Valid values are:

- 1 — Default setting. Enables the use of netif_carrier_ok().

- 0 — Enables the use of MII/ETHTOOL ioctls.

Tip

If the bonding interface insists that the link is up when it should not be, it is possible that your network device driver does not support netif_carrier_on/off .

xmit_hash_policy=<value>

Selects the transmit hash policy used for slave selection in balance-xor and 802.3ad modes. Possible values are:

- 0 or layer2 — Default setting. This option uses the XOR of hardware MAC addresses to generate the hash. The formula used is:

  ```
  (<source_MAC_address> XOR <destination_MAC>) MODULO <slave_count>
  ```

  This algorithm will place all traffic to a particular network peer on the same slave, and is 802.3ad compliant.

- 1 or layer3+4 — Uses upper layer protocol information (when available) to generate the hash. This allows for traffic to a particular network peer to span multiple slaves, although a single connection will not span multiple slaves.

  The formula for unfragmented TCP and UDP packets used is:

  ```
  ((<source_port> XOR <dest_port>) XOR
    ((<source_IP> XOR <dest_IP>) AND 0xffff)
      MODULO <slave_count>
  ```

  For fragmented TCP or UDP packets and all other IP protocol traffic, the source and destination port information is omitted. For non-IP traffic, the formula is the same as the layer2 transmit hash policy.

  This policy intends to mimic the behavior of certain switches; particularly, Cisco switches with PFC2 as well as some Foundry and IBM products.

  The algorithm used by this policy is not 802.3ad compliant.

- 2 or layer2+3 — Uses a combination of layer2 and layer3 protocol information to generate the hash.

  Uses XOR of hardware MAC addresses and IP addresses to generate the hash. The formula is:

  ```
  (((<source_IP> XOR <dest_IP>) AND 0xffff) XOR
    ( <source_MAC> XOR <destination_MAC> ))
      MODULO <slave_count>
  ```

  This algorithm will place all traffic to a particular network peer on the same slave. For non-IP traffic, the formula is the same as for the layer2 transmit hash policy.

  This policy is intended to provide a more balanced distribution of traffic than layer2 alone, especially in environments where a layer3 gateway device is required to reach most destinations.

  This algorithm is 802.3ad compliant.

43.6. Additional Resources

For more information on kernel modules and their utilities, refer to the following resources.

43.6.1. Installed Documentation

- lsmod man page — description and explanation of its output.
- insmod man page — description and list of command line options.
- modprobe man page — description and list of command line options.
- rmmod man page — description and list of command line options.
- modinfo man page — description and list of command line options.
- /usr/share/doc/kernel-doc-<version>/Documentation/kbuild/modules.txt — how to compile and use kernel modules. Note you must have the kernel-doc package installed to read this file.

43.6.2. Useful Websites

- http://tldp.org/HOWTO/Module-HOWTO/ — Linux Loadable Kernel Module HOWTO from the Linux Documentation Project.

The kdump Crash Recovery Service

kdump is an advanced crash dumping mechanism. When enabled, the system is booted from the context of another kernel. This second kernel reserves a small amount of memory, and its only purpose is to capture the core dump image in case the system crashes. Since being able to analyze the core dump helps significantly to determine the exact cause of the system failure, it is strongly recommended to have this feature enabled.

This chapter explains how to configure, test, and use the kdump service in Red Hat Enterprise Linux, and provides a brief overview of how to analyze the resulting core dump using the crash debugging utility.

44.1. Configuring the kdump Service

Note

To use the kdump service, you must have the kexec-tools package installed. Refer to II부. 패키지 관리 for more information on how to install new packages in Red Hat Enterprise Linux.

This section covers three common means of configuring the kdump service: at the first boot, using the Kernel Dump Configuration graphical utility, and doing so manually on the command line. It also describes how to test the configuration to verify that everything works as expected.

44.1.1. Configuring the kdump at First Boot

When the system boots for the first time, a firstboot application is launched allowing you to perform a basic configuration. This includes the kdump service.

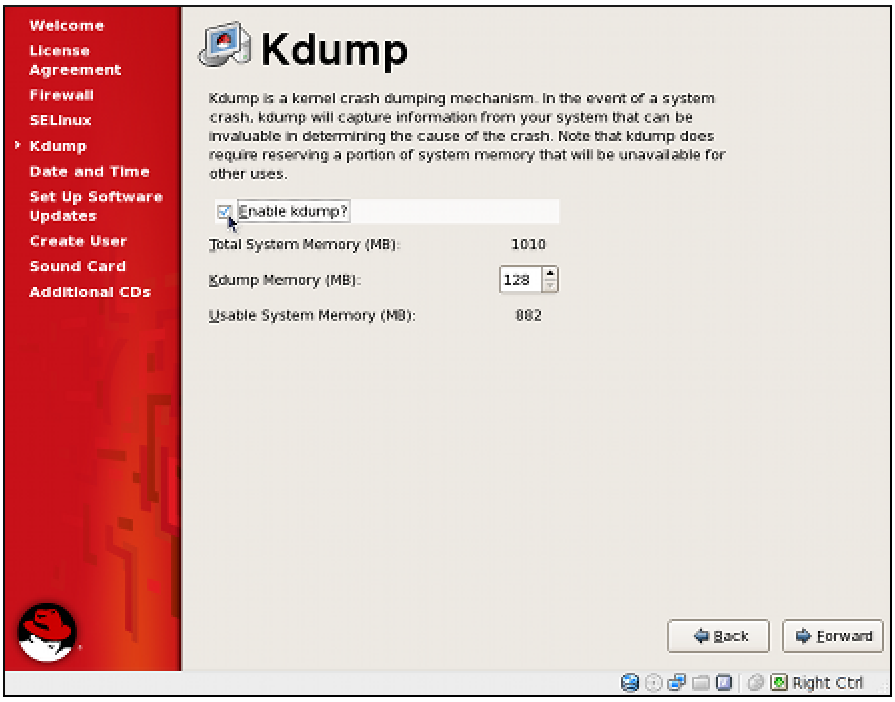

Kdump

Kdump is a kernel crash dumping mechanism. In the event of a system crash, kdump will capture information from your system that can be invaluable in determining the cause of the crash. Note that kdump does require reserving a portion of system memory that will be unavailable for other uses.

☑ Enable kdump?

Total System Memory (MB): 1010

Kdump Memory (MB): 128

Usable System Memory (MB): 882

그림 44.1. The kdump configuration screen

Important

Unless the system has enough memory, this option will not be available. For the information on minimum memory requirements, refer to the Required minimums section of the Red Hat Enterprise Linux comparison chart[1]. Note that when the kdump crash recovery is enabled, the minimum memory requirements increase by the amount of memory reserved for it. This value is determined by a user, and defaults to 128 MB.

44.1.1.1. Enabling the Service

To start the kdump daemon at boot time, select the Enable kdump? check box. This will enable the service for runlevels 2, 3, 4, and 5, and start it for the current session. Similarly, unselecting the check box will disable it for all runlevels and stop the service immediately.

44.1.1.2. Configuring the Memory Usage

To configure the amount of memory that is reserved for the kdump kernel, click the up and down arrow buttons next to the Kdump Memory field to increase or decrease the value. Notice that the Usable System Memory field changes accordingly showing you the remaining memory that will be available to the system.

44.1.2. Using the Kernel Dump Configuration Utility

To start the Kernel Dump Configuration utility, select Applications → System Tools → Kdump from the panel, or type system-config-kdump at a shell prompt (for example, xterm or GNOME Terminal). Unless you are already authenticated, you will be prompted to enter the superuser password.

[1] http://www.redhat.com/rhel/compare/

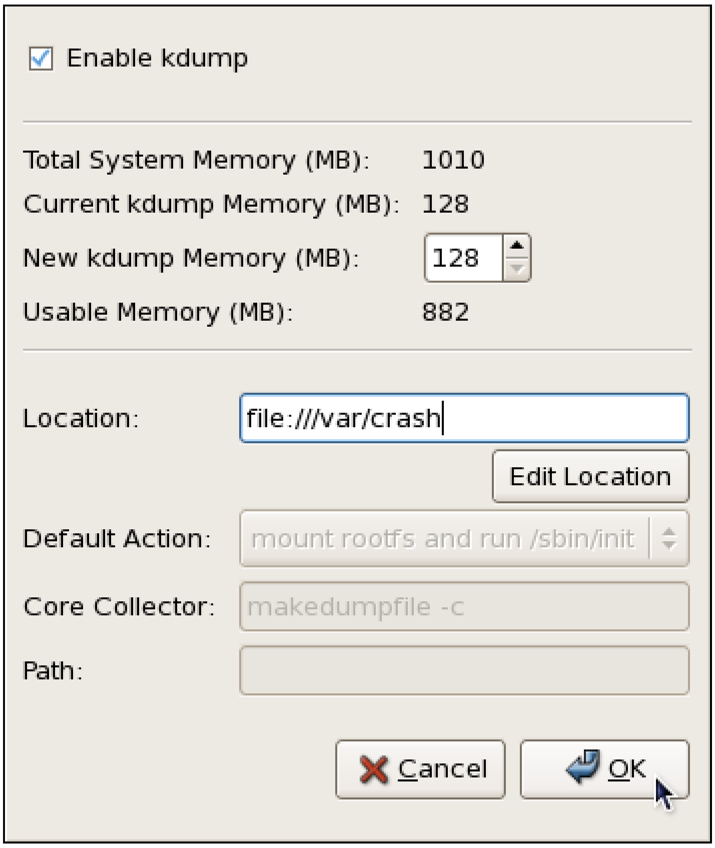

☑ Enable kdump

Total System Memory (MB): 1010

Current kdump Memory (MB): 128

New kdump Memory (MB): 128

Usable Memory (MB): 882

Location: file:///var/crash

Default Action: mount rootfs and run /sbin/init

Core Collector: makedumpfile -c

Path:

그림 44.2. The Kernel Dump Configuration utility

The utility allows you to configure kdump as well as to enable or disable starting the service at boot time. When you are done, click OK to save the changes. The system reboot will be requested.

Important

Unless the system has enough memory, the utility will not start, and you will be presented with an error message. For the information on minimum memory requirements, refer to the Required minimums section of the Red Hat Enterprise Linux comparison chart[2]. Note that when the kdump crash recovery is enabled, the minimum memory requirements increase by the amount of memory reserved for it. This value is determined by a user, and defaults to 128 MB.

44.1.2.1. Enabling the Service

To start the kdump daemon at boot time, select the Enable kdump check box. This will enable the service for runlevels 2, 3, 4, and 5, and start it for the current session. Similarly, unselecting the check box will disable it for all runlevels and stop the service immediately.

44.1.2.2. Configuring the Memory Usage

[2] http://www.redhat.com/rhel/compare/

To configure the amount of memory that is reserved for the kdump kernel, click the up and down arrow buttons next to the New kdump Memory field to increase or decrease the value. Notice that the Usable Memory field changes accordingly showing you the remaining memory that will be available to the system.

44.1.2.3. Configuring the Target Type

When a kernel crash is captured, the core dump can be either stored as a file in a local file system, written directly to a device, or sent over a network using the NFS (Network File System) or SSH (Secure Shell) protocol. To change this, click the Edit Location button, and select a location type as described below.

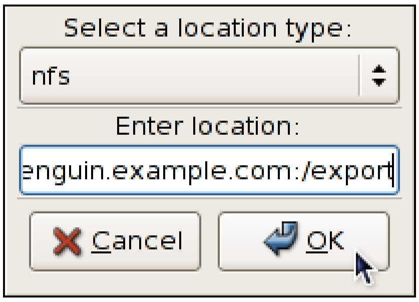

enguin.example.com:/export

그림 44.3. The Edit Location dialog

To save the dump to the local file system, select file from the pulldown list. Optionally, if you wish to write the file to a different partition, select ext3 or ext2 from the pulldown list according to the file system you are using, and enter a valid device name to the Enter location field. Note that after clicking OK, you can then customize the destination directory by changing the value in the Path field at the bottom.

To write the dump directly to a device, select raw from the pulldown list, and enter a valid device name (for example, /dev/sdb1). When you are done, click OK to confirm your choice.

To store the dump to a remote machine using the NFS protocol, select nfs from the pulldown list, and enter a valid target in the hostname:directory form (for example, penguin.example.com:/export). Clicking the OK button will confirm your changes. Finally, edit the value of the Path field to customize the destination directory (for instance, cores).

To store the dump to a remote machine using the SSH protocol, select ssh from the pulldown list, and enter a valid username and hostname in the username@hostname form (for example, john@penguin.example.com). Clicking the OK button will confirm your changes. Finally, edit the value of the Path field to customize the destination directory (for instance, /export/cores).

Refer to 19장. OpenSSH for information on how to configure an SSH server, and how to set up a key-based authentication.

44.1.2.4. Configuring the Core Collector

To reduce the size of the vmcore dump file, kdump allows you to specify an external application (that is, a core collector) to compress the data, and optionally leave out all irrelevant information. Currently, the only fully supported core collector is makedumpfile.

To enable the dump file compression, make sure the -c parameter is listed after the makedumpfile command in the Core Collector field (for example, makedumpfile -c).

To remove certain pages from the dump, add the -d value parameter after the makedumpfile command in the Core Collector field. The value is a sum of values of pages you want to omit as described in 표 44.1. “Supported filtering levels” . For example, to remove both zero and free pages, use makedumpfile -d 17.

Refer to the manual page for makedumpfile for a complete list of available options.

44.1.2.5. Changing the Default Action

To choose what action to perform when kdump fails to create a core dump, select the appropriate option from the Default Action pulldown list. Available options are mount rootfs and run /sbin/init (the default action), reboot (to reboot the system), shell (to present a user with an interactive shell prompt), and halt (to halt the system).

44.1.3. Configuring kdump on the Command Line

To perform actions described in this section, you have to be logged in as a superuser:

```
~]$ su -
Password:
```

44.1.3.1. Configuring the Memory Usage

To configure the amount of memory that is reserved for the kdump kernel, open the /boot/grub/grub.conf file in a text editor and add the crashkernel=<size>M@16M parameter to the list of kernel options as shown in 예 44.1. “A sample /boot/grub/grub.conf file” .

예 44.1. A sample /boot/grub/grub.conf file

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/sda3
#          initrd /initrd-version.img
#boot=/dev/sda
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux Server (2.6.18-274.3.1.el5)
        root (hd0,0)
        kernel /vmlinuz-2.6.18-274.3.1.el5 ro root=/dev/sda3 crashkernel=128M@16M
        initrd /initrd-2.6.18-274.3.1.el5.img
```

Important

When the kdump crash recovery is enabled, the minimum memory requirements increase by the amount of memory reserved for it. This value is determined by a user, and defaults to 128 MB, as lower values proved to be unreliable. For more information on minimum memory requirements for Red Hat Enterprise Linux, refer to the Required minimums section of the Red Hat Enterprise Linux comparison chart[3].

44.1.3.2. Configuring the Target Type

When a kernel crash is captured, the core dump can be either stored as a file in a local file system, written directly to a device, or sent over a network using the NFS (Network File System) or SSH (Secure Shell) protocol. Note that only one of these options can be set at the moment. The default option is to store the vmcore file in the /var/crash/ directory of the local file system. To change this, open the /etc/kdump.conf configuration file in a text editor and edit the options as described below.

To change the local directory in which the core dump is to be saved, remove the hash sign (“#”) from the beginning of the #path /var/crash line, and replace the value with a desired directory path. Optionally, if you wish to write the file to a different partition, follow the same procedure with the #ext3 /dev/sda3 line as well, and change both the file system type and the device (a device name, a file system label, and UUID are all supported) accordingly. For example:

```
ext3 /dev/sda4
path /usr/local/cores
```

To write the dump directly to a device, remove the hash sign (“#”) from the beginning of the #raw /dev/sda5 line, and replace the value with a desired device name. For example:

```
raw /dev/sdb1
```

To store the dump to a remote machine using the NFS protocol, remove the hash sign (“#”) from the beginning of the #net my.server.com:/export/tmp line, and replace the value with a valid hostname and directory path. For example:

```
net penguin.example.com:/export/cores
```

To store the dump to a remote machine using the SSH protocol, remove the hash sign (“#”) from the beginning of the #net user@my.server.com line, and replace the value with a valid username and hostname. For example:

```
net john@penguin.example.com
```

Refer to 19장. OpenSSH for information on how to configure an SSH server, and how to set up a key-based authentication.

[3] http://www.redhat.com/rhel/compare/

44.1.3.3. Configuring the Core Collector

To reduce the size of the vmcore dump file, kdump allows you to specify an external application (that is, a core collector) to compress the data, and optionally leave out all irrelevant information. Currently, the only fully supported core collector is makedumpfile.

To enable the core collector, open the /etc/kdump.conf configuration file in a text editor, remove the hash sign (“#”) from the beginning of the #core_collector makedumpfile -c --message-level 1 line, and edit the command line options as described below.

To enable the dump file compression, add the -c parameter. For example:

```
core_collector makedumpfile -c
```

To remove certain pages from the dump, add the -d value parameter, where value is a sum of values of pages you want to omit as described in 표 44.1. “Supported filtering levels” . For example, to remove both zero and free pages, use the following:

```
core_collector makedumpfile -d 17 -c
```

Refer to the manual page for makedumpfile for a complete list of available options.

표 44.1. Supported filtering levels

Option	Description
1	Zero pages
2	Cache pages
4	Cache private
8	User pages
16	Free pages

44.1.3.4. Changing the Default Action

By default, when kdump fails to create a core dump, the root file system is mounted and /sbin/init is run. To change this behavior, open the /etc/kdump.conf configuration file in a text editor, remove the hash sign (“#”) from the beginning of the #default shell line, and replace the value with a desired action as described in 표 44.2. “Supported actions” . For example:

```
default halt
```

표 44.2. Supported actions

Option	Action
reboot	Reboot the system, losing the core in the process.
halt	After failing to capture a core, halt the system.
shell	Run the msh session from within the initramfs, allowing a user to record the core manually.

44.1.3.5. Enabling the Service

To start the kdump daemon at boot time, type the following at a shell prompt:

```
~]# chkconfig kdump on
```

This will enable the service for runlevels 2, 3, 4, and 5. Similarly, typing chkconfig kdump off will disable it for all runlevels. To start the service in the current session, use the following command:

```
~]# service kdump start
No kdump initial ramdisk found.                    [WARNING]
Rebuilding /boot/initrd-2.6.18-194.8.1.el5kdump.img
Starting kdump:                                    [  OK  ]
```

For more information on runlevels and configuring services in general, refer to 17장. 서비스로의 접근 통제.

44.1.4. Testing the Configuration

Caution

The commands below will cause the kernel to crash. Use caution when following these steps, and by no means use them on a production machine.

To test the configuration, reboot the system with kdump enabled, and make sure that the service is running:

```
~]# service kdump status
Kdump is operational
```

Then type the following commands at a shell prompt:

```
~]# echo 1 > /proc/sys/kernel/sysrq
~]# echo c > /proc/sysrq-trigger
```

This will force the Linux kernel to crash, and the YYYY-MM-DD-HH:MM/vmcore file will be copied to the location you have selected in the configuration (that is, to /var/crash/ by default).

44.2. Analyzing the Core Dump

Note

To analyze the vmcore dump file, you must have the crash and kernel-debuginfo packages installed. To do so, type the following at a shell prompt:

```
~]# yum install --enablerepo=rhel-debuginfo crash kernel-debuginfo
```

Refer to II부. 패키지 관리 for more information on how to install new packages in Red Hat Enterprise Linux.

To determine the cause of the system crash, you can use the crash utility. This utility allows you to interactively analyze a running Linux system as well as a core dump created by netdump, diskdump, xendump, or kdump. When started, it presents you with an interactive prompt very similar to the GNU Debugger (GDB).

To start the utility, type the command in the following form at a shell prompt:

```
crash /var/crash/timestamp/vmcore /usr/lib/debug/lib/modules/kernel/vmlinux
```

Note that the kernel version should be the same as the one that was captured by kdump. To find out which kernel you are currently running, use the uname -r command.

예 44.2. Running the crash utility

```
~]# crash /var/crash/2010-08-04-17\:55/vmcore \
/usr/lib/debug/lib/modules/2.6.18-194.8.1.el5/vmlinux

crash 4.1.2-4.el5_5.1
Copyright (C) 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009  Red Hat, Inc.
Copyright (C) 2004, 2005, 2006  IBM Corporation
Copyright (C) 1999-2006  Hewlett-Packard Co
Copyright (C) 2005, 2006  Fujitsu Limited
Copyright (C) 2006, 2007  VA Linux Systems Japan K.K.
Copyright (C) 2005  NEC Corporation
Copyright (C) 1999, 2002, 2007  Silicon Graphics, Inc.
Copyright (C) 1999, 2000, 2001, 2002  Mission Critical Linux, Inc.
This program is free software, covered by the GNU General Public License,
and you are welcome to change it and/or distribute copies of it under
certain conditions.  Enter "help copying" to see the conditions.
This program has absolutely no warranty.  Enter "help warranty" for details.

GNU gdb 6.1
Copyright 2004 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License, and you are
welcome to change it and/or distribute copies of it under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB.  Type "show warranty" for details.
This GDB was configured as "i686-pc-linux-gnu"...

      KERNEL: /usr/lib/debug/lib/modules/2.6.18-194.8.1.el5/vmlinux
    DUMPFILE: /var/crash/2010-08-04-17:55/vmcore
        CPUS: 1
        DATE: Wed Aug  4 17:50:41 2010
      UPTIME: 00:56:53
LOAD AVERAGE: 0.47, 0.47, 0.55
```

```
       TASKS: 128
    NODENAME: localhost.localdomain
     RELEASE: 2.6.18-194.el5
     VERSION: #1 SMP Tue Mar 16 21:52:43 EDT 2010
     MACHINE: i686  (2702 Mhz)
      MEMORY: 1 GB
       PANIC: "SysRq : Trigger a crashdump"
         PID: 6042
     COMMAND: "bash"
        TASK: f09c7000  [THREAD_INFO: e1ba9000]
         CPU: 0
       STATE: TASK_RUNNING (SYSRQ)

crash>
```

To exit the interactive prompt and terminate crash, type exit.

44.2.1. Displaying the Message Buffer

To display the kernel message buffer, type the log command at the interactive prompt.

예 44.3. Displaying the kernel message buffer

```
crash> log
Linux version 2.6.18-194.el5 (mockbuild@x86-007.build.bos.redhat.com) (gcc version 4.1.2 20080704 (Red Hat 4.1.2-48))
 #1 SMP Tue Mar 16 21:52:43 EDT 2010
BIOS-provided physical RAM map:
 BIOS-e820: 0000000000010000 - 000000000009fc00 (usable)
 BIOS-e820: 000000000009fc00 - 00000000000a0000 (reserved)
 BIOS-e820: 00000000000f0000 - 0000000000100000 (reserved)
 BIOS-e820: 0000000000100000 - 000000003fff0000 (usable)
 BIOS-e820: 000000003fff0000 - 0000000040000000 (ACPI data)
 BIOS-e820: 00000000fffc0000 - 0000000100000000 (reserved)
127MB HIGHMEM available.
896MB LOWMEM available.
Using x86 segment limits to approximate NX protection
On node 0 totalpages: 262128
  DMA zone: 4096 pages, LIFO batch:0
  Normal zone: 225280 pages, LIFO batch:31
  HighMem zone: 32752 pages, LIFO batch:7
DMI 2.5 present.
Using APIC driver default
... several lines omitted ...
SysRq : Trigger a crashdump
```

Type help log for more information on the command usage.

44.2.2. Displaying a Backtrace

To display the kernel stack trace, type the bt command at the interactive prompt. You can use bt pid to display the backtrace of the selected process.

예 44.4. Displaying the kernel stack trace

```
crash> bt
PID: 6042   TASK: f09c7000  CPU: 0   COMMAND: "bash"
 #0 [e1ba9d10] schedule at c061c738
 #1 [e1ba9d28] netlink_getsockopt at c05d50bb
 #2 [e1ba9d34] netlink_queue_skip at c05d40d5
 #3 [e1ba9d40] netlink_sock_destruct at c05d506d
 #4 [e1ba9d84] sock_recvmsg at c05b6cc8
 #5 [e1ba9dd4] enqueue_task at c041eed5
 #6 [e1ba9dec] try_to_wake_up at c041f798
 #7 [e1ba9e10] vsnprintf at c04efef2
 #8 [e1ba9ec0] machine_kexec at c0419bf0
 #9 [e1ba9f04] sys_kexec_load at c04448a1
#10 [e1ba9f4c] tty_audit_exit at c0549f06
#11 [e1ba9f50] tty_audit_add_data at c0549d5d
#12 [e1ba9f84] do_readv_writev at c0476055
#13 [e1ba9fb8] system_call at c0404f10
    EAX: fffffffda  EBX: 00000001  ECX: b7f7f000  EDX: 00000002
    DS:  007b       ESI: 00000002  ES:  007b      EDI: b7f7f000
    SS:  007b       ESP: bf83f478  EBP: bf83f498
    CS:  0073       EIP: 009ac402  ERR: 00000004  EFLAGS: 00000246
```

Type help bt for more information on the command usage.

44.2.3. Displaying a Process Status

To display a status of processes in the system, type the ps command at the interactive prompt. You can use ps pid to display the status of the selected process.

예 44.5. Displaying status of processes in the system

```
crash> ps
   PID    PPID  CPU   TASK    ST  %MEM     VSZ    RSS  COMM
      0      0   0  c068a3c0  RU   0.0       0      0  [swapper]
      1      0   0  f7c81aa0  IN   0.1    2152    616  init
... several lines omitted ...
   6017      1   0  e39f6550  IN   1.2   40200  13000  gnome-terminal
   6019   6017   0  e39f6000  IN   0.1    2568    708  gnome-pty-helpe
   6020   6017   0  f0421550  IN   0.1    4620   1480  bash
   6021      1   0  f7f69aa0  ??   1.2   40200  13000  gnome-terminal
   6039   6020   0  e7e84aa0  IN   0.1    5004   1300  su
>  6042   6039   0  f09c7000  RU   0.1    4620   1464  bash
```

Type help ps for more information on the command usage.

44.2.4. Displaying Virtual Memory Information

To display basic virtual memory information, type the vm command at the interactive prompt. You can use vm pid to display information on the selected process.

예 44.6. Displaying virtual memory information of the current context

```
crash> vm
PID: 6042   TASK: f09c7000  CPU: 0   COMMAND: "bash"
   MM       PGD      RSS    TOTAL_VM
e275ee40  e2b08000  1464k    4620k
  VMA       START      END     FLAGS  FILE
e315d764    1fe000    201000      75  /lib/libtermcap.so.2.0.8
e315de9c    201000    202000  100073  /lib/libtermcap.so.2.0.8
c9b040d4    318000    46a000      75  /lib/libc-2.5.so
```

```
e315da04    46a000    46c000  100071   /lib/libc-2.5.so
e315d7b8    46c000    46d000  100073   /lib/libc-2.5.so
e315de48    46d000    470000  100073
e315dba8    9ac000    9ad000  8040075
c9b04a04    a2f000    a4a000     875   /lib/ld-2.5.so
c9b04374    a4a000    a4b000  100871   /lib/ld-2.5.so
e315d6bc    a4b000    a4c000  100873   /lib/ld-2.5.so
e315d908    fa1000    fa4000      75   /lib/libdl-2.5.so
e315db00    fa4000    fa5000  100071   /lib/libdl-2.5.so
e315df44    fa5000    fa6000  100073   /lib/libdl-2.5.so
e315d320    ff0000    ffa000      75   /lib/libnss_files-2.5.so
e315d668    ffa000    ffb000  100071   /lib/libnss_files-2.5.so
e315def0    ffb000    ffc000  100073   /lib/libnss_files-2.5.so
e315d374   8048000   80f5000    1875   /bin/bash
c9b045c0   80f5000   80fa000  101873   /bin/bash
... several lines omitted ...
```

Type help vm for more information on the command usage.

44.2.5. Displaying Open Files

To display information about open files, type the files command at the interactive prompt. You can use files pid to display files opened by the selected process.

예 44.7. Displaying information about open files of the current context

```
crash> files
PID: 6042   TASK: f09c7000  CPU: 0   COMMAND: "bash"
ROOT: /    CWD: /root
 FD     FILE      DENTRY     INODE     TYPE  PATH
  0  e33be480  e609bf70  f0e1d880  CHR   /dev/pts/1
  1  e424d8c0  d637add8  f7809b78  REG   /proc/sysrq-trigger
  2  e33be480  e609bf70  f0e1d880  CHR   /dev/pts/1
 10  e33be480  e609bf70  f0e1d880  CHR   /dev/pts/1
255  e33be480  e609bf70  f0e1d880  CHR   /dev/pts/1
```

Type help files for more information on the command usage.

44.3. Additional Resources

44.3.1. Installed Documentation

man kdump.conf

The manual page for the /etc/kdump.conf configuration file containing the full documentation of available options.

man kexec

The manual page for kexec containing the full documentation on its usage.

man crash

The manual page for the crash utility containing the full documentation on its usage.

/usr/share/doc/kexec-tools-version/kexec-kdump-howto.txt

An overview of the kdump and kexec installation and usage.

44.3.2. Useful Websites

https://access.redhat.com/kb/docs/DOC-6039
: The Red Hat Knowledgebase article about the kexec and kdump configuration.

http://people.redhat.com/anderson/
: The crash utility homepage.

부 VII. 보안 및 인증

시스템 관리자는 필요 시스템, 서비스, 데이터에 대한 보안이 필요하며 Red Hat Enterprise Linux는 포괄적인 보안 전략의 일부분으로서 그에 맞는 다양한 도구 및 방식을 제공합니다.

이 장에서는 Red Hat Enterprise Linux 의 관점에서 보안에 관해 전반적으로 소개하고 있습니다. 이는 보안 평가, 일반적 사용, 침입 및 사고 대응에 관한 개념적 정보를 제공함은 물론, 워크스테이션, 서버, VPN, 방화벽 및 기타 다른 기능 실행을 강화하기 위해 SELinux 를 사용하는 방법에 관한 개념적이고 자세한 설정 정보도 제공합니다.

이 장에서는 IT 보안에 관한 기본적인 지식을 갖고 있다고 간주하여 물리적 접근 제어하기, 철저한 계정 유지 정책 및 절차, 감사 (auditing) 등과 같이 일반적인 보안 사항에 대해 최소한의 정보만을 제공합니다. 이와 관련된 정보는 외부 참고 자료로 마련되어 있습니다.

보안 개요

기업을 운영하고 개인 정보를 관리함에 있어서 네트워크로 연결된 강력한 컴퓨터의 사용이 나날이 증가하고 있습니다. 대부분의 모든 산업 생산이 네트워크와 컴퓨터 보안을 중심으로 이루어 집니다. 기업체들은 보안 전문가의 지식과 기술을 바탕으로 시스템을 점검하고 각 기업체에서 필요로 하는 운영 체제 요건에 맞게 이를 적절히 설정합니다. 대부분의 기업체는 직원들이 회사 IT 자원을 지역적으로나 원격적으로 액세스하는 끊임없이 변화하는 환경을 갖추고 있습니다. 따라서 그 어느 때 보다 컴퓨터 시스템 보안에 대한 필요성이 매우 강조되고 있습니다.

그러나 불행히도 개인 사용자는 물론 대부분의 기업에서는 컴퓨터 용량을 늘여 생산성을 높이는데 더 중점을 두고 예산 문제로 컴퓨터 보안의 중요성을 간과하고 있습니다. 종종 postmortem (사후 검토) —로써 이미 시스템에 제 3자가 침입한 후에 적절한 보안 조치를 취하는 경우가 많습니다. 컴퓨터 보안 전문가들은 인터넷과 같이 신뢰할 수 없는 네트워크에 연결하기 이전에 적절한 보안 조치를 취하는 것이 대부분의 침입을 방지할 수 있는 최선의 방법이라고 동의합니다.

45.1. Introduction to Security

45.1.1. What is Computer Security?

Computer security is a general term that covers a wide area of computing and information processing. Industries that depend on computer systems and networks to conduct daily business transactions and access crucial information regard their data as an important part of their overall assets. Several terms and metrics have entered our daily business vocabulary, such as total cost of ownership (TCO) and quality of service (QoS). In these metrics, industries calculate aspects such as data integrity and high-availability as part of their planning and process management costs. In some industries, such as electronic commerce, the availability and trustworthiness of data can be the difference between success and failure.

45.1.1.1. How did Computer Security Come about?

Information security has evolved over the years due to the increasing reliance on public networks not to disclose personal, financial, and other restricted information. There are numerous instances such as the Mitnick and the Vladimir Levin cases that prompted organizations across all industries to rethink the way they handle information transmission and disclosure. The popularity of the Internet was one of the most important developments that prompted an intensified effort in data security.

An ever-growing number of people are using their personal computers to gain access to the resources that the Internet has to offer. From research and information retrieval to electronic mail and commerce transaction, the Internet has been regarded as one of the most important developments of the 20th century.

The Internet and its earlier protocols, however, were developed as a trust-based system. That is, the Internet Protocol was not designed to be secure in itself. There are no approved security standards built into the TCP/IP communications stack, leaving it open to potentially malicious users and processes across the network. Modern developments have made Internet communication more secure, but there are still several incidents that gain national attention and alert us to the fact that nothing is completely safe.

45.1.1.2. Security Today

In February of 2000, a Distributed Denial of Service (DDoS) attack was unleashed on several of the most heavily-trafficked sites on the Internet. The attack rendered yahoo.com, cnn.com, amazon.com, fbi.gov, and several other sites completely unreachable to normal users, as it tied up routers for several hours with large-byte ICMP packet transfers, also called a ping flood. The attack was brought on by unknown assailants using specially created, widely available programs that scanned vulnerable network servers, installed client applications called Trojans on the servers, and timed an attack with every infected server flooding the victim sites and rendering them unavailable. Many blame the attack on fundamental flaws in the way routers and the protocols used are structured to accept all incoming data, no matter where or for what purpose the packets are sent.

Currently, an estimated 945 million people use or have used the Internet worldwide (Computer Industry Almanac, 2004). At the same time:

- On any given day, there are approximately 225 major incidences of security breach reported to the CERT Coordination Center at Carnegie Mellon University.[1]

- In 2003, the number of CERT reported incidences jumped to 137,529 from 82,094 in 2002 and from 52,658 in 2001.[2]

- The worldwide economic impact of the three most dangerous Internet Viruses of the last three years was estimated at US$13.2 Billion.[3]

Computer security has become a quantifiable and justifiable expense for all IT budgets. Organizations that require data integrity and high availability elicit the skills of system administrators, developers, and engineers to ensure 24x7 reliability of their systems, services, and information. Falling victim to malicious users, processes, or coordinated attacks is a direct threat to the success of the organization.

Unfortunately, system and network security can be a difficult proposition, requiring an intricate knowledge of how an organization regards, uses, manipulates, and transmits its information. Understanding the way an organization (and the people that make up the organization) conducts business is paramount to implementing a proper security plan.

45.1.1.3. Standardizing Security

Enterprises in every industry rely on regulations and rules that are set by standards making bodies such as the American Medical Association (AMA) or the Institute of Electrical and Electronics Engineers (IEEE). The same ideals hold true for information security. Many security consultants and vendors agree upon the standard security model known as CIA, or Confidentiality, Integrity, and Availability. This three-tiered model is a generally accepted component to assessing risks of sensitive information and establishing security policy. The following describes the CIA model in further detail:

- Confidentiality — Sensitive information must be available only to a set of pre-defined individuals. Unauthorized transmission and usage of information should be restricted. For example, confidentiality of information ensures that a customer's personal or financial information is not obtained by an unauthorized individual for malicious purposes such as identity theft or credit fraud.

- Integrity — Information should not be altered in ways that render it incomplete or incorrect. Unauthorized users should be restricted from the ability to modify or destroy sensitive information.

- Availability — Information should be accessible to authorized users any time that it is needed. Availability is a warranty that information can be obtained with an agreed-upon frequency and

[1] Source: http://www.cert.org

[2] Source: http://www.cert.org/stats/

[3] Source: http://www.newsfactor.com/perl/story/16407.html

timeliness. This is often measured in terms of percentages and agreed to formally in Service Level Agreements (SLAs) used by network service providers and their enterprise clients.

45.1.2. Security Controls

Computer security is often divided into three distinct master categories, commonly referred to as controls:

- Physical
- Technical
- Administrative

These three broad categories define the main objectives of proper security implementation. Within these controls are sub-categories that further detail the controls and how to implement them.

45.1.2.1. Physical Controls

Physical control is the implementation of security measures in a defined structure used to deter or prevent unauthorized access to sensitive material. Examples of physical controls are:

- Closed-circuit surveillance cameras
- Motion or thermal alarm systems
- Security guards
- Picture IDs
- Locked and dead-bolted steel doors
- Biometrics (includes fingerprint, voice, face, iris, handwriting, and other automated methods used to recognize individuals)

45.1.2.2. Technical Controls

Technical controls use technology as a basis for controlling the access and usage of sensitive data throughout a physical structure and over a network. Technical controls are far-reaching in scope and encompass such technologies as:

- Encryption
- Smart cards
- Network authentication
- Access control lists (ACLs)
- File integrity auditing software

45.1.2.3. Administrative Controls

Administrative controls define the human factors of security. It involves all levels of personnel within an organization and determines which users have access to what resources and information by such means as:

- Training and awareness
- Disaster preparedness and recovery plans
- Personnel recruitment and separation strategies
- Personnel registration and accounting

45.1.3. Conclusion

Now that you have learned about the origins, reasons, and aspects of security, you can determine the appropriate course of action with regard to Red Hat Enterprise Linux. It is important to know what factors and conditions make up security in order to plan and implement a proper strategy. With this information in mind, the process can be formalized and the path becomes clearer as you delve deeper into the specifics of the security process.

45.2. 취약성 평가

충분한 시간, 자원과 동기만 있다면 크래커는 거의 어느 시스템이든 침입 가능합니다. 결국 현재 사용되는 어느 보안 절차와 기술로도 시스템이 침입당하지 않으리라고 보장하지 못합니다. 라우터는 인터넷으로부터 게이트웨이를 보호해주며 방화벽은 네트워크의 보안을 도와줍니다. 가상 사설 네트워크는 데이터를 암호화된 스트림으로 안전하게 전달해주며 수상한 활동이 탐지되었을 경우 여러분께 경고해주도록 침입 탐지 시스템을 사용할 수 있습니다. 그러나 다음과 같은 여러 변수에 따라 이러한 기술의 성공적인 사용 여부가 좌우됩니다:

- 기술을 설정하고 모니터하며 관리할 수 있는 전문 기술자
- 서비스와 커널을 신속하고 효율적으로 패치하고 업데이트할 수 있는 능력
- 네트워크 활동을 계속적으로 경계할 수 있는 담당자

데이터 시스템과 기술이 매우 동적으로 변화되기 때문에 회사 자원을 안전하게 지키는 작업은 결코 쉽지가 않습니다. 이러한 복잡한 문제로 인하여 회사 시스템 전부를 잘 이해하는 전문가를 찾기가 힘든 경우가 있습니다. 고급 수준의 다양한 영역의 정보 보안에 대하여 알고 있는 직원을 채용하는 것은 가능하지만, 여러 보안 영역의 전문가를 채용하는 것은 쉽지 않습니다. 그 이유는 각 정보 보안 분야가 멈추치 않고 변화하기 때문에 계속적인 관심과 집중을 필요로하기 때문입니다.

45.2.1. 적의 마음으로 생각하기

Suppose that you administer an enterprise network. Such networks are commonly comprised of operating systems, applications, servers, network monitors, firewalls, intrusion detection systems, and more. Now imagine trying to keep current with each of these. Given the complexity of today's software and networking environments, exploits and bugs are a certainty. Keeping current with patches and updates for an entire network can prove to be a daunting task in a large organization with heterogeneous systems.

이처럼 시스템을 계속적으로 감시하는 작업과 여러 분야의 보안 전문가를 찾는 어려움을 결합하여 보았을때 시스템 보안 침해가 발생하여 데이터가 손상되고 서비스가 중단되는 문제를 피할 수 없습니다.

보안 기술을 증대시키고 시스템, 네트워크 및 데이터 보안을 돕기 위하여 자신이 크래커라면 어떠한 시스템 헛점을 악용하여 보안을 침해할 것인지 한번 생각해 보십시오. 여러분의 시스템과 네트워크 자원에 예방적인 취약성 평가를 해봄으로서 잠재적으로 문제가 될 만한 사항들을 크래커가 헛점으로 사용하기 전에 미리 발견할 수 있습니다.

A vulnerability assessment is an internal audit of your network and system security; the results of which indicate the confidentiality, integrity, and availability of your network (as explained in 45.1.1.3절. “Standardizing Security”). Typically, vulnerability assessment starts with a reconnaissance phase, during which important data regarding the target systems and resources is gathered. This phase leads to the system readiness phase, whereby the target is essentially checked for all known vulnerabilities. The readiness phase culminates in the reporting phase, where the findings are classified into categories of high, medium, and low risk; and methods for improving the security (or mitigating the risk of vulnerability) of the target are discussed.

만일 집안의 취약성 평가를 수행하신다면 문들이 제대로 닫혀있는지 잠겨있는지를 확인하실 것입니다. 또한 창문들도 모두 완전히 잠겼는지 제대로 빗장이 걸려있는지 확인하실 것입니다. 이러한 동일한 개념이 시스템, 네트워크 및 컴퓨터 데이터에도 적용됩니다. 악의를 지닌 시스템 침입자는 여러분의 데이터를 훔쳐가는 도둑입니다. 침입자가 사용하는 도구가 무엇인지, 어떠한 생각을 하고 있는지 침입한 동기는 무엇인지를 먼저 생각해보신 후 즉각적으로 조치를 취하시기 바랍니다.

45.2.2. 평가와 테스팅 정의하기

취약성 평가는 다음과 같은 두가지 유형으로 구분될 수 있습니다: 외부에서 내부를 평가하기 및 내부에서 외부를 평가하기.

When performing an outside looking in vulnerability assessment, you are attempting to compromise your systems from the outside. Being external to your company provides you with the cracker's viewpoint. You see what a cracker sees — publicly-routable IP addresses, systems on your DMZ, external interfaces of your firewall, and more. DMZ stands for "demilitarized zone", which corresponds to a computer or small subnetwork that sits between a trusted internal network, such as a corporate private LAN, and an untrusted external network, such as the public Internet. Typically, the DMZ contains devices accessible to Internet traffic, such as Web (HTTP) servers, FTP servers, SMTP (e-mail) servers and DNS servers.

내부에서 취약성 평가를 수행하시는 경우 여러분이 내부에 위치하며 신뢰받는 위치에 있으므로 어느정도 유리한 입장에 있습니다. 이것이 여러분이나 함께 근무하는 직원이 시스템에 로그인시 위치하게 되는 지점입니다. 프린트 서버, 파일 서버, 데이터베이스 및 다른 자원을 살펴보는 것이 가능합니다.

이러한 두 가지 유형의 취약성 평가 사이에는 두드러진 차이점이 있습니다. 회사 내부에 위치하면 외부 사용자보다 많은 권한을 갖게 됩니다. 오늘날 대부분의 기업체에서는 외부인의 침입을 막을 수 있도록 보안을 설정합니다. 그러나 기업 내부에서의 침입을 막을 수 있는 조치를 거의 취하지 않고 있습니다 (예, 부서마다 방화벽 설치, 사용자-수준 접근 제어, 내부 자원을 위한 인증 절차 등). 일반적으로 대부분의 시스템은 회사 내부에 위치하므로 내부에서 보다 많은 자원을 찾을 수 있습니다. 만일 회사 외부에서 접근을 시도하시면 즉시 신뢰할 수 없는 상태로 간주됩니다. 회사 외부에서 접근 가능한 시스템과 자원은 일반적으로 매우 제한되어 있습니다.

취약성 평가와 침입 데스트 간의 차이점을 생각해 보십시오. 취약성 평가를 침입 데스트를 위한 첫번째 단계로 간주할 수 있습니다. 취약성 평가를 통해 수집한 정보는 테스팅에 사용됩니다. 취약성 평가에서는 보안 헛점과 잠재적 취약점을 찾는 반면, 침입 데스팅에서는 실제로 발견된 취약점을 사용하여 침입을 시도합니다.

네트워크 기반 구조를 평가하는 작업은 동적인 과정입니다. 정보 보안 및 물리적 보안 모두 동적이라 할 수 있습니다. 네트워크 평가를 수행함으로서 존재하지 않는 위험성이 보고되거나 존재하는 위험성이 보고되지 않는 경우를 발견할 수 있습니다.

보안 관리자는 자신의 지식과 보안 관리를 위해 사용되는 도구에 의존하고 있습니다. 현재 사용 가능한 평가 도구 중 하나를 한번 시스템 상에 실행시켜 보십시오. 거의 매번 보고된 문제점 중 최소한 몇 개는 존재하지 않는 문제점입니다. 프로그램이 잘못되었거나 사용자의 잘못인지 여부를 떠나서 결과는 언제나 같습니다. 도구는 실제로는 존재하지 않는 취약점을 찾아내거나 (false positive); 또는 실제로는 존재하는 취약점을 찾아내지 못하는 경우 (false negative)가 있습니다.

이제 취약성 평가와 침입 테스트 간의 차이점을 알아보았으니 평가 후 찾아낸 사항들을 주의깊게 검토한 후 침투 테스트를 실행해 보시기 바랍니다.

> **경고**
>
> 생산 자원의 취약점을 이용하여 침입 테스트를 시도하시면 회사 시스템과 네트워크의 생산성과 효율성에 반대 영향을 미칠 수 있으니 주의하십시오.

다음 목록에서는 취약성 평가를 수행함으로서 받을 수 있는 여러 가지 혜택을 설명하고 있습니다:

- 정보 보안에 대해 사전 대처할 수 있음
- 크래커가 찾아내기 전에 잠재적 취약점을 찾아낼 수 있음
- 시스템이 항상 업데이트되고 패치될 수 있음
- 직원의 전문적 기술을 증대시키는데 도움이 됨
- 재정적 손실과 부정적인 회사 이미지를 줄일 수 있음

45.2.2.1. 방법론 수립

취약성 평가에 사용될 도구 선택을 위하여 취약성 평가 방법론을 먼저 수립하시기 바랍니다. 불행히 아직 미리 정의되거나 산업체에서 인증받은 방법이 존재하지 않지만 일반 상식과 실행 결과만으로도 평가 방법론에 대한 충분한 길잡이가 될 수 있습니다.

어느 시스템을 상대로 하는가? 한 개의 서버를 상대하는가 또는 전체 네트워크 및 그 네트워크 내의 모든 시스템을 상대로 하는가? 현재 회사 외부에 위치하고 있는가 내부에 위치하고 있는가? 이러한 질문에 대한 해답을 찾는 것은 사용할 적절한 도구를 찾는데 도움이 될 뿐만 아니라 그 도구를 어떠한 방법으로 사용할 지 결정 가능하게 합니다.

방법론 수립에 대한 보다 많은 정보를 원하신다면 다음 웹사이트를 참조하시기 바랍니다:

- http://www.isecom.org/projects/osstmm.htm The Open Source Security Testing Methodology Manual (OSSTMM)
- http://www.owasp.org/ The Open Web Application Security Project

45.2.3. 도구를 평가하기

평가 과정은 정보 수집 도구를 사용함으로서 시작됩니다. 전체 네트워크를 평가하실 때에는 우선 실행 중인 호스트를 찾아내기 위해 네트워크 배치를 자세히 살펴보십시오. 일단 호스트를 찾으면 각 호스트를 개별적으로 검사해보십시오. 호스트 검사를 위해서는 또 다른 도구를 사용하셔야 합니다. 어떠한 도구를 사용하느냐에 따라서 호스트의 취약성을 찾아내는데 중요한 역할을 합니다.

일상 생활에서와 마찬가지로 동일한 작업을 수행하는데 여러 다른 도구를 사용할 수 있습니다. 취약성 평가를 수행시에도 마찬가지 입니다. 운영 체제에 특별히 사용되는 도구가 있으며 또한 응용

프로그램 및 심지어는 사용된 프로토콜에 따라서 네트워크에 특별히 사용되는 도구도 따로 존재합니다. 일부 도구는 사용이 무료이며 유료인 도구도 있습니다. 어떠한 도구는 직관적으로 이해가 가능하며 사용하기 쉽지만, 일부 다른 도구는 애매하며 제대로 문서화되어 있지 않아 사용하기 힘들지만 다른 도구가 가지지 않는 특별한 기능을 갖추고 있기도 합니다.

올바른 도구를 찾는 것은 어려운 작업일 수 있지만, 결국 얼마나 경험이 있느냐에 달려있습니다. 가능하면 실험실을 설립하여 최대한 많은 도구를 시험하여 각 도구의 장점과 단점을 기록해 놓으십시오. 각 도구의 README 파일이나 메뉴얼 페이지를 검토하는 것도 잊지 마십시오. 마지막으로 인터넷에서 기사, 단계별 설명서 또는 메일링 리스트에 이르기까지 특정 도구에 대한 추가 정보를 얻으시기 바랍니다.

다음에 설명되는 도구는 사용 가능한 도구 중 일부 예시입니다:

45.2.3.1. nmap을 사용하여 호스트 스캐닝하기

nmap은 Red Hat Enterprise Linux에 포함된 네트워크 배치를 찾아내기 위해 자주 사용되는 도구입니다. nmap은 수년간 사용되어져 왔으며 아마도 가장 흔히 사용되는 정보 수집용 도구입니다. 메뉴얼 페이지를 보시면 옵션과 사용법에 대한 자세한 정보를 찾으실 수 있습니다. 관리자는 네트워크 상에서 nmap을 사용하여 호스트 시스템과 호스트 시스템 상에서 열려진 포트를 찾아낼 수 있습니다.

nmap은 취약성 평가의 첫단계로 사용하기에 모자람이 없는 훌륭한 도구입니다. 네트워크 내의 모든 호스트를 찾고 심지어는 특정 호스트에서 실행 중인 운영 체제를 찾기 위한 옵션도 전달 가능합니다. nmap은 보안 서비스를 사용하고 사용되지 않는 서비스는 멈추는 정책을 수립하는데 좋은 기반이 됩니다.

45.2.3.1.1. nmap 사용법

nmap은 쉘 프롬프트에서 실행 가능합니다. 쉘 프롬프트에서 nmap 명령과 스캔할 시스템의 호스트명이나 IP 주소를 입력하십시오.

```
nmap foo.example.com
```

스캔 결과는 다음과 같이 나타날 것입니다 (호스트의 위치에 따라서 몇 분이 소요될 수도 있습니다):

```
Starting nmap V. 3.50 ( www.insecure.org/nmap/ )
Interesting ports on localhost.localdomain (127.0.0.1):
(The 1591 ports scanned but not shown below are in state: closed)
Port        State       Service
22/tcp      open        ssh
25/tcp      open        smtp
111/tcp     open        sunrpc
443/tcp     open        https
515/tcp     open        printer
950/tcp     open        oftep-rpc
6000/tcp    open        X11

Nmap run completed -- 1 IP address (1 host up) scanned in 71.825 seconds
```

nmap은 대부분의 일반 네트워크 통신 포트를 테스트하여 서비스를 청취하거나 대기 중인 포트가 있는지 찾아냅니다. 이는 관리자가 불필요하거나 사용되지 않는 서비스를 닫는데 매우 유용합니다.

nmap과 관련된 보다 많은 정보를 원하신다면 다음 URL에서 공식 홈페이지를 찾아보시기 바랍니다:

http://www.insecure.org/

45.2.3.2. Nessus

Nessus는 전체 서비스 보안 스캐너입니다. Nessus는 플러그인 구조로 되어있기 때문에 사용자가 자신의 시스템과 네트워크에 맞게 사용자 정의 가능합니다. 다른 스캐너와 마찬가지로 Nessus는 침입 탐지 패턴 데이터베이스에 의존하지만, 다행히도 Nessus는 자주 업데이트됩니다. 완전한 보고 기능, 호스트 스캐닝, 실시간 취약성 검색과 같은 기능을 제공합니다. Nessus처럼 자주 업데이트되고 강력한 도구라 해도 잘못된 결과 (존재하지 않는 취약점을 보고하거나 존재하는 취약점을 발견하지 못하는 결과)를 보고할 가능성이 있습니다.

> 알림
>
> Nessus는 Red Hat Enterprise Linux에 포함되어 있지 않으며 지원되지 않습니다. 이 문서에서는 이 응용 프로그램을 사용하고자 하시는 사용자를 위한 참고 자료로서 언급되었습니다.

Nessus와 관련된 보다 많은 정보를 원하신다면 다음 URL에서 공식 홈페이지를 찾아보시기 바랍니다:

http://www.nessus.org/

45.2.3.3. Nikto

Nikto은 훌륭한 CGI(common gateway interface) 스크립트 스캐너입니다. Nikto는 CGI 취약점을 확인하는 기능을 갖추고 있을 뿐만 아니라 침입 탐지 시스템의 눈에 띄지 않고 찾아내기 어려운 방법을 사용합니다. 이 프로그램은 훌륭한 문서 자료가 함께 나와 있으므로 프로그램을 실행하시기 전에 주의깊게 읽어보시기 바랍니다. CGI 스크립트를 사용하는 웹 서버를 찾으신다면 Nikto를 사용하여 이 서버의 보안을 확인해보실 수 있습니다.

> 알림
>
> Nikto는 Red Hat Enterprise Linux에 포함되어 있지 않으며 지원되지 않습니다. 이 문서에서는 이 응용 프로그램을 사용하고자 하시는 사용자를 위한 참고 자료로서 언급되었습니다.

Nikto에 대한 보다 자세한 정보는 다음 URL에서 찾으실 수 있습니다:

http://www.cirt.net/code/nikto.shtml

45.2.3.4. VLAD the Scanner

VLAD는 Bindview, Inc 사의 RAZOR 팀에서 개발한 스캐너로서 취약성을 검사하는데 사용됩니다. 이 프로그램은 가장 흔한 보안 문제점 중 SANS 최상위 열가지 문제점 (SNMP 문제, 파일 공유 문제점 등)을 찾아냅니다. VLAD는 Nessus 만큼 완전한 기능을 갖추고 있지는 않지만 조사해 볼만한 가치가 있습니다.

알림

VLAD는 Red Hat Enterprise Linux에 포함되어 있지 않으며 지원되지 않습니다. 이 문서에서는 이 응용 프로그램을 사용하고자 하시는 사용자를 위한 참고 자료로서 언급되었습니다.

VLAD에 대한 보다 자세한 정보는 RAZOR 팀 웹사이트인 다음 URL에서 찾으실 수 있습니다:

http://www.bindview.com/Support/Razor/Utilities/

45.2.3.5. 향후 필요한 사항을 미리 준비하십시오.

Depending upon your target and resources, there are many tools available. There are tools for wireless networks, Novell networks, Windows systems, Linux systems, and more. Another essential part of performing assessments may include reviewing physical security, personnel screening, or voice/PBX network assessment. New concepts, such as war walking scanning the perimeter of your enterprise's physical structures for wireless network vulnerabilities are some emerging concepts that you can investigate and, if needed, incorporate into your assessments. Imagination and exposure are the only limits of planning and conducting vulnerability assessments.

45.3. 공격자와 취약점

좋은 보안 정책을 계획하고 구현하기 위해서는 먼저 공격자가 시스템에 침입하고자 결정하게된 계기와 동기를 알아보아야 합니다. 이를 알아보기에 앞서, 우선 공격자를 지칭하는 여러 용어에 대하여 설명해 보겠습니다.

45.3.1. 간략한 해커 역사

오늘날 우리가 알고 있는 해커 (hacker)의 기원은 1960년대 MIT 테크 모델 철도 클럽(TMRC, Tech Model Railroad Club)에서 시작되었습니다. 복잡한 구조의 대형 기차 모형을 제작하는 이 동호회에서는 독창적인 요령이나 문제 해결책을 발견한 동호회 회원을 지칭하는데 해커라는 이름을 사용하기 시작했습니다.

그 이후 해커라는 용어는 컴퓨터 애호가로부터 능력이 뛰어난 프로그래머까지 모든 것을 포함하는 의미를 갖게 되었습니다. 대부분 해커들의 일반적인 특징은 다른 사람의 영향을 받지 않고 스스로 컴퓨터 시스템과 네트워크 기능의 작업 방식에 대하여 자세히 알아보고자 하는 정신입니다. 오픈 소스 소프트웨어 개발자들은 종종 자기 자신과 또한 함께 일하는 동료들을 해커로 지칭하며, 해커 정신을 존경할 것을 요구합니다.

일반적으로 해커들은 정보와 전문적 기술을 탐구하고 이 정보를 공유하는 것이 사회에 대한 해커의 근본 윤리임을 제시하는 해커 윤리 (hacker ethic)를 따릅니다. 이러한 지식을 탐구하는 일환으로 일부 해커들은 컴퓨터 시스템의 보안 침투를 시도하기도 합니다. 이러한 이유로 매체에서는 종종 해커라는 용어를 비도덕적이고 악의를 가지고 범죄를 저지를 생각으로 불법으로 시스템과 네트워크에 침입한 사람을 지칭하는데 사용합니다. 이러한 유형의 컴퓨터 해커에 대한 보다 적절한 용어는 크래커 (cracker) 입니다 — 1980년대 중반에 해커들이 해커와 크래커에 차이를 두기 위해 만들어낸 용어입니다.

45.3.1.1. 다양한 종류의 해커

Within the community of individuals who find and exploit vulnerabilities in systems and networks are several distinct groups. These groups are often described by the shade of hat that they "wear" when performing their security investigations and this shade is indicative of their intent.

white hat 해커란 네트워크와 시스템을 검사하여 얼마나 외부 침입에 취약한지를 연구하는 사람입니다. 일반적으로 white hat 해커는 자신의 시스템이나 보안 검사를 위해 자신을 고용한 고객의 시스템에 침투하여 연구합니다. white hat 해커의 대표적인 예로서는 대학 연구원이나 전문 보안 상담자를 들 수 있습니다.

black hat 해커는 크래커와 동일한 의미입니다. 일반적으로 크래커는 연구나 프로그래밍이 보다는 크래킹 프로그램을 사용하거나 잘 알려진 취약점을 이용하여 시스템에 침입한 후 기밀 정보를 빼내거나 목표 시스템이나 네트워크를 손상시킵니다.

반면 gray hat 해커는 대부분의 경우 white hat 해커의 지식과 의도를 갖추고 있지만, 가끔씩 자신의 지식을 보다 정당하지 못한 의도로 사용하는 사람을 지칭합니다. gray hat 해커는 white hat 해커이지만 가끔씩 자신의 욕심을 채우기 위해 black hat으로 변하는 해커라고 말할 수 있습니다.

Grey hat 해커는 일반적으로 다른 유형의 해커 윤리를 따릅니다. 이 해커 윤리에 따르면 해커가 도둑질이나 기밀을 유포하지 않는 한 시스템에 침입하는 것을 허용합니다. 그러나 일부에서는 시스템에 침입하는 것 자체가 윤리적이지 못하다고 비난합니다.

침입자의 의도에 상관없이 크래커가 시스템 침입에 사용할 취약점을 알아내는 것이 중요합니다. 이 장의 나머지 부분에서는 이러한 취약점에 대하여 중점적으로 설명해 보겠습니다.

45.3.2. 네트워크 보안 위협

다음과 같은 네트워크 설정시 잘못된 설정으로 인해 침입의 위험이 증가될 수 있습니다.

45.3.2.1. 불안정한 구조

잘못 설정된 네트워크는 허가 없는 사용자들이 시스템에 침입할 수 있게 해주는 주요 시작 지점입니다. 로컬 네트워크를 인터넷에 공개해 놓는 것은 마치 우범 지역에서 집의 문을 활짝 열어놓는 것과 같습니다 — 얼마 동안은 아무런 일도 일어나지 않을지 몰라도, 결국에는 누군가가 그 기회를 이용하여 침입할 것입니다.

45.3.2.1.1. 브로드캐스트 네트워크

시스템 관리자는 종종 보안 계획을 구상시 네트워크 하드웨어의 중요성을 간과하는 경우가 있습니다. 허브(hub)와 라우터와 같이 브로드캐스트나 비교환 원칙에 의존하는 단순 하드웨어; 즉, 네트워크를 통하여 수신자 노드로 데이터를 전송시, 허브나 라우터는 수신자 노드가 데이터 패킷을 받아서 프로세스할 때까지 데이터 패킷을 브로드캐스트합니다. 이러한 방법은 지역 네트워크 상에서 허가 없는 사용자나 외부 침입자가 주소 결정 프로토콜 (arp)이나 MAC (media access control) 주소 스푸핑을 사용하여 쉽게 침입할 수 있게 해줍니다.

45.3.2.1.2. 중앙 집중형 서버

또 다른 네트워크 보안 위협으로서 중앙 집중식 컴퓨팅을 들 수 있습니다. 많은 사업체에서 경비를 절감하는 방법으로서 한 개의 강력한 컴퓨터에 모든 서비스를 통합하는 경우가 있습니다. 여러 개의 서버를 설정하는 것 보다 상당히 경비도 절감되고 관리하기에 편한 이점이 있지만, 이러한 중앙 집중형 서버를 사용하는 경우 만일 중앙 서버가 손상되면 네트워크 전체가 정지되거나 또는 데이터 조작이나 도난당하기 쉽습니다. 이러한 경우 침입자는 중앙 서버를 이용하여 전체 네트워크에 접속 가능합니다.

45.3.3. 서버 보안 위협

Server security is as important as network security because servers often hold a great deal of an organization's vital information. If a server is compromised, all of its contents may become available for the cracker to steal or manipulate at will. The following sections detail some of the main issues.

45.3.3.1. 사용되지 않은 서비스와 공개 포트

Red Hat Enterprise Linux를 전체 설치하시면 1000여개에 이르는 응용 프로그램과 라이브러리 패키지가 설치됩니다. 그러나 대부분의 서버 관리자 분들은 배포판에 포함된 모든 개별 패키지를 설치하기 보다는, 대신 여러 서버 응용 프로그램을 포함한 기본 패키지 설치를 선호합니다.

A common occurrence among system administrators is to install the operating system without paying attention to what programs are actually being installed. This can be problematic because unneeded services may be installed, configured with the default settings, and possibly turned on. This can cause unwanted services, such as Telnet, DHCP, or DNS, to run on a server or workstation without the administrator realizing it, which in turn can cause unwanted traffic to the server, or even, a potential pathway into the system for crackers. Refer To 46.2절. "서버 보안" for information on closing ports and disabling unused services.

45.3.3.2. 패치가 설치되지 않은 서비스

기본 설치에 포함된 대부분의 서버 응용 프로그램들은 철저하게 테스트와 검증을 거친 소프트웨어입니다. 여러 해를 거쳐 생산 환경에서 사용되면서, 이 소프트웨어의 코드가 보다 개선되었고 다수의 문제점이 발견되어 수정되었습니다.

그러나 완벽한 소프트웨어란 있을 수 없으며 언제든지 개선할 요소가 있기 마련입니다. 더우기 새로운 소프트웨어는 기대하는 만큼 엄격하게 테스트되지 않는 경우가 종종 있습니다. 그 이유는 이 소프트웨어가 제품 환경에 출시된지 얼마되지 않아서 이거나 또는 다른 서버 소프트웨어 만큼 많이 사용되지 않기 때문입니다.

개발자와 시스템 관리자는 서버 응용 프로그램에서 문제점을 발견한 경우 그 정보를 Bugtraq 메일링 리스트 (http://www.securityfocus.com)나 컴퓨터 비상 대응팀 (Computer Emergency Response Team: CERT) 웹사이트 (http://www.cert.org)와 같은 버그 추적과 보안 관련 웹사이트에 공개합니다. 이러한 방법은 커뮤니티에 보안 취약점을 빠르게 알릴 수 있는 효율 적인 방법이기는 하지만, 시스템 관리자들은 즉시 시스템에 패치를 설치하셔야 합니다. 크래커는 동일한 취약점 추적 서비스를 볼 수 있기 때문에 재빠르게 패치가 설치되지 않은 시스템에 침입하여 정보를 빼내올 가능성이 있기 때문입니다. 훌륭한 시스템 관리자라면 보다 안전한 컴퓨팅 환경을 만들기 위하여 항상 경계하며 지속적으로 버그 (문제점)를 추적하고 적절한 시스템 관리 작업을 수행해야 합니다.

Refer to 45.5절. "보안 업데이트" for more information about keeping a system up-to-date.

45.3.3.3. 부주의한 관리

Administrators who fail to patch their systems are one of the greatest threats to server security. According to the System Administration Network and Security Institute (SANS), the primary cause of computer security vulnerability is to "assign untrained people to maintain security and provide neither the training nor the time to make it possible to do the job."[4] This applies as much to inexperienced administrators as it does to overconfident or amotivated administrators.

[4] Source: http://www.sans.org/security-resources/mistakes.php

일부 관리자들은 서버와 워크스테이션에 패치를 설치하는 것을 잊어버리는 경우가 있는 반면, 다른 어떤 관리자들은 시스템 커널의 로그 메시지나 네트워크 트래픽을 잊고 살펴보지 않는 경우도 있습니다. 또 다른 흔한 실수로 서비스의 기본 암호나 키를 변경하지 않고 그대로 사용하는 것을 들 수 있습니다. 예를 들어 일부 데이터베이스에는 시스템 관리자가 설치 후 암호를 즉시 변경할 것이라는 가정 하에 기본 관리 암호가 할당됩니다. 만일 데이터베이스 관리자가 이 암호를 변경하지 않으면, 경험이 없는 크래커도 잘 알려진 기본 암호를 사용하여 데이터베이스에 관리자 허가를 얻을 수 있습니다. 앞에서 설명된 것은 부주의한 관리가 서버 침입에 미치는 영향을 보여주는 몇 가지 예시일 뿐입니다.

45.3.3.4. 본질적으로 안전하지 못한 서비스

가장 경계가 투철한 기업체에서도 선택한 네트워크 서비스가 원래 안전하지 못하다면 공격당하기 쉽습니다. 예를 들어 신뢰하는 네트워크 하에서 사용될 것이라는 가정 하에서 개발된 서비스가 많습니다; 그러나 이러한 가정은 서비스가 본질적으로 신뢰할 수 없는 — 인터넷 상에서 사용 가능해지면 더 이상 적용되지 않습니다.

이러한 안전하지 못한 네트워크 서비스의 한 예로서 인증을 위해 암호화되지 않는 사용자명과 암호를 요구하는 서비스를 들 수 있습니다. Telnet과 FTP가 이러한 서비스의 두 예입니다. 만일 패킷 스니핑 소프트웨어가 원격 사용자와 이러한 서비스 사이의 트래픽을 감시 중이라면 서비스 사용자명과 암호를 쉽게 가로챌 수 있습니다.

Inherently, such services can also more easily fall prey to what the security industry terms the man-in-the-middle attack. In this type of attack, a cracker redirects network traffic by tricking a cracked name server on the network to point to his machine instead of the intended server. Once someone opens a remote session to the server, the attacker's machine acts as an invisible conduit, sitting quietly between the remote service and the unsuspecting user capturing information. In this way a cracker can gather administrative passwords and raw data without the server or the user realizing it.

Another category of insecure services include network file systems and information services such as NFS or NIS, which are developed explicitly for LAN usage but are, unfortunately, extended to include WANs (for remote users). NFS does not, by default, have any authentication or security mechanisms configured to prevent a cracker from mounting the NFS share and accessing anything contained therein. NIS, as well, has vital information that must be known by every computer on a network, including passwords and file permissions, within a plain text ASCII or DBM (ASCII-derived) database. A cracker who gains access to this database can then access every user account on a network, including the administrator's account.

By default, Red Hat Enterprise Linux is released with all such services turned off. However, since administrators often find themselves forced to use these services, careful configuration is critical. Refer to 46.2절. “서버 보안” for more information about setting up services in a safe manner.

45.3.4. 워크스테이션과 가정용 PC 보안 위협

Workstations and home PCs may not be as prone to attack as networks or servers, but since they often contain sensitive data, such as credit card information, they are targeted by system crackers. Workstations can also be co-opted without the user's knowledge and used by attackers as "slave" machines in coordinated attacks. For these reasons, knowing the vulnerabilities of a workstation can save users the headache of reinstalling the operating system, or worse, recovering from data theft.

45.3.4.1. 너무 단순한 암호

Bad passwords are one of the easiest ways for an attacker to gain access to a system. For more on how to avoid common pitfalls when creating a password, refer to 46.1.3절. "암호 보안".

45.3.4.2. 공격 당하기 쉬운 클라이언트 응용 프로그램

Although an administrator may have a fully secure and patched server, that does not mean remote users are secure when accessing it. For instance, if the server offers Telnet or FTP services over a public network, an attacker can capture the plain text usernames and passwords as they pass over the network, and then use the account information to access the remote user's workstation.

SSH와 같이 보안 프로토콜을 사용하는 경우에도, 원격 사용자가 클라이언트 응용 프로그램을 항상 업데이트하지 않는다면 이러한 공격을 당할 가능성이 있습니다. 예를 들어 v.1 SSH 클라이언트는 악의를 가진 SSH 서버로부터 X-forwarding 공격을 받을 수 있습니다. 클라이언트가 서버에 연결되면, 공격자는 조용히 클라이언트가 네트워크 상에서 누르는 키조합이나 마우스 클릭을 캡쳐할 수 있습니다. 이러한 문제점은 v.2 SSH 프로토콜에서는 고쳐졌지만, 이러한 응용 프로그램의 취약점을 알고 필요한 경우 업데이트하는 것은 사용자의 몫입니다.

46.1절. "워크스테이션 보안" discusses in more detail what steps administrators and home users should take to limit the vulnerability of computer workstations.

45.4. 일반 보안 취약점과 공격

표 45.1. "보안 취약점" details some of the most common exploits and entry points used by intruders to access organizational network resources. Key to these common exploits are the explanations of how they are performed and how administrators can properly safeguard their network against such attacks.

표 45.1. 보안 취약점

보안 취약점	설명	알림
암호가 없거나 디폴트 암호를 사용하는 경우	관리자 암호를 설정하지 않거나 판매업체에서 설정한 디폴트 암호를 사용하는 경우는 라우터나 방화벽에서 자주 발생하지만 일부 리눅스 서비스도 기본 관리자 암호를 그대로 사용하는 경우가 종종 있습니다 (Red Hat Enterprise Linux 5는 제외).	라우터, 방화벽, VPNs, NAS 어플라이언스와 같이 일반적으로 네트워킹 하드웨어와 관련되어 있습니다. 여러 레거시 운영 시스템, 특히 UNIX 및 Windows와 같이 서비스를 일괄해서 판매하는 OSes에서 일반적입니다. 때때로 시스템 관리자가 급하게 권한있는 사용자 계정을 만들고 패스워드를 지정하지 않아 침입자가 사용자 계정을 알아내어 침입할 수 있게 합니다.
디폴트 공유 키	보안 서비스는 종종 개발용이나 평가 테스팅에 사용되는 기본 보안 키를 포함하고 있습니다. 만일 이 보안 키가 변경되지 않은 채 인터넷 상 생산 환경에 저장된다면, 동일한 기본 키를 가진 사용자라면 누구든지 그 공유-키와 그 키에 포함된 기밀 정보를 볼 수 있습니다.	무선 액세스 지점 및 사전 설정된 보안 서버 어플라이언스에서 가장 일반적입니다.
IP 스푸핑 (Spoofing)	A remote machine acts as a node on your local network, finds vulnerabilities with your servers, and installs a backdoor program or Trojan horse	스푸핑은 목표 시스템에 연결하기 위해 침입자가 TCP/IP SYN-ACK 번호를 예측해야 하므로 매우 어렵습니다. 그러나 크래커는 여러 다른 도구를

보안 취약점	설명	알림
	to gain control over your network resources.	사용하여 이러한 취약점을 공격 가능합니다. rsh, telnet, FTP 등과 같이 소스 기반 인증을 사용하는 서비스의 목표 시스템에 따라 ssh이나 SSL/TLS에서 사용되는 암호화된 인증 방식이나 PKI와 비교했을 때 안전하지 않으므로 사용을 권장하지 않습니다.
도청 (Eavesdropping)	네트워크 상 두 개의 활성 노드 사이의 통신 선로 전선에 접속하여 정보를 빼내는 행위.	이러한 종류의 공격은 Telnet, FTP, HTTP 전송과 같이 평문으로 된 전송 프로토콜에서 자주 발생합니다. 원격 침입자는 시스템에 침입하기 위하여 LAN 상에 있는 목표 시스템에 접속해야 합니다; 일반적으로 크래커는 IP 스푸핑 또는 man-in-the-middle 공격과 같이 LAN 상의 시스템에 침입하기 위하여 능동적 공격을 사용합니다. 보안 방지법에는 패스워드 스누핑(snooping)을 방지하기 위해 암호화 키 교환, 일회용 패스워드 또는 암호화된 인증 방식과 같은 서비스가 있습니다; 네트워크 자료를 전송하는 동안 철저한 암호화 과정을 거치실 것을 권장합니다.
서비스 취약점	침입자는 인터넷 상에서 실행되는 서비스에서 허점을 찾아서 그 시스템에 칩입하여 저장된 정보를 가로챌 수 있습니다. 또한 동일한 네트워크 상에 위치한 다른 시스템에 침입하는 것도 가능합니다.	HTTP-based services such as CGI are vulnerable to remote command execution and even interactive shell access. Even if the HTTP service runs as a non-privileged user such as "nobody", information such as configuration files and network maps can be read, or the attacker can start a denial of service attack which drains system resources or renders it unavailable to other users. Services sometimes can have vulnerabilities that go unnoticed during development and testing; these vulnerabilities (such as buffer overflows, where attackers crash a service using arbitrary values that fill the memory buffer of an application, giving the attacker an interactive command prompt from which they may execute arbitrary commands) can give complete administrative control to an attacker. 관리자는 루트 사용자로서 서비스를 실행해서는 안되며, 제조업체나 CERT 및 CVE와 같은 보안 기관의

보안 취약점	설명	알림
		응용 프로그램에 대한 패치와 에라타 업데이트를 빈틈없이 해야 합니다.
응용 프로그램 보안 취약점	Attackers find faults in desktop and workstation applications (such as e-mail clients) and execute arbitrary code, implant Trojan horses for future compromise, or crash systems. Further exploitation can occur if the compromised workstation has administrative privileges on the rest of the network.	일반 직장인들은 보안에 대한 경험이나 지식이 부족하기 때문에 워크스테이션이나 데스크탑은 보안 취약성 공격의 목표가 되기 쉽습니다; 따라서 개인마다 허가를 받지 않은 소프트웨어나 검열받지 않은 이메일 첨부 파일을 열 경우 그 위험성을 알려야 합니다. 이메일 클라이언트 소프트웨어가 자동으로 첨부 파일을 열거나 실행하지 않게하는 안전 장치를 실행할 수 도 있습니다. 또한 Red Hat Network를 통한 워크스테이션 소프트웨어 자동 업데이트나 다른 시스템 관리 서비스를 사용하시면 보다 쉽게 다중 보안 작업을 실행하실 수 있습니다.
서비스 거부 (DoS) 공격	Attacker or group of attackers coordinate against an organization's network or server resources by sending unauthorized packets to the target host (either server, router, or workstation). This forces the resource to become unavailable to legitimate users.	미국에서 보고된 DoS 케이스는 대부분 2000년에 발생하였습니다. 정보 소통량이 많은 일부 상업 사이트와 정부 사이트를 표적으로 삼아 이미 해킹당한 여러 시스템을 좀비 또는 방향 변경 브로드캐스트 노드처럼 사용하여 그 시스템의 고대역폭 연결을 이용하여 표적 사이트에 핑플루드 (ping flood) 공격을 가한 후 접속 불능 상태로 만들었습니다. 일반적으로 소스 패킷은 위장되어 다시 브로드캐스트되었기 때문에, 공격의 원소스를 찾아 내기는 어렵습니다. iptables 및 snort와 같은 Network IDSes를 사용한 ingress filtering (IETF rfc2267)이 개선되면서, 관리자는 분산된 DoS 공격을 추격하고 방지할 수 있습니다.

45.5. 보안 업데이트

보안 상 허점이 발견된다면, 적절한 소프트웨어를 업데이트하여 보안 위험을 최소화 시켜야 합니다. 해당 소프트웨어가 현재 Red Hat Enterprise Linux 배포판에 포함된 지원 가능한 패키지의 일부라면, Red Hat, Inc.는 최대한 빨리 보안 취약점을 고칠 수 있는 패키지를 업데이트하여 보내드릴 것입니다. 종종 보안 취약점을 공개시 그에 상응하는 패치 (문제를 해결할 수 있는 소스 코드)를 함께 출시합니다. Red Hat QA 팀은 이 패치를 Red Hat Enterprise Linux 패키지에 적용하여 테스팅을 마친 후 에라타 업데이트로 출시할 것입니다. 그러나 패치가 없는 보안 취약점을 공개시, Red Hat 개발자는 소프트웨어 관리자와 함께 작업하여 문제 해결책을 찾아낼 것입니다. 일단 문제가 해결되면, 패키지를 테스트한 후 에라타 업데이트로 배포합니다.

고객 시스템에 사용된 소프트웨어에 대한 에라타 업데이트가 발표된다면, 최대한 빨리 패키지를 업데이트하여 시스템 보안 위험을 최소화하시기 바랍니다.

45.5.1. 패키지 업데이트하기

시스템에 소프트웨어를 업데이트하실 경우, 신뢰할 수 있는 소스에서 업데이트를 다운로드 받으셔야 합니다. 누구든 문제 해결 패치와 동일한 버전 번호을 가졌지만 또 다른 보안상 허점을 제공하는 패키지를 재구축하여 인터넷에 올려놓을 가능성도 있습니다. 이러한 경우, 파일을 기존 RPM과 비교 검증하는 것과 같은 보안 대책을 사용하여도 보안 허점이 발견되지 않습니다. 따라서 Red Hat, Inc.와 같은 신뢰할 수 있는 소스에서 RPM을 다운로드 받으시는 것이 매우 중요합니다. 패키지의 무결성을 검증하기 위해 패키지의 서명을 확인하시기 바랍니다.

다음과 같은 두가지 방법을 사용하여 Red Hat 보안 에라타 업데이트를 받으실 수 있습니다:

1. Red Hat Network에 기록 및 다운로드 가능

2. Red Hat 에라타 웹사이트에 기록되었으며 링트되어 있지 않음

> **알림**
>
> Red Hat Enterprise Linux 부터는 Red Hat Network를 통해서만 업데이트된 패키지를 다운받을 수 있습니다. Red Hat 에라타 웹페이지에 업데이트 정보는 있지만, 실제 패키지를 다운받을 수는 없습니다.

45.5.1.1. Using Automatic Updates with RHN Classic

> **Warning: Deprecate Feature**
>
> Automatic system updates are only available using RHN Classic, which basis subscription consumption on access to content repository channels. RHN Classic is available as a convenience for customer environments with legacy systems which have not updated to Certificate-Based Red Hat Network.
>
> The update and content stream is different for Certificate-Based Red Hat Network, so automatic updates are not used.
>
> The new Certificate-Based Red Hat Network and the differences between Certificate-Based Red Hat Network and RHN Classic are described in 14장. Product Subscriptions and Entitlements.

RHN Classic allows the majority of the update process to be automated. It determines which RPM packages are necessary for the system, downloads them from a secure repository, verifies the RPM signature to make sure they have not been tampered with, and updates them. The package install can occur immediately or can be scheduled during a certain time period.

RHN Classic requires a system profile for each machine, which contains hardware and software information about the system. This information is kept confidential and is not given to anyone else. It is only used to determine which errata updates are applicable to each system, and, without it, RHN Classic can not determine whether a given system needs updates. When a security errata (or any type of errata) is released, RHN Classic sends an email with a description of the errata as well as a list of systems which are affected. To apply the update, use the Red Hat Update Agent or schedule the package to be updated through the RHN Classic Subscription Management area of the Customer Portal.

중요

Before installing any security errata, be sure to read any special instructions contained in the errata report and execute them accordingly. Refer to 45.5.1.5절. "변경 사항 적용하기" for general instructions about applying the changes made by an errata update.

45.5.1.2. Red Hat 에라타 웹사이트 이용 방법

When security errata reports are released, they are published on the Red Hat Errata website available at http://www.redhat.com/security/. From this page, select the product and version for your system, and then select security at the top of the page to display only Red Hat Enterprise Linux Security Advisories. If the synopsis of one of the advisories describes a package used on your system, click on the synopsis for more details.

자세한 정보 페이지에는 보안 허점에 대한 정보와 함께 보안 상 문제점을 고치기 위한 패키지 업데이트 및 필요한 작업 실행 방법에 대한 지시 사항이 있습니다.

To download the updated package(s), click on the link to login to Red Hat Network, click the package name(s) and save to the hard drive. It is highly recommended that you create a new directory, such as /tmp/updates, and save all the downloaded packages to it.

45.5.1.3. 패키지 서명 검증하기

모든 Red Hat Enterprise Linux 패키지는 Red Hat, Inc. GPG키로 서명되었습니다. GPG는 GNU Privacy Guard (GnuPG)의 줄임말로서 배포 파일의 인증을 확인하는데 사용되는 자유 소프트웨어 패키지를 말합니다. 예를 들면 Red Hat은 비밀키 (개인키)를 사용하여 패키지를 잠근 후 공개키를 사용하여 패키지를 잠금 해제 후 검증합니다. 만일 RPM 검증 과정에서 Red Hat에서 배포한 공개키가 비밀키와 일치하지 않는다면, 패키지가 변경되었을 수 있으므로 신뢰할 수 없습니다.

Red Hat Enterprise Linux에서 RPM 유틸리티를 사용하면 RPM 패키지를 설치하기 전에 자동적으로 패키지의 GPG 서명을 검증합니다. 만일 Red Hat GPG 키를 아직 설치하지 않으셨다면, Red Hat Enterprise Linux 설치 CD-ROM과 같이 안전한 곳에서 GPG 키를 받아 설치하시기 바랍니다.

Assuming the CD-ROM is mounted in /mnt/cdrom, use the following command to import it into the keyring (a database of trusted keys on the system):

```
rpm --import /mnt/cdrom/RPM-GPG-KEY-redhat-release
```

RPM 검증을 위해 설치된 모든 키 목록을 보시려면, 다음 명령을 실행하십시오:

```
rpm -qa gpg-pubkey*
```

Red Hat키는 다음과 같이 출력됩니다:

```
gpg-pubkey-37017186-45761324
```

To display details about a specific key, use the rpm -qi command followed by the output from the previous command, as in this example:

```
rpm -qi gpg-pubkey-37017186-45761324
```

RPM 파일들을 설치하시기 전에 반드시 파일의 서명을 검증하여 파일이 Red Hat, Inc.에서 출시한 패키지에서 변경되지 않았음을 확인하시기 바랍니다. 다운로드 받은 패키지들을 동시에 모두 검증하시려면, 다음 명령을 입력하십시오:

```
rpm -K /tmp/updates/*.rpm
```

For each package, if the GPG key verifies successfully, the command returns gpg OK. If it doesn't, make sure you are using the correct Red Hat public key, as well as verifying the source of the content. Packages that do not pass GPG verifications should not be installed, as they may have been altered by a third party.

GPG 키 검색과 에라타 리포트와 관련된 모든 패키지를 다운로드 받으셨다면, 쉘 프롬프트에서 루트로 로그인 하신 후 패키지를 설치하시기 바랍니다.

45.5.1.4. 서명된 패키지 설치하기

다음 명령을 입력하여 패키지를 안전하게 설치하실 수 있습니다 (커널 패키지 제외):

```
rpm -Uvh /tmp/updates/*.rpm
```

커널 패키지의 경우 다음 명령을 사용하시기 바랍니다:

```
rpm -ivh /tmp/updates/<kernel-package>
```

Replace <kernel-package> in the previous example with the name of the kernel RPM.

새로운 커널을 사용하여 컴퓨터가 안전하게 재부팅한 후에는 다음 명령을 사용하여 이전 커널을 제거하셔도 됩니다:

```
rpm -e <old-kernel-package>
```

Replace <old-kernel-package> in the previous example with the name of the older kernel RPM.

알림

이전 커널을 반드시 삭제할 필요는 없습니다. 기본 부트로더인 GRUB을 이용하여 여러 커널을 설치한 후 부팅시 어느 커널을 부팅할 지 결정할 수 있습니다.

중요

Before installing any security errata, be sure to read any special instructions contained in the errata report and execute them accordingly. Refer to 45.5.1.5절. "변경 사항 적용하기" for general instructions about applying the changes made by an errata update.

45.5.1.5. 변경 사항 적용하기

Red Hat Network 또는 Red Hat 에라타 웹사이트를 통해 보안 에라타를 다운로드 받고 설치하신 후, 이전 소프트웨어 사용을 중단하시고 새로운 소프트웨어를 사용하셔야 합니다. 업데이트된 소프트웨어의 종류에 따라서 이 방법이 달라집니다. 다음 목록은 일반 소프트웨어 종류와 패키지 업그레이드 후 업데이트된 버전을 사용하는 방법을 보여줍니다.

알림

일반적으로 시스템을 재부팅하시는 것이 최신 버전의 소프트웨어 패키지를 사용할 수 있는 가장 쉬운 방법입니다; 그러나 시스템 관리자가 시스템을 재부팅할 수 없는 경우도 있습니다.

응용 프로그램

사용자-공간 응용 프로그램이란 시스템 사용자가 시작할 수 있는 모든 프로그램을 의미합니다. 일반적으로 이러한 응용 프로그램은 사용자나 스크립트 또는 자동화 작업 유틸리티에 의해 시작되어 사용되며, 오랫동안 실행되지 않습니다.

이러한 사용자-공간 응용 프로그램이 업데이트되면, 시스템 상 해당 응용 프로그램의 인스턴스를 멈춘 후 프로그램을 다시 시작하여 업데이트된 버전을 사용하십시오.

커널

커널은 Red Hat Enterprise Linux 운영 체제의 핵심 소프트웨어 요소입니다. 커널은 메모리, 프로세서 및 주변 기기의 사용을 관리할 뿐만 아니라 모든 작업을 계획하고 관리합니다.

커널의 중심적인 역할 때문에, 컴퓨터를 멈추지 않고서는 커널을 재시작할 수 없습니다. 따라서 커널을 업데이트한 경우 시스템을 재부팅하지 않고서는 업데이트된 버전의 커널을 사용할 수 없습니다.

공유 라이브러리

Shared libraries are units of code, such as glibc, which are used by a number of applications and services. Applications utilizing a shared library typically load the shared code when the application is initialized, so any applications using the updated library must be halted and relaunched.

To determine which running applications link against a particular library, use the lsof command as in the following example:

```
lsof /usr/lib/libwrap.so*
```

This command returns a list of all the running programs which use TCP wrappers for host access control. Therefore, any program listed must be halted and relaunched if the tcp_wrappers package is updated.

SysV 서비스

SysV services are persistent server programs launched during the boot process. Examples of SysV services include sshd, vsftpd, and xinetd.

Because these programs usually persist in memory as long as the machine is booted, each updated SysV service must be halted and relaunched after the package is upgraded. This can be done using the Services Configuration Tool or by logging into a root shell prompt and issuing the /sbin/service command as in the following example:

```
service <service-name> restart
```

In the previous example, replace <service-name> with the name of the service, such as sshd.

Refer to 16장. 네트워크 설정 for more information on the Services Configuration Tool.

xinetd Services

Services controlled by the xinetd super service only run when a there is an active connection. Examples of services controlled by xinetd include Telnet, IMAP, and POP3.

Because new instances of these services are launched by xinetd each time a new request is received, connections that occur after an upgrade are handled by the updated software. However, if there are active connections at the time the xinetd controlled service is upgraded, they are serviced by the older version of the software.

To kill off older instances of a particular xinetd controlled service, upgrade the package for the service then halt all processes currently running. To determine if the process is running, use the ps command and then use the kill or killall command to halt current instances of the service.

For example, if security errata imap packages are released, upgrade the packages, then type the following command as root into a shell prompt:

```
ps -aux | grep imap
```

이 명령은 모든 활성 IMAP 세션을 보여줍니다. 이제 다음 명령을 사용하여 개별 세션을 종료시킬 수 있습니다:

```
kill <PID>
```

이 세션을 종료하지 못하셨을 경우, 다음 명령을 사용하십시오:

```
kill -9 <PID>
```

In the previous examples, replace <PID> with the process identification number (found in the second column of the ps command) for an IMAP session.

모든 활성 IMAP 세션을 종료하시려면, 다음 명령을 입력하십시오:

```
killall imapd
```

네트워크 보안

46.1. 워크스테이션 보안

리눅스 환경 보안은 워크스테이션을 안전하게 보호하는 것부터 시작합니다. 개인용 컴퓨터를 잠그거나 엔터프라이즈 시스템을 보호하던 간에 훌륭한 보안 정책은 개인 컴퓨터에서 시작됩니다. 결국 가장 보안이 약한 컴퓨터 한 대가 전체 컴퓨터 네트워크의 보안을 결정합니다.

46.1.1. 워크스테이션 보안 평가하기

Red Hat Enterprise Linux 워크스테이션의 보안을 평가하실 때 다음과 같은 사항을 고려하십시오:

- BIOS와 부트로더 보안 — 허가가 없는 사용자가 컴퓨터에 직접 접근하여 암호가 필요없는 단독 사용자 모드나 복구 모드로 부팅할 가능성이 있는가?
- 암호 보안 — 컴퓨터의 사용자 계정 암호가 얼마나 안전한가?
- 관리 제어 — 시스템 계정을 가진 사람은 누구이며 얼마나 많은 관리 제어권을 가지고 있는가?
- 사용 가능한 네트워크 서비스 — 네트워크에서 오는 요청을 청취하고 있는 서비스는 무엇이며 이 서비스가 실행될 필요가 있는가?
- 개인 방화벽 — 어떠한 유형의 방화벽이 필요한가?
- 보안 강화된 통신 도구 — 워크스테이션 간에 통신하는데 어떠한 도구를 사용할 것이며 어떠한 도구의 사용을 피할 것인가?

46.1.2. BIOS와 부트로더 보안

BIOS (또는 BIOS 동급) 및 부트로더에 암호 보호 기능을 사용함으로서 시스템에 물리적으로 접근할 수 있는 사용자 중 허가가 없는 사용자가 이동 매체를 사용하여 부팅하거나 단독 사용자 모드를 사용하여 루트 권한을 갖게되는 것을 방지할 수 있습니다. 그러나 이러한 공격에 대비한 보안 척도는 워크스테이션에 저장된 정보의 중요성과 컴퓨터가 어디에 위치하고 있는지에 따라서 달라집니다.

For example, if a machine is used in a trade show and contains no sensitive information, then it may not be critical to prevent such attacks. However, if an employee's laptop with private, unencrypted SSH keys for the corporate network is left unattended at that same trade show, it could lead to a major security breach with ramifications for the entire company.

다른 한편으로 워크스테이션이 오직 허가가 있는 사용자와 신뢰할 수 있는 사용자만 접근할 수 있는 곳에 위치한다면 BIOS나 부트로더를 보안 조치를 취하는 것이 그리 필요하지 않을 수 도 있습니다.

46.1.2.1. BIOS 암호

다음은 컴퓨터의 BIOS에 암호를 지정하여 보호해야하는 두 가지 중요한 이유입니다[1]:

1. BIOS 설정을 변경하지 못하도록 함 — 만일 침입자가 BIOS에 들어갈 수 있다면 디스켓이나 CD-ROM을 사용하여 부팅하도록 설정할 수 있습니다. 이렇게 되면 침입자는 복구 모드나 단일 사용자 모드로 들어가서 시스템에 바이러스 프로그램을 퍼트리거나 중요한 자료를 복사할 수 있는 기회를 갖게됩니다.

2. 시스템 부팅 방지하기 — 일부 BIOS는 암호를 입력해야 부팅할 수 있도록 하는 기능을 제공합니다. 이 기능이 활성화된 경우 BIOS가 부트로더를 시작하기 전에 침입자는 암호를 입력하여야 시스템을 부팅 가능합니다.

Because the methods for setting a BIOS password vary between computer manufacturers, consult the computer's manual for specific instructions.

BIOS 암호를 잊으셨다면 마더보드에서 점퍼(jumper)를 사용하거나 CMOS 배터리를 빼서 재설정하실 수 있습니다. 이러한 이유로 가능한 컴퓨터 케이스를 잠그시는 것이 보안상 좋습니다. 그러나 CMOS 배터리 빼기를 시도하시기 전에 컴퓨터나 마더보드에 함께 포함된 설명서를 먼저 참조하시기 바랍니다.

46.1.2.1.1. x86가 아닌 플랫폼에서의 보안

다른 구조는 다른 프로그램을 사용하여 x86 시스템 상 BIOS와 유사한 하위-레벨 작업을 수행합니다. 예를 들면 Intel® Itanium™ 컴퓨터는 EFI (Extensible Firmware Interface) 쉘을 사용합니다.

For instructions on password protecting BIOS-like programs on other architectures, refer to the manufacturer's instructions.

46.1.2.2. 부트로더 암호

리눅스 부트로더를 암호를 사용하여 보호해야하는 중요한 이유는 다음과 같습니다:

1. 단독 사용자 모드로 들어가지 못하게 함 — 만일 침입자가 단독 사용자 모드로 들어갈 수 있다면, 그 침입자는 루트 암호를 입력하지 않고서도 루트 사용자 권한을 갖게 됩니다.

2. Preventing Access to the GRUB Console — If the machine uses GRUB as its boot loader, an attacker can use the GRUB editor interface to change its configuration or to gather information using the cat command.

3. 안전하지 않은 운영 체제에 접근하는 것을 방지함 — 이중 부트 시스템에서 침입자는 부팅시 접근 제어와 파일 권한과 같은 기능을 사용하지 않는 DOS와 같은 운영 체제를 선택할 가능성이 있습니다.

Red Hat Enterprise Linux ships with the GRUB boot loader on the x86 platform. For a detailed look at GRUB, refer to the Red Hat Installation Guide.

46.1.2.2.1. GRUB 암호 보호

You can configure GRUB to address the first two issues listed in 46.1.2.2절. "부트로더 암호" by adding a password directive to its configuration file. To do this, first choose a strong password, open a shell, log in as root, and then type the following command:

[1] 시스템 BIOS는 제조업체에 따라 다르기 때문에 일부 BIOS는 두 가지 중 한가지 유형의 암호 보호를 지원하지만 다른 유형은 지원하지 않는 반면 다른 BIOS는 두가지 유형 모두 지원하지 않을 수 도 있습니다

```
grub-md5-crypt
```

When prompted, type the GRUB password and press Enter. This returns an MD5 hash of the password.

Next, edit the GRUB configuration file /boot/grub/grub.conf. Open the file and below the timeout line in the main section of the document, add the following line:

```
password --md5 <password-hash>
```

Replace <password-hash> with the value returned by /sbin/grub-md5-crypt[2].

The next time the system boots, the GRUB menu prevents access to the editor or command interface without first pressing p followed by the GRUB password.

Unfortunately, this solution does not prevent an attacker from booting into an insecure operating system in a dual-boot environment. For this, a different part of the /boot/grub/grub.conf file must be edited.

Look for the title line of the operating system that you want to secure, and add a line with the lock directive immediately beneath it.

DOS 시스템에서 이 절은 다음과 같이 유사하게 시작해야 합니다:

```
title DOS lock
```

경고

A password line must be present in the main section of the /boot/grub/grub.conf file for this method to work properly. Otherwise, an attacker can access the GRUB editor interface and remove the lock line.

To create a different password for a particular kernel or operating system, add a lock line to the stanza, followed by a password line.

각 절은 다음 예시와 유사한 줄로 시작하는 고유 암호로 보호됩니다:

```
title DOS lock password --md5 <password-hash>
```

46.1.3. 암호 보안

Passwords are the primary method that Red Hat Enterprise Linux uses to verify a user's identity. This is why password security is so important for protection of the user, the workstation, and the network.

보안을 위하여 설치 프로그램은 시스템이 MD5 (Message-Digest Algorithm)와 섀도우 암호를 사용하도록 설정합니다. 이 설정을 변경하지 마시길 적극 권장합니다.

[2] GRUB also accepts unencrypted passwords, but it is recommended that an MD5 hash be used for added security.

만일 설치 과정에서 MD5 암호를 선택하지 않으시면 이전 DES (Data Encryption Standard) 형식이 사용됩니다. 이 형식은 암호가 문자와 숫자를 조합하여 (구두점과 그외 특수 문자를 제외하고) 8 글자까지만 허용하며 낮은 56 비트 수준의 암호화만 제공합니다.

If shadow passwords are deselected during installation, all passwords are stored as a one-way hash in the world-readable /etc/passwd file, which makes the system vulnerable to offline password cracking attacks. If an intruder can gain access to the machine as a regular user, they can copy the /etc/passwd file to their own machine and run any number of password cracking programs against it. If there is an insecure password in the file, it is only a matter of time before the password cracker discovers it.

Shadow passwords eliminate this type of attack by storing the password hashes in the file /etc/shadow, which is readable only by the root user.

This forces a potential attacker to attempt password cracking remotely by logging into a network service on the machine, such as SSH or FTP. This sort of brute-force attack is much slower and leaves an obvious trail as hundreds of failed login attempts are written to system files. Of course, if the cracker starts an attack in the middle of the night on a system with weak passwords, the cracker may have gained access before dawn and edited the log files to cover their tracks.

In addition to format and storage considerations is the issue of content. The single most important thing a user can do to protect their account against a password cracking attack is create a strong password.

46.1.3.1. 추측하기 힘든 암호 생성하기

보안 암호를 생성하시려면 다음과 같은 사항을 따르시기 바랍니다:

- 단어나 번호만 사용하시면 안됩니다 — 암호를 생성시 번호나 숫자만 사용하시면 안됩니다.

  안전하지 못한 암호의 예시:

  - 8675309
  - juan
  - hackme

- 알아내기 쉬운 단어를 사용하지 마십시오 — 고유 명사, 사전에 나온 단어나 심지어는 TV 프로그램이나 소설책 이름은 번호가 포함되더라도 암호로 사용하지 마십시오.

  안전하지 못한 암호의 예시:

  - john1
  - DS-9
  - mentat123

- 외국어로된 단어를 사용하지 마십시오 — 암호 크래킹 프로그램은 종종 외국어 사전을 포함한 단어 목록을 사용하여 암호를 찾아냅니다. 따라서 외국어를 보안 암호로 사용하시는 것은 소용이 없습니다.

  안전하지 못한 암호의 예시:

  - cheguevara
  - bienvenido1

  - 1dumbKopf

- 해커 용어를 사용하지 마십시오 — 해커 용어 — l337 (LEET) 용어 — 를 암호로 사용하니까 자신이 수준 높은 해커라고 생각하십니까? 결코 그렇지 않습니다. 많은 용어집에 LEET 용어가 포함되어 있어 암호를 알아내기 쉽습니다.

  안전하지 못한 암호의 예시:

  - H4X0R
  - 1337

- 개인 정보를 사용하지 마십시오 — 암호에 개인 정보를 사용하지 마십시오. 만일 침입자가 신상 정보를 알고 있다면 쉽게 암호를 추측해낼 수 있습니다. 다음은 암호 생성시 피하셔야 할 정보 유형 목록입니다:

  안전하지 못한 암호의 예시:

  - 이름
  - 애완 동물 이름
  - 가족 이름
  - 생년월일
  - 전화번호 또는 우편번호

- 알아내기 쉬운 단어를 역순으로 사용하지 마십시오 — 잘 고안된 암호 추측 프로그램은 항상 일반 단어를 역순으로 확인해 봅니다. 따라서 알아내기 쉬운 단어를 역순으로 한다고 해서 더 안전해지는 것은 아닙니다.

  안전하지 못한 암호의 예시:

  - R0X4H
  - nauj
  - 9-DS

- 암호를 적어두지 마십시오 — 암호를 종이에 적어두지 마십시오. 머리속으로 외우시는 것이 훨씬 안전합니다.

- 모든 기계에 똑같은 암호를 사용하지 마십시오 — 각 기계마다 별개의 암호를 만드시는 것이 중요합니다. 이렇게 하심으로서 만일 한 시스템에서 보안 침입이 발생할 경우 모든 시스템이 즉시 위험에 처할 가능성이 낮아집니다.

다음의 사항을 참조하여 추측하기 힘든 암호를 생성하시기 바랍니다:

- 최소한 8자 이상으로 암호를 만드십시오 — 암호가 길면 길수록 좋습니다. MD5 암호를 사용하시려면 암호는 15자 이상이어야 합니다. DES 암호를 사용하시는 경우 최대 길이 (8 글자)를 사용하십시오.

- 대문자와 소문자를 함께 사용하십시오 — Red Hat Enterprise Linux는 대/소문자를 구별합니다. 따라서 대문자와 소문자를 함께 사용하시면 암호를 추측하기 더욱 힘들어져 보안을 강화합니다.

- 글자와 숫자 조합을 사용하십시오 — 암호에 숫자를 추가시, 특히 암호 첫부분이나 마지막이 아닌 중간에 추가하시면 암호 보안을 강화할 수 있습니다.

- Include Non-Alphanumeric Characters — Special characters such as &, $, and > can greatly improve the strength of a password (this is not possible if using DES passwords).

- 기억할 수 있는 암호를 선택하십시오 — 가장 좋은 암호를 생성하신다고 해도 만일 여러분이 기억하지 못하신다면 무슨 소용입니까?; 암호를 기억하기 위해 머리 글자를 사용하거나 기억을 돕는 장치를 사용하시기 바랍니다.

앞서 설명된 모든 규칙을 사용하여 나쁜 암호의 약점을 피하여 좋은 암호를 만드는 작업이 힘들게 느껴질 수도 있습니다. 다행히 기억하기 쉽고 안전한 암호를 쉽게 생성할 수 있는 방법이 있습니다.

46.1.3.1.1. 보안 암호 생성 방법

여러가지 방법을 사용하여 보안 암호를 생성할 수 있습니다. 가장 자주 사용되는 방법은 다음과 같이 긴 이름이나 구문의 각 단어의 머리글자를 합성하여 만든 약어를 사용하는 방법입니다. 예:

- 기억하기 쉬운 문장을 생각해 내십시오. 예를 들면:

  "over the river and through the woods, to grandmother's house we go."

- 다음으로 이 문장의 각 단어의 머리글자 (구두점 포함)를 합성하여 약어로 만드십시오.

  otrattw,tghwg.

- Add complexity by substituting numbers and symbols for letters in the acronym. For example, substitute 7 for t and the at symbol (@) for a:

  o7r@77w,7ghwg.

- Add more complexity by capitalizing at least one letter, such as H.

  o7r@77w,7gHwg.

- 마지막으로 이 책에서 나온 예시 암호를 절대 여러분의 시스템에 사용하시면 안됩니다.

보안 암호를 만드는 것이 매우 중요하지만 만들어진 암호를 적절하게 관리하는 것 또한 중요합니다. 특히 대기업 시스템 관리자의 경우 더욱 그러합니다. 다음 부분에서는 기업체에서 사용자 암호를 생성하여 관리하는 방법에 대하여 설명해 보겠습니다.

46.1.3.2. 기업체에서 사용자 암호 생성하기

만일 기업체내 사용자 수가 매우 많다면, 시스템 관리자는 사용자에게 좋은 암호를 사용하도록 다음과 같은 두가지 방법을 사용할 수 있습니다. 첫번째 방법은 관리자가 직접 사용자 암호를 생성하는 것이며, 다른 방법은 사용자가 스스로 자신의 암호를 생성하도록 하지만 관리자가 그 암호가 안전한지 확인한 후 사용을 허가하는 방법입니다.

사용자를 위해 암호를 생성하시면 보안 암호를 보증할 수 있지만, 회사가 커지면서 너무 많은 시간과 노력이 낭비될 수 있습니다. 또한 사용자가 자신의 암호를 적어서 기억할 위험도 커집니다.

이러한 이유로 대다수 시스템 관리자 분들은 사용자가 스스로 자신의 암호를 생성하게 하지만, 그 암호가 사용해도 좋은지 계속적으로 확인하고 어떠한 경우에는 사용자가 주기적으로 암호를 변경하도록 암호 유효 기간을 설정하시는 것도 좋습니다.

46.1.3.2.1. 좋은 암호 사용을 엄격히 시행하기

To protect the network from intrusion it is a good idea for system administrators to verify that the passwords used within an organization are strong ones. When users are asked to create or change passwords, they can use the command line application passwd, which is Pluggable Authentication Manager (PAM) aware and therefore checks to see if the password is too short or otherwise easy to crack. This check is performed using the pam_cracklib.so PAM module. Since PAM is customizable, it is possible to add more password integrity checkers, such as pam_passwdqc (available from http://www.openwall.com/passwdqc/) or to write a new module. For a list of available PAM modules, refer to http://www.kernel.org/pub/linux/libs/pam/modules.html. For more information about PAM, refer to 46.4절. "PAM (Pluggable Authentication Modules)" .

암호 생성시 사용되는 암호 확인 기능은 암호 크래킹 프로그램을 실행하여 암호를 알아내는 것만큼 효율적으로 잘못 고안된 암호를 발견하지 못 합니다.

비록 운영 체제에 포함되어 있지는 않지만, 많은 암호 크래킹 프로그램을 Red Hat Enterprise Linux에서 실행 가능합니다. 다음은 가장 자주 사용되는 암호 크래킹 프로그램 목록입니다:

알림

다음 도구들은 Red Hat Enterprise Linux에 포함되지 않으며 따라서 Red Hat, Inc에서는 어떠한 지원도 제공하지 않습니다.

- John The Ripper — A fast and flexible password cracking program. It allows the use of multiple word lists and is capable of brute-force password cracking. It is available online at http://www.openwall.com/john/.
- Crack — Perhaps the most well known password cracking software, Crack is also very fast, though not as easy to use as John The Ripper. It can be found online at http://www.openwall.com/john/.
- Slurpie — Slurpie is similar to John The Ripper and Crack, but it is designed to run on multiple computers simultaneously, creating a distributed password cracking attack. It can be found along with a number of other distributed attack security evaluation tools online at http://www.ussrback.com/distributed.htm.

경고

회사에서 암호 크래킹 프로그램을 사용하시기 전에 항상 먼저 서면으로 허가를 받으셔야 합니다.

46.1.3.2.2. 암호 유효 기간 설정

Password aging is another technique used by system administrators to defend against bad passwords within an organization. Password aging means that after a specified period (usually 90 days), the user is prompted to create a new password. The theory behind this is that if a user is forced to change their password periodically, a cracked password is only useful to an intruder for a limited amount of time. The downside to password aging, however, is that users are more likely to write their passwords down.

There are two primary programs used to specify password aging under Red Hat Enterprise Linux: the chage command or the graphical User Manager (system-config-users) application.

The -M option of the chage command specifies the maximum number of days the password is valid. For example, to set a user's password to expire in 90 days, use the following command:

```
chage -M 90 <username>
```

In the above command, replace <username> with the name of the user. To disable password expiration, it is traditional to use a value of 99999 after the -M option (this equates to a little over 273 years).

You can also use the chage command in interactive mode to modify multiple password aging and account details. Use the following command to enter interactive mode:

```
chage <username>
```

다음은 이러한 명령을 사용한 상호 대화식 세션의 예입니다:

```
~]# chage davido
Changing the aging information for davido
Enter the new value, or press ENTER for the default

        Minimum Password Age [0]: 10
        Maximum Password Age [99999]: 90
        Last Password Change (YYYY-MM-DD) [2006-08-18]:
        Password Expiration Warning [7]:
        Password Inactive [-1]:
        Account Expiration Date (YYYY-MM-DD) [1969-12-31]:
~]#
```

사용 가능한 옵션에 대해 변경된 자세한 정보를 원하시면 메뉴얼 페이지를 참조하시기 바랍니다.

You can also use the graphical User Manager application to create password aging policies, as follows. Note: you need Administrator privileges to perform this procedure.

1. Click the System menu on the Panel, point to Administration and then click Users and Groups to display the User Manager. Alternatively, type the command system-config-users at a shell prompt.

2. Click the Users tab, and select the required user in the list of users.

3. Click Properties on the toolbar to display the User Properties dialog box (or choose Properties on the File menu).

4. Click the Password Info tab, and select the check box for Enable password expiration.

5. Enter the required value in the Days before change required field, and click OK.

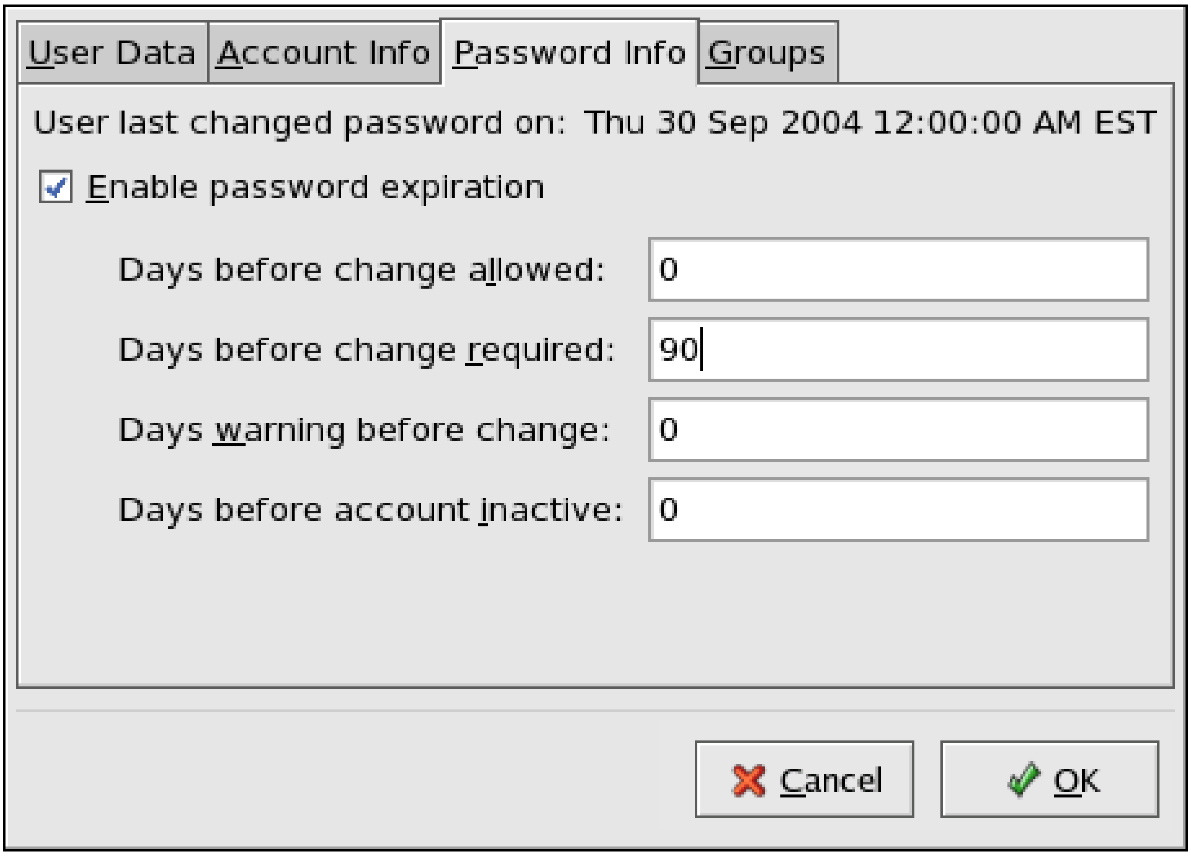

User last changed password on: Thu 30 Sep 2004 12:00:00 AM EST

☑ Enable password expiration

Days before change allowed: 0

Days before change required: 90

Days warning before change: 0

Days before account inactive: 0

그림 46.1. Specifying password aging options

For more information about user and group configuration (including instructions on forcing first time passwords), refer to 35장. 사용자 및 그룹.

46.1.4. 관리 제어

When administering a home machine, the user must perform some tasks as the root user or by acquiring effective root privileges via a setuid program, such as sudo or su. A setuid program is one that operates with the user ID (UID) of the program's owner rather than the user operating the program. Such programs are denoted by an s in the owner section of a long format listing, as in the following example:

```
-rwsr-xr-x    1 root     root        47324 May   1 08:09 /bin/su
```

알림

The s may be upper case or lower case. If it appears as upper case, it means that the underlying permission bit has not been set.

For the system administrators of an organization, however, choices must be made as to how much administrative access users within the organization should have to their machine. Through a PAM module called pam_console.so, some activities normally reserved only for the root user, such as rebooting and mounting removable media are allowed for the first user that logs in at the physical console (refer to 46.4절. "PAM (Pluggable Authentication Modules)" for more information about the pam_console.so module.) However, other important system administration tasks, such as altering network settings, configuring a new mouse, or mounting network devices, are not possible without administrative privileges. As a result, system administrators must decide how much access the users on their network should receive.

46.1.4.1. 루트 액세스 허용하기

만일 기업체 내의 사용자들이 신뢰할 수 있고 컴퓨터에 대한 많은 지식을 갖추고 있다면 루트 액세스를 허용해도 괜찮습니다. 사용자에게 루트 액세스를 허용하시면 개인 사용자가 장치를 추가하거나 네트워크 인터페이스를 설정하는 것과 같은 사소한 작업을 스스로 처리할 수 있게되며 시스템 관리자는 네트워크 보안과 다른 중요한 문제에 보다 많은 시간과 노력을 할애할 수 있게 됩니다.

다른 한편으로 개인 사용자에게 루트 액세스를 할당함으로서 다음과 같은 문제가 발생할 수 있습니다:

- 잘못된 시스템 설정 — 루트 액세스를 가진 사용자가 시스템을 잘못 설정하여 도움을 요청하거나, 심각한 경우 보안 약점이 생성될 수도 있습니다.

- 비보안 서비스 실행 — 루트 액세스를 가진 사용자가 FTP와 Telnet과 같이 비보안 서버를 시스템 상에서 실행하여 네트워크 상에서 사용자명과 암호를 암호화되지 않은 상태로 전달하여 누출될 위험에 처할 가능성이 있습니다.

- 루트 사용자로 이메일 첨부 파일을 실행 — 드문 예이기는 하지만 리눅스에 영향을 미치는 이메일 바이러스가 존재하기도 합니다. 그러나 이러한 바이러스는 루트 사용자가 실행할 경우에만 위험합니다.

46.1.4.2. 루트 액세스를 허가하지 않기

If an administrator is uncomfortable allowing users to log in as root for these or other reasons, the root password should be kept secret, and access to runlevel one or single user mode should be disallowed through boot loader password protection (refer to 46.1.2.2절. "부트로더 암호" for more information on this topic.)

The following are four different ways that an administrator can further ensure that root logins are disallowed:

Changing the root shell

To prevent users from logging in directly as root, the system administrator can set the root account's shell to /sbin/nologin in the /etc/passwd file.

표 46.1. 루트 쉘 사용 금지하기

영향을 미치는 사항	영향을 미치지 않는 사항
Prevents access to the root shell and logs any such attempts. The following programs are prevented from accessing the root account: • login • gdm • kdm • xdm	Programs that do not require a shell, such as FTP clients, mail clients, and many setuid programs. The following programs are not prevented from accessing the root account: • sudo • FTP clients • Email clients

영향을 미치는 사항	영향을 미치지 않는 사항
• su • ssh • scp • sftp	

Disabling root access via any console device (tty)

To further limit access to the root account, administrators can disable root logins at the console by editing the /etc/securetty file. This file lists all devices the root user is allowed to log into. If the file does not exist at all, the root user can log in through any communication device on the system, whether via the console or a raw network interface. This is dangerous, because a user can log in to their machine as root via Telnet, which transmits the password in plain text over the network.

By default, Red Hat Enterprise Linux's /etc/securetty file only allows the root user to log in at the console physically attached to the machine. To prevent the root user from logging in, remove the contents of this file by typing the following command at a shell prompt as root:

```
echo > /etc/securetty
```

To enable securetty support in the KDM, GDM, and XDM login managers, add the following line:

```
auth [user_unknown=ignore success=ok ignore=ignore default=bad] pam_securetty.so
```

to the files listed below:

- /etc/pam.d/gdm
- /etc/pam.d/gdm-autologin
- /etc/pam.d/gdm-fingerprint
- /etc/pam.d/gdm-password
- /etc/pam.d/gdm-smartcard
- /etc/pam.d/kdm
- /etc/pam.d/kdm-np
- /etc/pam.d/xdm

경고

A blank /etc/securetty file does not prevent the root user from logging in remotely using the OpenSSH suite of tools because the console is not opened until after authentication.

표 46.2. 루트 로그인 금지하기

영향을 미치는 사항	영향을 미치지 않는 사항
콘솔 또는 네트워크를 통하여 루트 계정에 액세스하지 못하게 합니다. 다음 프로그램을 사용하여 루트 계정에 액세스하지 못하게 합니다: • login • gdm • kdm • xdm • Other network services that open a tty	Programs that do not log in as root, but perform administrative tasks through setuid or other mechanisms. The following programs are not prevented from accessing the root account: • su • sudo • ssh • scp • sftp

Disabling root SSH logins

To prevent root logins via the SSH protocol, edit the SSH daemon's configuration file, /etc/ssh/sshd_config, and change the line that reads:

```
#PermitRootLogin yes
```

다음과 같이 변경하십시오:

```
PermitRootLogin no
```

표 46.3. 루트 SSH 로그인 금지하기

영향을 미치는 사항	영향을 미치지 않는 사항
OpenSSH suite 도구를 사용하여 루트 액세스를 하지 못하게 합니다. 다음 프로그램을 사용하여 루트 계정으로 액세스하지 못하게 합니다: • ssh • scp • sftp	Programs that are not part of the OpenSSH suite of tools.

Using PAM to limit root access to services

PAM, through the /lib/security/pam_listfile.so module, allows great flexibility in denying specific accounts. The administrator can use this module to reference a list of users who are not allowed to log in. To limit root access to a system service, edit the file for the target service in the /etc/pam.d/ directory and make sure the pam_listfile.so module is required for authentication.

The following is an example of how the module is used for the vsftpd FTP server in the /etc/pam.d/vsftpd PAM configuration file (the \ character at the end of the first line is not necessary if the directive is on a single line):

```
auth    required    /lib/security/pam_listfile.so    item=user \
 sense=deny file=/etc/vsftpd.ftpusers onerr=succeed
```

This instructs PAM to consult the /etc/vsftpd.ftpusers file and deny access to the service for any listed user. The administrator can change the name of this file, and can keep separate lists for each service or use one central list to deny access to multiple services.

If the administrator wants to deny access to multiple services, a similar line can be added to the PAM configuration files, such as /etc/pam.d/pop and /etc/pam.d/imap for mail clients, or /etc/pam.d/ssh for SSH clients.

For more information about PAM, refer to 46.4절. "PAM (Pluggable Authentication Modules)" .

표 46.4. PAM을 사용한 루트 계정 로그인 금지하기

영향을 미치는 사항	영향을 미치지 않는 사항
Prevents root access to network services that are PAM aware. The following services are prevented from accessing the root account: • login • gdm • kdm • xdm • ssh • scp • sftp • FTP clients • Email clients • Any PAM aware services	PAM을 인식하지 않는 프로그램과 서비스.

46.1.4.3. 루트 액세스 제한하기

Rather than completely denying access to the root user, the administrator may want to allow access only via setuid programs, such as su or sudo.

46.1.4.3.1. The su Command

When a user executes the su command, they are prompted for the root password and, after authentication, is given a root shell prompt.

Once logged in via the su command, the user is the root user and has absolute administrative access to the system[3]. In addition, once a user has become root, it is possible for them to use the su command to change to any other user on the system without being prompted for a password.

이 명령은 매우 강력하므로, 시스템 관리자는 기업 내에서 이 명령을 사용할 수 있는 권한을 제한하시는 것이 좋습니다.

[3] This access is still subject to the restrictions imposed by SELinux, if it is enabled.

가장 쉬운 방법은 사용자를 특별 관리 그룹인 wheel에 추가하시는 것입니다. 루트로 로그인하신 후 다음 명령을 입력하십시오:

```
usermod -G wheel <username>
```

In the previous command, replace <username> with the username you want to add to the wheel group.

You can also use the User Manager to modify group memberships, as follows. Note: you need Administrator privileges to perform this procedure.

1. Click the System menu on the Panel, point to Administration and then click Users and Groups to display the User Manager. Alternatively, type the command system-config-users at a shell prompt.

2. Click the Users tab, and select the required user in the list of users.

3. Click Properties on the toolbar to display the User Properties dialog box (or choose Properties on the File menu).

4. Click the Groups tab, select the check box for the wheel group, and then click OK. Refer to 그림 46.2. "Adding users to the "wheel" group." .

5. Open the PAM configuration file for su (/etc/pam.d/su) in a text editor and remove the comment # from the following line:

   ```
   auth   required  /lib/security/$ISA/pam_wheel.so  use_uid
   ```

   This change means that only members of the administrative group wheel can use this program.

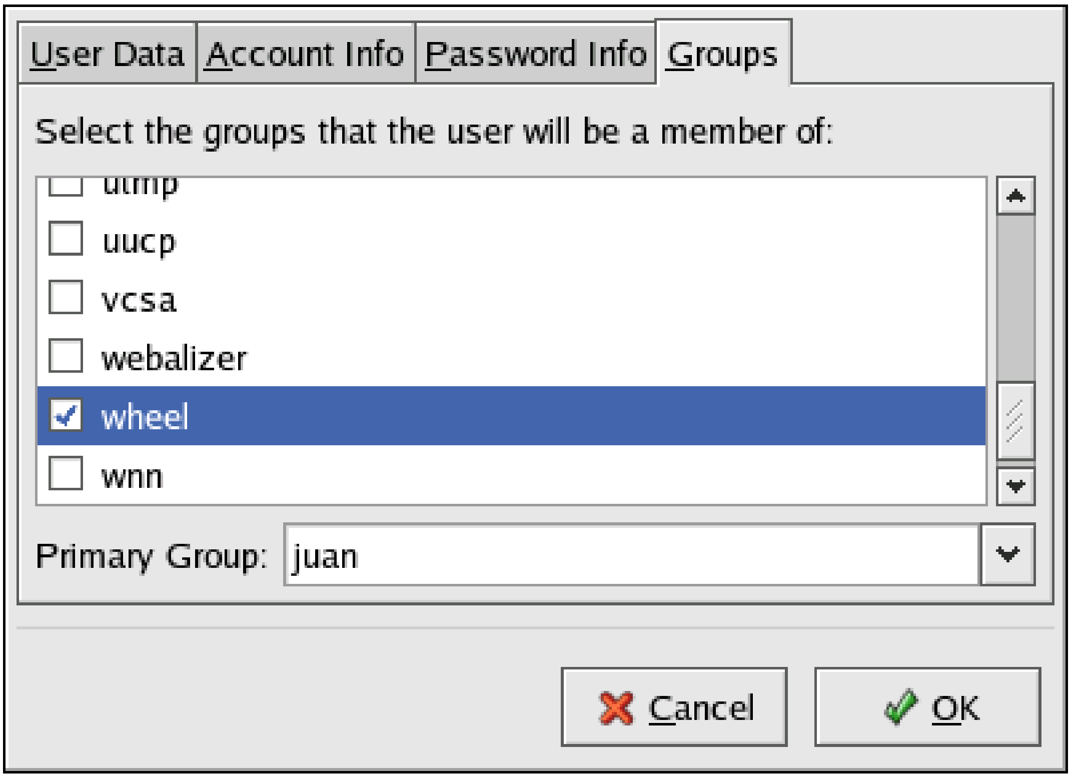

Select the groups that the user will be a member of:

□ utmp
□ uucp
□ vcsa
□ webalizer
☑ wheel
□ wnn

Primary Group: juan

그림 46.2. Adding users to the "wheel" group.

알림

The root user is part of the wheel group by default.

46.1.4.3.2. The sudo Command

The sudo command offers another approach to giving users administrative access. When trusted users precede an administrative command with sudo, they are prompted for their own password. Then, when they have been authenticated and assuming that the command is permitted, the administrative command is executed as if they were the root user.

The basic format of the sudo command is as follows:

```
sudo <command>
```

In the above example, <command> would be replaced by a command normally reserved for the root user, such as mount.

중요

Users of the sudo command should take extra care to log out before walking away from their machines since sudoers can use the command again without being asked for a password within a five minute period. This setting can be altered via the configuration file, /etc/sudoers.

The sudo command allows for a high degree of flexibility. For instance, only users listed in the /etc/sudoers configuration file are allowed to use the sudo command and the command is executed in the user's shell, not a root shell. This means the root shell can be completely disabled, as shown in 46.1.4.2절. "루트 액세스를 허가하지 않기".

The sudo command also provides a comprehensive audit trail. Each successful authentication is logged to the file /var/log/messages and the command issued along with the issuer's user name is logged to the file /var/log/secure.

Another advantage of the sudo command is that an administrator can allow different users access to specific commands based on their needs.

Administrators wanting to edit the sudo configuration file, /etc/sudoers, should use the visudo command.

To give someone full administrative privileges, type visudo and add a line similar to the following in the user privilege specification section:

```
juan ALL=(ALL) ALL
```

This example states that the user, juan, can use sudo from any host and execute any command.

The example below illustrates the granularity possible when configuring sudo:

```
%users  localhost=/sbin/shutdown -h now
```

This example states that any user can issue the command /sbin/shutdown -h now as long as it is issued from the console.

The man page for sudoers has a detailed listing of options for this file.

46.1.5. 사용 가능한 네트워크 서비스

기업체 시스템을 담당하는 관리자에게는 사용자의 관리 제어 액세스를 조정하는 것이 중요하지만 어느 네트워크 서비스가 활성화되었는지 확인하는 것도 리눅스 시스템을 관리하고 운영하는데 매우 중요합니다.

Red Hat Enterprise Linux 에서 많은 서비스는 네트워크 서버로 작동합니다. 만일 네트워크 서비스가 시스템에서 실행 중이라면, 데몬이라고 부르는 서버 응용 프로그램은 한 개 이상의 네트워크 포트를 청취합니다. 이러한 서버는 잠재적으로 침입을 당할 수 있는 약점이 됩니다.

46.1.5.1. 서비스로 인한 보안 위험

네트워크 서비스는 리눅스 시스템에 많은 위험을 주며, 주요 문제점은 다음과 같습니다:

- 서비스 거부 공격 (DoS) — 수많은 요청을 한 서비스에 집중하여 보내는 서비스 거부 공격을 사용하여 시스템이 각 요청을 기록하고 응답하기 위한 작업 과부하로 멈추는 상태가 발생합니다.

- 스크립트 취약성 공격 — 만일 웹 서버와 같이 서버가 스크립트를 사용하여 서버쪽 작업을 실행한다면, 크래커는 잘못 작성된 스크립트에 공격을 가할 수 있습니다. 이러한 스크립트 취약성 공격으로 인해 버퍼 초과 현상이 발생하거나 침입자는 시스템 상의 파일을 변경할 수 있게 됩니다.

- 버퍼 초과 공격 — 번호가 0에서 1023 사이인 포트에 연결하는 서비스는 관리자로 실행되어야 합니다. 만일 이 프로그램에 버퍼 초과 공격이 가해진 경우 침입자는 데몬을 실행 중인 사용자인 것처럼 시스템에 침입할 수 있습니다. 크래커는 이용 가능한 초과 버퍼가 존재하는 것을 알고 있기 때문에, 취약한 시스템을 찾아내기 위해 자동화된 도구를 사용하여 일단 침입에 성공하면 자동 루트킷(rootkit)을 사용하여 시스템 관리자용 권한을 유지합니다.

알림

Red Hat Enterprise Linux에서는 x86-호환 가능한 유니프로세서 및 멀티프로세서 커널에 의해 지원되는 실행 가능한 메모리를 분리하여 보호해주는 기술인 ExecShield가 사용되어 버퍼 초과 취약성 공격을 방지해줍니다. ExecShield는 가상 메모리를 실행 가능한 부분과 실행 불가능한 부분으로 구분하여 버퍼 초과 공격으로 인한 위험을 줄여줍니다. 실행 가능한 부분 외부에서 프로그램 코드 (예, 버퍼 초과 공격을 이용하여 들어온 악의성 코드)를 실행하려고 한다면, 세그멘테이션 오류가 발생하여 바로 정지됩니다.

Execshield also includes support for No eXecute (NX) technology on AMD64 platforms and eXecute Disable (XD) technology on Itanium and Intel® 64 systems. These technologies work in conjunction with ExecShield to prevent malicious code from running in the executable portion of virtual memory with a granularity of 4KB of executable code, lowering the risk of attack from stealthy buffer overflow exploits.

Tip

네트워크 상 공격을 방지하기 위해서는 사용되지 않은 서비스를 모두 정지시켜야 합니다.

46.1.5.2. 서비스 식별과 설정

보안을 강화하기 위해서 Red Hat Enterprise Linux에 설치된 대부분의 네트워크 서비스는 꺼져있도록 기본 설정되었습니다. 그러나 여기에는 중요한 예외가 있습니다:

- cupsd — The default print server for Red Hat Enterprise Linux.
- lpd — An alternative print server.
- xinetd — A super server that controls connections to a range of subordinate servers, such as gssftp and telnet.
- sendmail — The Sendmail Mail Transport Agent (MTA) is enabled by default, but only listens for connections from the localhost.
- sshd — The OpenSSH server, which is a secure replacement for Telnet.

When determining whether to leave these services running, it is best to use common sense and err on the side of caution. For example, if a printer is not available, do not leave cupsd running. The same is true for portmap. If you do not mount NFSv3 volumes or use NIS (the ypbind service), then portmap should be disabled.

Red Hat Enterprise Linux ships with three programs designed to switch services on or off. They are the Services Configuration Tool (system-config-services), ntsysv, and chkconfig. For information on using these tools, refer to 17장. 서비스로의 접근 통제.

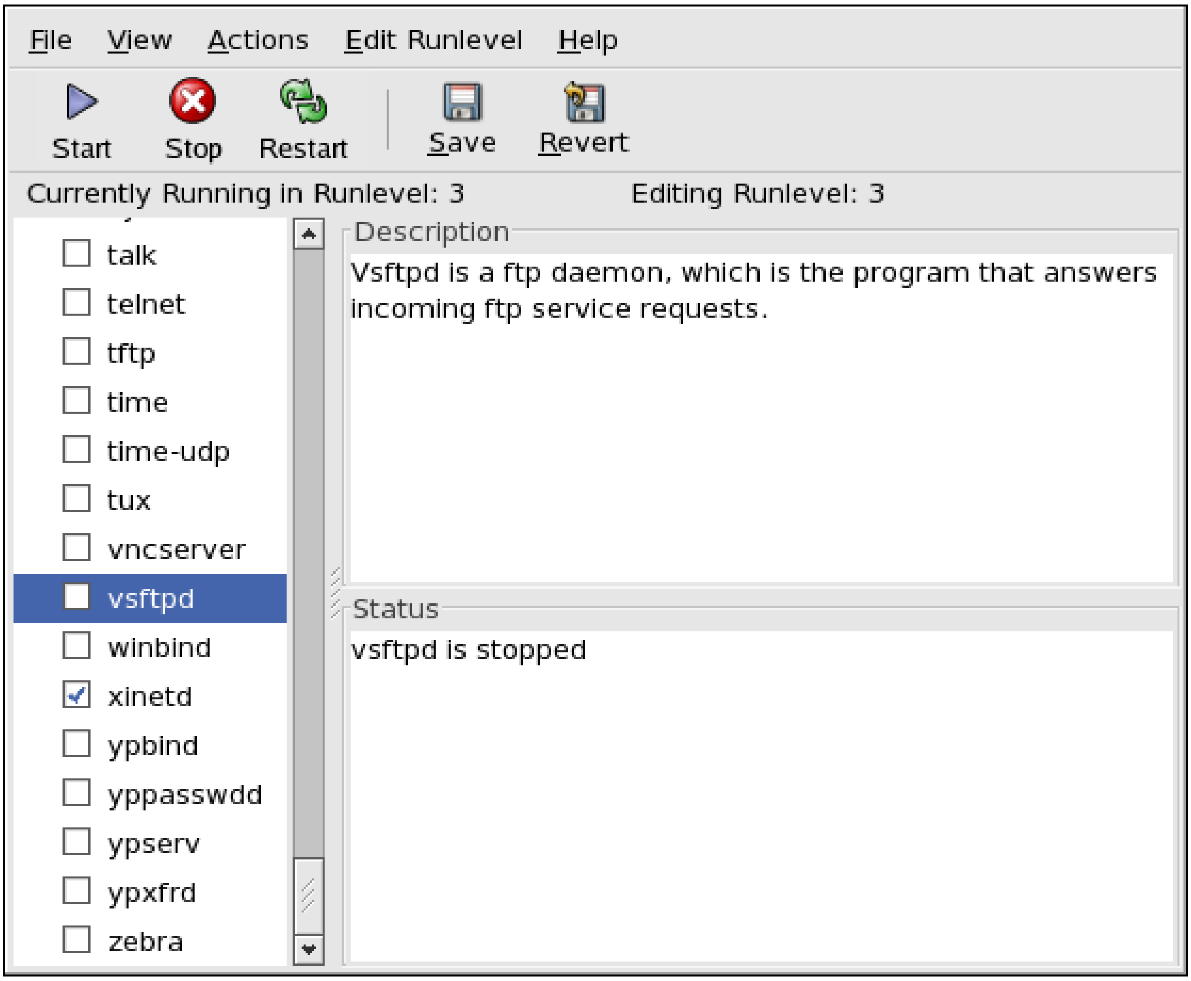

Vsftpd is a ftp daemon, which is the program that answers incoming ftp service requests.

vsftpd is stopped

그림 46.3. Services Configuration Tool

If unsure of the purpose for a particular service, the Services Configuration Tool has a description field, illustrated in 그림 46.3. "Services Configuration Tool" , that provides additional information.

Checking which network services are available to start at boot time is only part of the story. You should also check which ports are open and listening. Refer to 46.2.8절. "청취 중인 포트 확인하기" for more information.

46.1.5.3. 비보안 서비스

Potentially, any network service is insecure. This is why turning off unused services is so important. Exploits for services are routinely revealed and patched, making it very important to regularly update packages associated with any network service. Refer to 45.5절. "보안 업데이트" for more information.

일부 네트워크 프로토콜은 다른 프로토콜에 비해 더욱 비안전적입니다. 다음과 같은 특성을 지닌 서비스가 그러합니다:

- 암호화되지 않은 상태에서 네트워크를 통해 사용자명과 암호를 전달 — Telnet과 FTP와 같이 이전 프로토콜은 인증 세션을 암호화하지 않으므로 가능한 사용하지 않으시는 것이 좋습니다.

- Transmit Sensitive Data Over a Network Unencrypted — Many protocols transmit data over the network unencrypted. These protocols include Telnet, FTP, HTTP, and SMTP. Many network file systems, such as NFS and SMB, also transmit information over the network unencrypted. It is the user's responsibility when using these protocols to limit what type of data is transmitted.

Remote memory dump services, like netdump, transmit the contents of memory over the network unencrypted. Memory dumps can contain passwords or, even worse, database entries and other sensitive information.

Other services like finger and rwhod reveal information about users of the system.

Examples of inherently insecure services include rlogin, rsh, telnet, and vsftpd.

All remote login and shell programs (rlogin, rsh, and telnet) should be avoided in favor of SSH. Refer to 46.1.7절. "보안 강화된 통신 도구" for more information about sshd.

FTP is not as inherently dangerous to the security of the system as remote shells, but FTP servers must be carefully configured and monitored to avoid problems. Refer to 46.2.6절. "FTP 보안 강화" for more information about securing FTP servers.

다음과 같은 서비스는 신중하게 방화벽과 함께 구현되어야 합니다:

- finger
- authd (this was called identd in previous Red Hat Enterprise Linux releases.)
- netdump
- netdump-server
- nfs
- rwhod
- sendmail
- smb (Samba)
- yppasswdd
- ypserv
- ypxfrd

More information on securing network services is available in 46.2절. "서버 보안".

다음 부분에서는 간단한 방화벽을 설정하는데 사용되는 도구에 대하여 설명해 보겠습니다.

46.1.6. 개인 방화벽

필수 네트워크 서비스를 설정하신 후에는 방화벽을 실행하셔야 합니다.

중요

인터넷 또는 신뢰하지 않는 네트워크에 연결하기 전에 필수 서비스를 설정하고 방화벽을 실행하셔야 합니다.

Firewalls prevent network packets from accessing the system's network interface. If a request is made to a port that is blocked by a firewall, the request is ignored. If a service is listening on one of these blocked ports, it does not receive the packets and is effectively disabled. For this reason, care should be taken when configuring a firewall to block access to ports not in use, while not blocking access to ports used by configured services.

For most users, the best tool for configuring a simple firewall is the graphical firewall configuration tool which ships with Red Hat Enterprise Linux: the Security Level Configuration Tool (system-config-securitylevel). This tool creates broad iptables rules for a general-purpose firewall using a control panel interface.

Refer to 46.8.2절. "Basic Firewall Configuration" for more information about using this application and its available options.

For advanced users and server administrators, manually configuring a firewall with iptables is probably a better option. Refer to 46.8절. "Firewalls" for more information. Refer to 46.9절. "IPTables" for a comprehensive guide to the iptables command.

46.1.7. 보안 강화된 통신 도구

인터넷의 인기와 크기가 날로 증가됨에 따라서 통신을 중간에서 가로채기할 위협도 증가되었습니다. 시간이 지나면서 네트워크 상에서 전송되는 통신을 암호화할 수 있는 도구가 개발되었습니다.

Red Hat Enterprise Linux에는 네트워크 상에서 전송되는 정보를 보호하기 위해 고수준, 공개키 암호화에 기반한 암호화 알고리즘 기법을 사용하는 두가지 기본 도구가 포함되어 있습니다.

- OpenSSH — 네트워크 통신을 암호화하는데 사용되는 공개 소스 SSH 프로토콜입니다.
- GPG (Gnu Privacy Guard) — 데이터를 암호화하는데 사용되는 공개 소스 PGP (Pretty Good Privacy) 암호화 프로그램입니다.

OpenSSH is a safer way to access a remote machine and replaces older, unencrypted services like telnet and rsh. OpenSSH includes a network service called sshd and three command line client applications:

- ssh — A secure remote console access client.
- scp — A secure remote copy command.
- sftp — A secure pseudo-ftp client that allows interactive file transfer sessions.

> **중요**
>
> Although the sshd service is inherently secure, the service must be kept up-to-date to prevent security threats. Refer to 45.5절. "보안 업데이트" for more information.

GPG는 이메일을 주고 받을때 데이터를 보호하기 위한 좋은 방법입니다. 공중 네트워크 상에서 기밀 데이터를 이메일로 보내거나 하드 드라이브 상에서 기밀 데이터를 보존할 경우에 모두 사용 가능합니다.

46.2. 서버 보안

시스템이 공중 네트워크에서 서버로 사용될 경우 공격의 대상이 되기 쉽습니다. 이러한 이유로 시스템 보안을 보강하고 서비스를 잠그는 것은 시스템 관리자에게 무엇보다 중요합니다.

특정 사항에 대하여 깊이 파고들기 이전에 서버 보안을 강화시킬 수 있는 일반적인 힌트를 다음에서 간략히 살펴보도록 하겠습니다:

- 최신 침입 유형에 대비하여 모든 서비스를 항상 업데이트 시키십시오.
- 가능한 보안 프로토콜을 사용하십시오.
- 가능한 한 기계당 한가지 유형의 네트워크 서비스를 사용하십시오.
- 모든 서버에서 수상한 행동이 발견되는지 주의깊게 감시하십시오.

46.2.1. TCP 래퍼와 xinetd를 사용하여 서비스 보안 강화하기

TCP 래퍼(Wrappers)는 다양한 서비스에 접근 제어를 제공합니다. SSH, Telnet, FTP와 같은 대부분의 최신 네트워크 서비스는 들어오는 요청과 요청된 서비스 사이에서 감시 역할을 하는 TCP 래퍼를 사용합니다.

추가 액세스, 기록, 바인딩, 방향 전환 및 자원 활용 제어와 같은 기능을 제공하는 수퍼 서버인 xinetd를 함께 사용하면 TCP 래퍼가 제공하는 보안 기능이 보다 강화됩니다.

Tip

It is a good idea to use iptables firewall rules in conjunction with TCP Wrappers and xinetd to create redundancy within service access controls. Refer to 46.8절. "Firewalls" for more information about implementing firewalls with iptables commands.

Refer to 17.2절. "TCP 래퍼 (Wrappers)" for more information on configuring TCP Wrappers and xinetd.

다음 부분에서는 여러분이 각 주제에 대한 기본적인 지식을 갖추고 계신다고 간주하고 특정 보안 옵션에 중점을 두고 설명해 보겠습니다.

46.2.1.1. TCP 래퍼를 사용하여 보안 강화하기

TCP 래퍼는 서비스로의 액세스를 거부하는 것 이외에도 다른 많은 기능을 제공합니다. 이 부분에서는 TCP 래퍼를 사용하여 연결 배너를 보내고, 특정 호스트에서 침입자에게 경고 메시지를 보내며, 기록 기능을 강화하는 방법에 대하여 설명하고 있습니다. TCP 래퍼의 기능과 제어 언어에 대한 전체적인 목록을 보시려면 hosts_options 메뉴얼 페이지를 참조하시기 바랍니다.

46.2.1.1.1. TCP 래퍼와 연결 배너

서비스에 접속하는 클라이언트에 경고성 배너를 보내는 것이 서버가 어떠한 시스템을 운영 중인지 보여주지 않으면서 동시에 침입자에게 시스템 관리자가 감시 중이라고 알려줄 수 있는 좋은 방법입니다. 서비스에 TCP 래퍼 배너를 구현하시려면 banner 옵션을 사용하십시오.

이 예시는 vsftpd에 배너를 사용합니다. 먼저 배너 파일을 생성하셔야 합니다. 시스템 상 어디에서든 생성하실 수 있지만 이 파일은 사용될 데몬과 동일한 이름을 가져야 합니다. 이 예시에서 파일 이름은 /etc/banners/vsftpd 입니다:

```
220-Hello, %c
220-All activity on ftp.example.com is logged.
220-Inappropriate use will result in your access privileges being removed.
```

%c 토큰은 사용자명, 호스트명 또는 연결 메시지에 보다 효과가 있도록 사용자명과 IP 주소와 같은 다양한 클라이언트 정보를 제공합니다.

이 배너가 들어오는 접속에 보여지도록 하시려면 /etc/hosts.allow 파일에서 다음 줄을 추가하시면 됩니다:

```
vsftpd : ALL : banners /etc/banners/
```

46.2.1.1.2. TCP 래퍼와 침입 경고

만일 특정 호스트나 네트워크가 서버를 침입하는 것이 발견되었다면 TCP 래퍼에 spawn 지시자를 사용하여 침입이 시도된 호스트나 네트워크의 관리자에게 경고 메시지를 보낼 수 있습니다.

예를 들어 206.182.68.0/24 네트워크에서 크래커가 서버에 침입 시도하려는 것이 발견되었다고 가정해봅니다. /etc/hosts.deny 파일에 다음과 같은 줄을 추가하시면, 연결 시도가 거부되며 특별 파일에 기록될 것입니다:

```
ALL : 206.182.68.0 : spawn /bin/ 'date' %c %d >> /var/log/intruder_alert
```

%d 토큰은 침입자가 접근하려고 시도한 서비스의 이름을 제공합니다.

연결을 허용 후 기록하기 위해서는 /etc/hosts.allow 파일에 spawn 지시자를 추가하시기 바랍니다.

알림

spawn 지시자는 모든 쉘 명령을 실행하므로, 특정 클라이언트가 서버에 접속을 시도할 경우 관리자에게 알리거나 여러 명령을 수행할 특수 스크립트를 작성하십시오.

46.2.1.1.3. TCP 래퍼와 향상된 기록 기능

만일 특정 유형의 접속이 다른 유형 보다 중요하다면 severity 옵션을 사용하여 해당 서비스에 대한 여러 다른 기록 수준을 설정하실 수 있습니다.

이 예시에서는 FTP 서버 포트 23 (Telnet 포트)로 접속을 시도하는 사용자를 크래커라고 가정합니다. 크래커가 침입하는 것을 방지하기 위하여 로그 파일에서 기본 플래그(flag)인 info 대신 emerg 플래그를 지정하시고 이 포트로 들어오는 연결을 거부합니다.

연결을 거부하기 위해서는 /etc/hosts.deny 파일에 다음 줄을 추가하시면 됩니다:

```
in.telnetd : ALL : severity emerg
```

이러한 설정은 기본 authpriv 기록 기능을 사용하지만 기록 심각성 수준을 기본 값인 info에서 emerg 수준으로 높여서 로그 메시지를 콘솔에 바로 보여줍니다.

46.2.1.2. xinetd를 사용하여 보안 강화하기

이 부분에서는 xinetd를 사용하여 트랩(trap) 서비스를 설정하는 방법과 xinetd 서비스가 사용할 수 있는 자원의 양을 제어하는 방법에 대하여 중점적으로 설명해 보겠습니다. 서비스가 사용할 수 있는 자원의 한계를 설정하면 서비스 거부 (DoS) 공격을 좌절시킬 수 있습니다.모든 사용 가능한 옵션의 목록을 보시려면 xinetd 및 xinetd.conf의 메뉴얼 페이지를 참조하시기 바랍니다.

46.2.1.2.1. 트랩(Trap) 설정하기

xinetd의 중요한 기능 중 하나는 전역 no_access 목록에 호스트를 추가할 수 있는 기능입니다. 이 목록에 포함된 호스트는 xinetd가 관리하는 서비스에 정해진 기간 동안 또는 xinetd가 재시작될 때까지 연결을 거부당합니다. 이 기능은 SENSOR 속성을 통해 실행 가능하며, 서버에서 포트를 스캔하려고 시도하는 호스트를 손쉽게 막을 수 있는 방법입니다.

SENSOR를 설정하기 위한 첫번째 단계는 사용할 계획이 없는 서비스를 선택하는 것입니다. 이 예에서는 Telnet이 사용됩니다.

/etc/xinetd.d/telnet 파일에서 flags 줄을 다음과 같이 수정하시기 바랍니다:

```
flags           = SENSOR
```

다음 줄을 추가하십시오:

```
deny_time       = 30
```

이 설정은 포트로 연결을 시도하는 호스트를 30 분 동안 거부할 것입니다. deny_time 속성에 사용 가능한 다른 값에는 FOREVER와 NEVER가 있습니다. FORVER는 xinetd가 재시작될 때까지 연결을 거부하며, NEVER는 연결을 허용한 후 기록합니다.

마지막 줄을 다음과 같이 수정하십시오:

```
disable         = no
```

이는 트랩 자체를 활성화시킵니다.

SENSOR를 사용하여 보안을 위협하는 호스트로부터 연결을 검색하여 정지시키는 것이 좋은 방법이기는 하지만, 다음과 같은 두가지 결점이 있습니다:

- 스텔스 스캔 (쉽게 발견되지 않도록 한 스캔)을 찾아내지 못합니다.
- 만일 침입자가 SENSOR가 실행 중인 사실을 이미 알고 있다면 자신의 IP 주소를 위장하여 특정 호스트에 서비스 거부 공격을 마운트한 후 금지된 포트에 연결할 수 있습니다.

46.2.1.2.2. 서버 자원을 제어하기

xinetd의 또 다른 중요한 기능은 서비스가 활용 가능한 자원의 양을 제어할 수 있는 기능입니다.

다음 지시자를 통하여 이 기능을 사용 가능합니다:

- cps = <number_of_connections> <wait_period> — Limits the rate of incoming connections. This directive takes two arguments:
  - <number_of_connections> — The number of connections per second to handle. If the rate of incoming connections is higher than this, the service is temporarily disabled. The default value is fifty (50).

  - <wait_period> — The number of seconds to wait before re-enabling the service after it has been disabled. The default interval is ten (10) seconds.

- instances = <number_of_connections> — Specifies the total number of connections allowed to a service. This directive accepts either an integer value or UNLIMITED.
- per_source = <number_of_connections> — Specifies the number of connections allowed to a service by each host. This directive accepts either an integer value or UNLIMITED.
- rlimit_as = <number[K|M]> — Specifies the amount of memory address space the service can occupy in kilobytes or megabytes. This directive accepts either an integer value or UNLIMITED.
- rlimit_cpu = <number_of_seconds> — Specifies the amount of time in seconds that a service may occupy the CPU. This directive accepts either an integer value or UNLIMITED.

이러한 지시자를 사용하시면, 서비스 거부 공격을 통해 xinetd 서비스가 시스템을 마비시키는 상황을 방지하는데 도움이 됩니다.

46.2.2. Portmap 보안 강화

portmap 서비스는 NIS와 NFS와 같은 RPC 서비스에 사용되는 동적 포트 할당 데몬입니다. 이 데몬은 허술한 인증 메커니즘을 갖추고 있으며 데몬이 제어하는 서비스에 광범위한 포트를 할당 가능합니다. 따라서 보안 관리가 쉽지 않습니다.

> **알림**
>
> portmap을 보안 강화하게 되면 NFSv2와 NFSv3만 영향을 받습니다. NFSv4는 더 이상 portmap을 사용하지 않으므로 영향을 받지 않습니다. NFSv2 이나 NFSv3 서버를 구현할 계획이라면, portmap이 사용되므로 다음 부분에서 설명된 내용을 따르십시오.

RPC 서비스를 실행하신다면 다음과 같은 기본 규칙을 따르십시오.

46.2.2.1. TCP 래퍼를 사용하여 portmap 보호

portmap 서비스에는 내장된 인증 방식이 없으므로 TCP 래퍼를 사용하여 이 서비스를 사용할 수 있는 네트워크나 호스트를 제한하는 것이 중요합니다.

또한 서비스로 접근을 제한하실 때는 IP 주소만 사용하셔야 합니다. 호스트명은 DNS poisoning이나 다른 방법으로 위조가 가능하므로 사용하지 마십시오.

46.2.2.2. IPTables를 사용하여 portmap 보호

portmap 서비스로 접근을 더 제한하시려면 서버에 iptables 규칙을 추가하여 특정 네트워크로 접근하는 것을 제한하시는 것이 좋습니다.

다음은 (포트 111을 청취하는) portmap 서비스로 192.168.0/24 네트워크와 로컬호스트에서 TCP 연결을 허용하는 두가지 IPTables 명령 예시입니다. Nautilus가 sgi_fam 서비스를 사용하는데 필요한 설정입니다. 모든 다른 패킷은 버립니다(drop).

```
iptables -A INPUT -p tcp -s! 192.168.0.0/24 --dport 111 -j DROP
iptables -A INPUT -p tcp -s 127.0.0.1 --dport 111 -j ACCEPT
```

이와 유사한 방식으로 UDP 트래픽을 제한하기 위해서는 다음 명령을 사용하십시오.

```
iptables -A INPUT -p udp -s! 192.168.0.0/24 --dport 111 -j DROP
```

Tip

Refer to 46.8절. "Firewalls" for more information about implementing firewalls with iptables commands.

46.2.3. NIS 보안 강화

The Network Information Service (NIS) is an RPC service, called ypserv,--> which is used in conjunction with portmap and other related services to distribute maps of usernames, passwords, and other sensitive information to any computer claiming to be within its domain.

NIS 서버는 다음과 같은 여러가지 응용 프로그램으로 구성되어 있습니다:

- /usr/sbin/rpc.yppasswdd — yppasswdd 서비스로도 불리우는 이 데몬은 사용자가 NIS 암호를 변경할 수 있게 해줍니다.
- /usr/sbin/rpc.ypxfrd — ypxfrd 서비스라고도 불리는 이 데몬은 네트워크 상에서 NIS 맵(map)을 전송합니다.
- /usr/sbin/yppush — 이 프로그램은 수정된 NIS 데이터베이스를 다수의 NIS 서버에 전달하는 역할을 합니다.
- /usr/sbin/ypserv — NIS 서버 데몬입니다.

NIS is somewhat insecure by today's standards. It has no host authentication mechanisms and transmits all of its information over the network unencrypted, including password hashes. As a result, extreme care must be taken when setting up a network that uses NIS. This is further complicated by the fact that the default configuration of NIS is inherently insecure.

It is recommended that anyone planning to implement an NIS server first secure the portmap service as outlined in 46.2.2절. "Portmap 보안 강화", then address the following issues, such as network planning.

46.2.3.1. 네트워크를 신중하게 설정하기

NIS는 네트워크 상에서 기밀 정보를 암호화되지 않은 상태에서 전달하기 때문에 서비스를 분할되고 안전한 네트워크 상에서 방화벽을 사용한 상태에서 서비스를 실행해야 합니다. 비보안 네트워크 상에서 NIS 정보가 전달될 때마다 누군가 정보를 가로챌 위협이 있습니다. 이러한 의미에서 네트워크를 신중히 설정함으로서 심각한 보안 침입 위협을 방지할 수 있습니다.

46.2.3.2. 암호와 같이 추측하기 힘든 NIS 도메인 이름과 호스트명 사용하기

Any machine within an NIS domain can use commands to extract information from the server without authentication, as long as the user knows the NIS server's DNS hostname and NIS domain name.

예를 들어 만일 누군가 네트워크에 랩탑 컴퓨터를 연결하거나 외부에서 네트워크에 침입하여 내부 IP 주소를 위장할 수 있다면 다음 명령을 사용하여/etc/passwd 파일의 내용을 보는 것이 가능합니다:

```
ypcat -d <NIS_domain> -h <DNS_hostname> passwd
```

만일 이 침입자가 루트 사용자 권한을 가지고 있다면 다음과 같은 명령을 입력하여 /etc/shadow 파일의 내용을 볼 수 있습니다:

```
ypcat -d <NIS_domain> -h <DNS_hostname> shadow
```

알림

커베로스가 사용된 경우 /etc/shadow 파일은 NIS map에 저장되지 않습니다.

침입자가 NIS map에 액세스하는 것을 보다 힘들게 하기 위하여 DNS 호스트명을 o7hfawtgmhwg.domain.com와 같은 임의 문자열로 생성하시는 것이 좋습니다. 유사한 방법으로 NIS 도메인 이름도 임의 문자열로 생성하시되 호스트명과 다른 이름을 생성하십시오. 이렇게 하시면 침입자가 NIS 서버에 액세스하는 것이 더욱 힘들어 집니다.

46.2.3.3. /var/yp/securenets 파일을 수정하기

/var/yp/securenets 파일이 공백으로 비어있거나 (기본 설치를 수행한 후) 파일이 존재하지 않는 경우 NIS는 모든 네트워크를 청취합니다. 이러한 경우 가장 먼저 하실 것은 ypserv가 적절한 네트워크에서 들어오는 요청만 응답하도록 이 파일에 넷마스크/네트워크 쌍을 입력하셔야 합니다.

다음은 /var/yp/securenets 파일 예제입니다:

```
255.255.255.0     192.168.0.0
```

경고

/var/yp/securenets 파일을 생성하지 않은 상태에서 NIS 서버를 처음으로 시작하시면 안됩니다.

이 기술은 IP 스푸핑 공격에 대한 보호를 제공하지는 못하지만 최소한 NIS 서비스가 청취할 네트워크를 제한해줍니다.

46.2.3.4. 정적 포트를 할당하고 iptables 규칙 사용하기

NIS와 관련된 모든 서버에 특정 포트를 할당하는 것이 가능하지만 사용자가 로그인 암호를 변경할 수 있게 해주는 데몬인 rpc.yppasswdd는 예외입니다. 다른 두 개의 NIS 서버 데몬인 rpc.ypxfrd와 ypserv에 포트를 할당함으로서 방화벽 규칙을 생성하여 침입자가 NIS 서버 데몬에 침입하지 못하도록 보안을 강화할 수 있습니다.

이러한 설정을 위해 /etc/sysconfig/network 파일에 다음과 같은 줄을 삽입하시기 바랍니다:

```
YPSERV_ARGS="-p 834" YPXFRD_ARGS="-p 835"
```

다음 iptables 규칙을 입력하여 이 포트에서 서버가 청취할 네트워크를 제한 설정하실 수 있습니다:

```
iptables -A INPUT -p tcp -s! 192.168.0.0/24 --dport 834 -j DROP
iptables -A INPUT -p tcp -s! 192.168.0.0/24 --dport 835 -j DROP
iptables -A INPUT -p udp -s! 192.168.0.0/24 --dport 834 -j DROP
iptables -A INPUT -p udp -s! 192.168.0.0/24 --dport 835 -j DROP
```

This means that the server only allows connections to ports 834 and 835 if the requests come from the 192.168.0.0/24 network.

Tip

Refer to 46.8절. "Firewalls" for more information about implementing firewalls with iptables commands.

46.2.3.5. 커베로스 인증 사용하기

NIS 인증이 사용될 경우 가장 심각한 결점은 사용자가 시스템에 로그인할 때마다 /etc/shadow 파일의 암호 해시가 네트워크 상에서 전달되는 것입니다. 만일 침입자가 NIS 도메인에 침입하여 네트워크 트래픽을 훔쳐보고 있다면 사용자명과 암호 해시를 모르게 수집할 수 있습니다. 충분한 시간이 주어진다면 암호 크래킹 프로그램을 사용하여 추측하기 쉬운 암호를 알아낸 후 침입자는 유효한 계정을 사용하여 네트워크에 액세스 가능합니다.

Since Kerberos uses secret-key cryptography, no password hashes are ever sent over the network, making the system far more secure. Refer to 46.6절. "Kerberos" for more information about Kerberos.

46.2.4. NFS 보안 강화

The Network File System (NFS) is a service that provides network accessible file systems for client machines. Refer to 20장. 네트워크 파일 시스템 (NFS) for more information about NFS. The following subsections assume a basic knowledge of NFS.

중요

The version of NFS included in Red Hat Enterprise Linux, NFSv4, no longer requires the portmap service as outlined in 46.2.2절. "Portmap 보안 강화" . NFS traffic now utilizes TCP in all versions, rather than UDP, and requires it when using NFSv4. NFSv4 now includes Kerberos user and group authentication, as part of the RPCSEC_GSS kernel module. Information on portmap is still included, since Red Hat Enterprise Linux supports NFSv2 and NFSv3, both of which utilize portmap.

46.2.4.1. 네트워크를 신중하게 설정하기

NFSv4는 네트워크 상에서 모든 정보를 커베로스를 사용하여 암호화하여 전달할 수 있으므로, 방화벽이나 분활된 네트워크 상에서 올바르게 서비스를 설정하셔야 합니다. NFSv2와 NFSv3는 여전히 정보를 암호화되지 않은 상태로 전달하기 때문에 이러한 점이 고려되어야 합니다. 따라서 네트워크를 신중히 설정함으로서 심각한 보안 침입 위협을 방지할 수 있습니다.

46.2.4.2. 구문 오류에 주의하십시오.

NFS 서버는 /etc/exports 파일을 통해 어느 파일 시스템은 익스포트할 것이고 이 디렉토리를 익스포트할 호스트는 무엇인지 결정합니다. 따라서 이 파일을 수정하실 때 불필요한 공간을 추가하지 않도록 주의하셔야 합니다.

예를 들어 /etc/exports 파일에서 다음과 같은 줄은 bob.example.com 호스트에서 /tmp/nfs/ 디렉토리를 읽고 쓸 수 있도록 공유합니다.

```
/tmp/nfs/     bob.example.com(rw)
```

반면 /etc/exports 파일에 이 줄을 삽입하시면 호스트명 다음에 삽입된 빈 공간 때문에 동일한 디렉토리를 bob.example.com에 읽기 전용 허가를 가지고 공유하고 그 외 다른 전체 호스트에 읽기 쓰기 허가를 가지고 이 디렉토리를 공유합니다.

```
/tmp/nfs/     bob.example.com (rw)
```

따라서 showmount 명령을 사용하여 어떠한 디렉토리가 어떻게 공유되고 있는지 NFS 공유 설정을 확인해보시기 바랍니다:

```
showmount -e <hostname>
```

46.2.4.3. no_root_squash 옵션을 사용하지 마십시오

NFS 공유는 루트 사용자를 특별한 권한이 없는 사용자 계정인 nfsnobody로 변경하도록 기본 설정되어 있습니다. 따라서 루트 사용자가 생성한 파일은 모두 nfsnobody가 소유하게 되어 사용자 아이디 비트를 재설정하여 프로그램을 업로드하지 못하게 됩니다.

If no_root_squash is used, remote root users are able to change any file on the shared file system and leave applications infected by Trojans for other users to inadvertently execute.

46.2.5. Apache HTTP 서버 보안 강화

Apache HTTP 서버는 Red Hat Enterprise Linux에 포함된 서비스 중 가장 안정적이고 안전한 서비스 중 하나입니다. Apache HTTP 서버 보안을 강화하기 위해 사용 가능한 매우 다양한 옵션과 기술이 존재합니다 — 이 메뉴얼에서 깊게 다루기에는 너무 광범위합니다.

When configuring the Apache HTTP Server, it is important to read the documentation available for the application. This includes 23장. Apache HTTP Server, and the Stronghold manuals, available at http://www.redhat.com/docs/manuals/stronghold/.

시스템 관리자는 다음의 설정 옵션을 사용할 때 유의하셔야 합니다:

46.2.5.1. FollowSymLinks

이 지시자는 기본값으로 활성화되어 있습니다. 따라서 웹 서버의 문서 루트로 심볼릭 링크를 생성하실 때는 주의하시기 바랍니다. 예를 들어 /로 심볼릭 링크를 제공하는 것은 좋은 생각이 아닙니다.

46.2.5.2. Indexes 지시자

이 지시자는 기본 값으로 활성화되어 있지만, 그리 바람직하지 않습니다. 침입자가 서버에서 파일을 검색하는 것을 방지하기 위해 이 지시자를 삭제하시기 바랍니다.

46.2.5.3. UserDir 지시자

UserDir 지시자는 침입자가 시스템 상에 사용자 계정이 존재하는지 확인할 수 있기 때문에 비활성화되도록 기본 설정되어 있습니다. 서버에서 사용자 디렉토리 검색 기능을 활성화하시려면 다음 지시자를 사용하시기 바랍니다:

```
UserDir enabled
UserDir disabled root
```

이 지시자는 사용자 디렉토리 검색 기능을 /root/를 제외한 모든 사용자 디렉토리에서 활성화할 것입니다. 비활성 계정 목록에 사용자를 추가하시려면 UserDir disabled 줄에 사용자 이름을 빈칸으로 구분하여 추가하십시오.

46.2.5.4. IncludesNoExec 지시자를 삭제하지 마십시오

Server-Side Includes (SSI) 모듈에서는 명령을 실행하지 못하도록 기본 설정되어 있습니다. 만일 절대적으로 필요한 경우가 아니라면 이 설정을 변경하지 마십시오. 이 설정을 변경하시면 침입자가 시스템에서 명령을 실행할 위험이 높아집니다.

46.2.5.5. 실행 가능 디렉토리의 허가 제한하기

스크립트나 CGI를 포함한 디렉토리에는 루트 사용자에게만 쓰기 허가를 부여하셔야 합니다. 다음 명령을 입력하시기 바랍니다:

```
chown root <directory_name>
chmod 755 <directory_name>
```

중요

또한 시스템 상에서 실행될 스크립트는 의도하는 대로 작동하는지 미리 확인하신 후 생산 환경에서 사용하셔야 합니다.

46.2.6. FTP 보안 강화

The File Transfer Protocol (FTP) is an older TCP protocol designed to transfer files over a network. Because all transactions with the server, including user authentication, are unencrypted, it is considered an insecure protocol and should be carefully configured.

Red Hat Enterprise Linux는 3가지 FTP 서버를 제공합니다.

- gssftpd — 커베로스를 사용하는 xinetd-기반 FTP 데몬으로서 인증 정보를 네트워크 상에서 전달하지 않습니다.
- Red Hat 콘텐츠 가속기(Content Accelerator) (tux) — FTP 기능을 갖춘 커널 영역 웹 서버.
- vsftpd — 독립형, 보안 FTP 서비스

다음은 vsftpd FTP 서비스를 설정하는데 사용되는 보안 정책입니다.

46.2.6.1. FTP 환영 배너

사용자명과 암호를 입력하기 전에 환경 배너가 나타납니다. 이 배너에는 버전 정보가 포함되어 있으며, 이 정보는 크래커가 시스템 약점을 찾아내는데 유용하게 사용됩니다.

따라서 vsftpd의 환영 배너를 변경하시려면 /etc/vsftpd/vsftpd.conf 파일에 다음 지시자를 추가하시기 바랍니다:

```
ftpd_banner=<insert_greeting_here>
```

Replace <insert_greeting_here> in the above directive with the text of the greeting message.

여러 개의 줄로 이루어진 배너 메시지를 입력하시려면 배너 파일을 사용하시는 것이 좋습니다 여러 배너를 손쉽게 관리하기 위하여 /etc/banners/라는 새 디렉토리를 만드신 후 모든 패너 파일을 이 디렉토리에 저장하십시오. 이 예시에서 FTP 접속에 사용되는 배너 파일은 /etc/banners/ftp.msg 입니다. 다음은 이 파일의 내용 예제입니다:

```
######### # 안녕하세요, ftp.example.com에 있는 모든 작업은 기록됩니다. #########
```

알림

It is not necessary to begin each line of the file with 220 as specified in 46.2.1.1.1절. "TCP 래퍼와 연결 배너" .

vsftpd에 이 환경 배너 파일을 사용하기 위해서는 /etc/vsftpd/vsftpd.conf 파일에 다음과 같은 지시자를 추가하십시오:

```
banner_file=/etc/banners/ftp.msg
```

중요

Make sure that you specify the path to the banner file correctly in /etc/vsftpd/vsftpd.conf, or else every attempt to connect to vsftpd will result in the connection being closed immediately and a 500 OOPS: cannot open banner <path_to_banner_file> error message.

Note that the banner_file directive in /etc/vsftpd/vfsftpd.conf takes precedence over any ftpd_banner directives in the configuration file: if banner_file is specified, then ftpd_banner is ignored.

It also is possible to send additional banners to incoming connections using TCP Wrappers as described in 46.2.1.1.1절. "TCP 래퍼와 연결 배너".

46.2.6.2. 익명 계정(anonymous) 접속

/var/ftp/ 디렉토리를 생성하시면 익명 계정이 활성화됩니다.

이 디렉토리를 생성할 수 있는 가장 쉬운 방법은 vsftpd 패키지를 설치하는 것입니다. 이 패키지는 익명 사용자를 위한 디렉토리 구조를 설정하고, 익명 사용자에게 이 디렉토리를 읽기만 할 수 있는 허가를 설정합니다.

익명 사용자는 어느 디렉토리에도 쓰기 작업을 할 수 없도록 기본 설정되어 있습니다.

주의

FTP 서버에 익명 계정으로 접속을 활성화하실 경우 기밀 데이터가 저장된 디렉토리를 염두하시기 바랍니다.

46.2.6.2.1. 익명 사용자 계정으로 업로드

익명 사용자가 파일을 업로드하는 것을 허용하시려면 /var/ftp/pub/에 쓰기 전용 디렉토리를 생성하시기를 권장합니다.

이를 실행하기 위해 다음 명령을 입력하십시오:

```
mkdir /var/ftp/pub/upload
```

다음으로 익명 계정 사용자가 디렉토리 내의 내용을 보지 못하도록 허가를 변경하기 위해 다음 명령을 입력하십시오:

```
chmod 730 /var/ftp/pub/upload
```

디렉토리 목록은 다음과 같이 나타나야 합니다:

```
drwx-wx---    2 root     ftp          4096 Feb 13 20:05 upload
```

경고

관리자가 익명 사용자가 디렉토리에 읽고 쓸 수 있는 권한을 부여한 경우 종종 도단당한 소프트웨어가 서버에 저장되어 있는 것을 발견하게 됩니다.

vsftpd에서 추가적으로 다음 줄을 /etc/vsftpd/vsftpd.conf 파일에 첨가하십시오:

```
anon_upload_enable=YES
```

46.2.6.3. 사용자 계정

FTP는 비보안 네트워크 상에서 인증을 위해 암호화되지 않은 사용자명과 암호를 전달하기 때문에 시스템 사용자가 사용자 계정을 통해 서버에 접속하는 것을 거부하도록 설정하는 것이 좋습니다.

vsftpd에서 사용자 계정을 비활성화하시려면 /etc/vsftpd/vsftpd.conf 파일에 다음 지시자를 추가하십시오:

```
local_enable=NO
```

46.2.6.3.1. 사용자 계정 제한하기

To disable FTP access for specific accounts or specific groups of accounts, such as the root user and those with sudo privileges, the easiest way is to use a PAM list file as described in 46.1.4.2절. "루트 액세스를 허가하지 않기". The PAM configuration file for vsftpd is /etc/pam.d/vsftpd.

각 서비스에서 직접 사용자 계정을 비활성화시키는 것도 가능합니다.

vsftpd에서 특정 사용자 계정을 비활성화시키려면 /etc/vsftpd.ftpusers에 해당 사용자명을 추가하시면 됩니다.

46.2.6.4. 접근 제어를 위해 TCP 래퍼 사용하기

Use TCP Wrappers to control access to either FTP daemon as outlined in 46.2.1.1절. "TCP 래퍼를 사용하여 보안 강화하기".

46.2.7. Sendmail 보안 강화

Sendmail은 메일 전송 에이전트 (MTA)로서 SMTP (Simple Mail Transport Protocol)을 사용하여 다른 MTA와 이메일 클라이언트나 배달 에이전트 사이에서 전자 메시지를 배달하는 역할을 합니다. 많은 MTA가 서로 주고 받는 트래픽을 암호화 가능하지만, 대부분의 MTA는 그렇지 않습니다. 따라서 공중 네트워크 상에서 이메일을 주고받는 것은 안전하지 못한 통신 방식으로 간주됩니다.

Refer to 25장. Email for more information about how email works and an overview of common configuration settings. This section assumes a basic knowledge of how to generate a valid /etc/mail/sendmail.cf by editing the /etc/mail/sendmail.mc and using the m4 command.

Sendmail 서버를 사용하려고 계획하신다면 다음에 언급된 사항들을 해결하셔야 합니다.

46.2.7.1. 서비스 거부 공격 제한하기

이메일의 특성상, 침입자는 비교적 간단히 서버에 다량의 이메일을 집중적으로 보내어 서비스 거부 현상을 야기시킬 수 있습니다. /etc/mail/sendmail.mc에 다음과 같은 지시자를 제한 설정하여 이와 같은 서비스 거부 공격이 발생할 가능성을 줄일 수 있습니다.

- confCONNECTION_RATE_THROTTLE — 일초당 서버가 받을 수 있는 연결 수를 지정합니다. Sendmail은 기본적으로 연결 수를 제한하지 않습니다. 만일 제한이 설정된 경우 한계수에 이르게 되면 그 후에 들어오는 연결은 지연됩니다.

- confMAX_DAEMON_CHILDREN — 서버에서 배출할 수 있는 최대 자식 프로세스 수. Sendmail은 기본적으로 자식 프로세스의 수를 제한하지 않습니다. 만일 제한이 설정된 경우 한계수에 이르게 되면 그 후에 들어오는 연결은 지연됩니다.

- confMIN_FREE_BLOCKS — 서버가 메일을 수용하는데 필요한 최소 여유 블록의 수. 기본값은 100 블록입니다.

- confMAX_HEADERS_LENGTH — 메시지 헤더의 최대 용량 (바이트 단위)

- confMAX_MESSAGE_SIZE — 한 메시지의 최대 용량 (바이트 단위)

46.2.7.2. NFS와 Sendmail

절대로 메일 스풀 디렉토리인 /var/spool/mail/를 NFS 공유 볼륨에 놓지 마십시오.

Because NFSv2 and NFSv3 do not maintain control over user and group IDs, two or more users can have the same UID, and receive and read each other's mail.

알림

그러나 커베로스를 사용하는 NFSv4의 SECRPC_GSS 커널 모듈이 UID 기반 인증을 사용하지 않으므로 다릅니다. 메일 스풀 디렉토리를 NFS 공유 볼륨에 놓지 마십시오.

46.2.7.3. 메일 전용 사용자

이렇게 로컬 사용자가 Sendmail 서버를 악용하는 것을 방지하기 위하여 메일 사용자는 이메일 프로그램을 사용하여 Sendmail 서버만 사용하도록 설정하시는 것이 좋습니다. 메일 사용자는 메일 서버에서 쉘 계정을 갖지 못하고 /etc/passwd 파일에서 모든 메일 사용자 쉘은 /sbin/nologin으로 설정하셔야 합니다. (루트 사용자 예외)

46.2.8. 청취 중인 포트 확인하기

After configuring network services, it is important to pay attention to which ports are actually listening on the system's network interfaces. Any open ports can be evidence of an intrusion.

네트워크 상에서 청취 중인 포트를 찾아낼 수 있는 두가지 방법이 있습니다. 보다 덜 안정적인 방법으로 netstat -an 또는 lsof -i와 같은 명령을 입력하여 네트워크 스택을 질의하실 수 있습니다. 이 프로그램은 네트워크의 시스템에 연결하지 않고 시스템 상에 무엇이 실행 중인지 확인하기 때문에 신뢰성이 떨어집니다. 따라서 침입자는 종종 이 프로그램을 상대로 침입을 시도합니다. 침입자가 netstat 및 lsof를 자신의 수정된 버전으로 교체하여 권한이 없는 네트워크 포트를 열게된 경우 침입한 자취를 감추는데 이러한 방법을 사용합니다.

보다 안전하게 네트워크 상에서 청취 중인 포트를 확인할 수 있는 방법은 nmap과 같은 포트 스캐너를 사용하는 것입니다.

콘솔에서 다음 명령을 입력하시면 네트워크에서 어느 포트가 TCP 연결을 청취하고 있는지 확인할 수 있습니다:

```
nmap -sT -O localhost
```

이 명령의 결과는 다음과 같이 출력될 것입니다:

```
Starting nmap 3.55 ( http://www.insecure.org/nmap/ ) at 2004-09-24 13:49 EDT
Interesting ports on localhost.localdomain (127.0.0.1):
(The 1653 ports scanned but not shown below are in state: closed)
PORT       STATE SERVICE
22/tcp     open   ssh
25/tcp     open   smtp
111/tcp    open   rpcbind
113/tcp    open   auth
631/tcp    open   ipp
834/tcp    open   unknown
2601/tcp   open   zebra
32774/tcp open   sometimes-rpc11
Device type: general purpose
Running: Linux 2.4.X|2.5.X|2.6.X OS details: Linux 2.5.25 - 2.6.3 or Gentoo 1.2 Linux 2.4.19 rc1-rc7)
Uptime 12.857 days (since Sat Sep 11 17:16:20 2004)
Nmap run completed -- 1 IP address (1 host up) scanned in 5.190 seconds
```

이 출력 결과는 시스템이 sunrpc 서비스가 존재하기 때문에 portmap을 실행 중인것을 보여줍니다. 그러나 포트 834에서 수상한 서비스를 발견할 수 있습니다. 이 포트가 공식적으로 알려진 서비스와 관계있는지 확인해보시려면 다음 명령을 입력하시기 바랍니다:

```
cat /etc/services | grep 834
```

이 명령이 아무런 결과도 출력하지 않습니다. 즉, 포트는 0에서 1023 사이의 범위에 속하지만 루트 권한이 있어야 열 수 있습니다. 따라서 이 포트는 알려진 서비스와 관련되지 않습니다.

다음으로 netstat이나 lsof를 사용하여 포트에 대한 정보를 확인해보시기 바랍니다. netstat을 사용하여 포트 834를 확인하시려면 다음 명령을 사용하십시오:

```
netstat -anp | grep 834
```

명령은 다음과 같은 결과를 출력할 것입니다:

```
tcp    0     0 0.0.0.0:834     0.0.0.0:*     LISTEN     653/ypbind
```

netstat을 사용하여 열려진 포트를 발견하시면 이 포트가 안전하다고 안심하실 수 있습니다. 그 이유는 크래커가 침입한 시스템에서 은밀하게 포트를 연 경우에는 이 명령을 사용하여 발견되지 않도록 설정할 것이기 때문에 포트가 열려져 있다는 것은 이 포트가 안전하다는 것을 의미합니다. 또한 [p] 옵션은 포트를 연 서비스의 프로세스 ID (PID)를 보여줍니다. 이 예시에서 열려진 포트는 portmap 서비스와 함께 사용되는 RPC 서비스인 ypbind (NIS)에 사용됩니다.

lsof 명령도 열려진 포트와 관련된 서비스를 보여주는 기능을 갖추고 있으므로 netstat와 유사한 정보를 보여줍니다:

```
lsof -i | grep 834
```

다음은 출력 결과에서 이 명령과 관련있는 부분입니다:

```
ypbind     653     0    7u   IPv4      1319                TCP *:834 (LISTEN)
ypbind     655     0    7u   IPv4      1319                TCP *:834 (LISTEN)
ypbind     656     0    7u   IPv4      1319                TCP *:834 (LISTEN)
ypbind     657     0    7u   IPv4      1319                TCP *:834 (LISTEN)
```

이러한 도구를 사용하여 시스템 상에서 실행 중인 서비스의 상태에 대한 많은 정보를 얻을 수 있습니다. 이 도구들은 사용이 유연하며 네트워크 서비스와 설정에 대한 광범위한 정보를 제공합니다. 보다 자세한 정보는 lsof, netstat, nmap, services의 메뉴얼 페이지를 참조하시기 바랍니다.

46.3. Single Sign-on (SSO)

46.3.1. Introduction

The Red Hat Enterprise Linux SSO functionality reduces the number of times Red Hat Enterprise Linux desktop users have to enter their passwords. Several major applications leverage the same underlying authentication and authorization mechanisms so that users can log in to Red Hat Enterprise Linux from the log-in screen, and then not need to re-enter their passwords. These applications are detailed below.

In addition, users can log in to their machines even when there is no network (offline mode) or where network connectivity is unreliable, for example, wireless access. In the latter case, services will degrade gracefully.

46.3.1.1. Supported Applications

The following applications are currently supported by the unified log-in scheme in Red Hat Enterprise Linux:

- Login
- Screensaver
- Firefox and Thunderbird

46.3.1.2. Supported Authentication Mechanisms

Red Hat Enterprise Linux currently supports the following authentication mechanisms:

- Kerberos name/password login
- Smart card/PIN login

46.3.1.3. Supported Smart Cards

Red Hat Enterprise Linux has been tested with the Cyberflex e-gate card and reader, but any card that complies with both Java card 2.1.1 and Global Platform 2.0.1 specifications should operate correctly, as should any reader that is supported by PCSC-lite.

Red Hat Enterprise Linux has also been tested with Common Access Cards (CAC). The supported reader for CAC is the SCM SCR 331 USB Reader.

As of Red Hat Enterprise Linux 5.2, Gemalto smart cards (Cyberflex Access 64k v2, standard with DER SHA1 value configured as in PKCSI v2.1) are now supported. These smart cards now use readers compliant with Chip/Smart Card Interface Devices (CCID).

46.3.1.4. Advantages of Red Hat Enterprise Linux Single Sign-on

Numerous security mechanisms currently exist that utilize a large number of protocols and credential stores. Examples include SSL, SSH, IPsec, and Kerberos. Red Hat Enterprise Linux SSO aims to unify these schemes to support the requirements listed above. This does not mean replacing Kerberos with X.509v3 certificates, but rather uniting them to reduce the burden on both system users and the administrators who manage them.

To achieve this goal, Red Hat Enterprise Linux:

- Provides a single, shared instance of the NSS crypto libraries on each operating system.

- Ships the Certificate System's Enterprise Security Client (ESC) with the base operating system. The ESC application monitors smart card insertion events. If it detects that the user has inserted a smart card that was designed to be used with the Red Hat Enterprise Linux Certificate System server product, it displays a user interface instructing the user how to enroll that smart card.

- Unifies Kerberos and NSS so that users who log in to the operating system using a smart card also obtain a Kerberos credential (which allows them to log in to file servers, etc.)

46.3.2. Getting Started with your new Smart Card

Before you can use your smart card to log in to your system and take advantage of the increased security options this technology provides, you need to perform some basic installation and configuration steps. These are described below.

> **Note**
>
> This section provides a high-level view of getting started with your smart card. More detailed information is available in the Red Hat Certificate System Enterprise Security Client Guide.

1. Log in with your Kerberos name and password

2. Make sure you have the nss-tools package loaded.

3. Download and install your corporate-specific root certificates. Use the following command to install the root CA certificate:

   ```
   certutil -A -d /etc/pki/nssdb -n "root ca cert" -t "CT,C,C" \
    -i ./ca_cert_in_base64_format.crt
   ```

4. Verify that you have the following RPMs installed on your system: esc, pam_pkcs11, coolkey, ifd-egate, ccid, gdm, authconfig, and authconfig-gtk.

5. Enable Smart Card Login Support

   a. On the Gnome Title Bar, select System->Administration->Authentication.

   b. Type your machine's root password if necessary.

   c. In the Authentication Configuration dialog, click the Authentication tab.

   d. Select the Enable Smart Card Support check box.

   e. Click the Configure Smart Card... button to display the Smartcard Settings dialog, and specify the required settings:

      - Require smart card for login — Clear this check box. After you have successfully logged in with the smart card you can select this option to prevent users from logging in without a smart card.

      - Card Removal Action — This controls what happens when you remove the smart card after you have logged in. The available options are:

        - Lock — Removing the smart card locks the X screen.

    - Ignore — Removing the smart card has no effect.

6. If you need to enable the Online Certificate Status Protocol (OCSP), open the /etc/pam_pkcs11/pam_pkcs11.conf file, and locate the following line:

   ```
   enable_ocsp = false;
   ```

   Change this value to true, as follows:

   ```
   enable_ocsp = true;
   ```

7. Enroll your smart card

8. If you are using a CAC card, you also need to perform the following steps:

   a. Change to the root account and create a file called /etc/pam_pkcs11/cn_map.

   b. Add the following entry to the cn_map file:

      ```
      MY.CAC_CN.123454 -> myloginid
      ```

      where MY.CAC_CN.123454 is the Common Name on your CAC and myloginid is your UNIX login ID.

9. Logout

46.3.2.1. Troubleshooting

If you have trouble getting your smart card to work, try using the following command to locate the source of the problem:

```
pklogin_finder debug
```

If you run the pklogin_finder tool in debug mode while an enrolled smart card is plugged in, it attempts to output information about the validity of certificates, and if it is successful in attempting to map a login ID from the certificates that are on the card.

46.3.3. How Smart Card Enrollment Works

Smart cards are said to be enrolled when they have received an appropriate certificate signed by a valid Certificate Authority (CA). This involves several steps, described below:

1. The user inserts their smart card into the smart card reader on their workstation. This event is recognized by the Enterprise Security Client (ESC).

2. The enrollment page is displayed on the user's desktop. The user completes the required details and the user's system then connects to the Token Processing System (TPS) and the CA.

3. The TPS enrolls the smart card using a certificate signed by the CA.

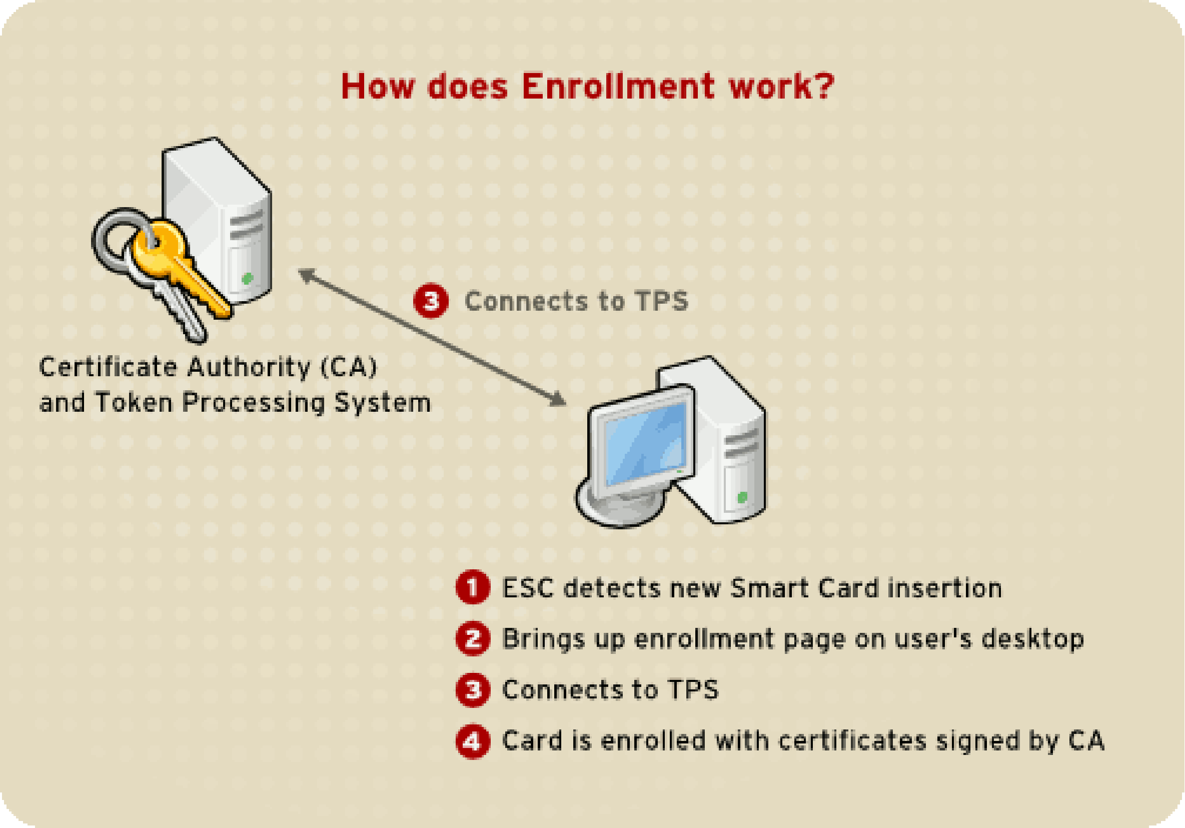
How does Enrollment work?
Certificate Authority (CA)
and Token Processing System
3 Connects to TPS
1 ESC detects new Smart Card insertion
2 Brings up enrollment page on user's desktop
3 Connects to TPS
4 Card is enrolled with certificates signed by CA

그림 46.4. How Smart Card Enrollment Works

46.3.4. How Smart Card Login Works

This section provides a brief overview of the process of logging in using a smart card.

1. When the user inserts their smart card into the smart card reader, this event is recognized by the PAM facility, which prompts for the user's PIN.

2. The system then looks up the user's current certificates and verifies their validity. The certificate is then mapped to the user's UID.

3. This is validated against the KDC and login granted.

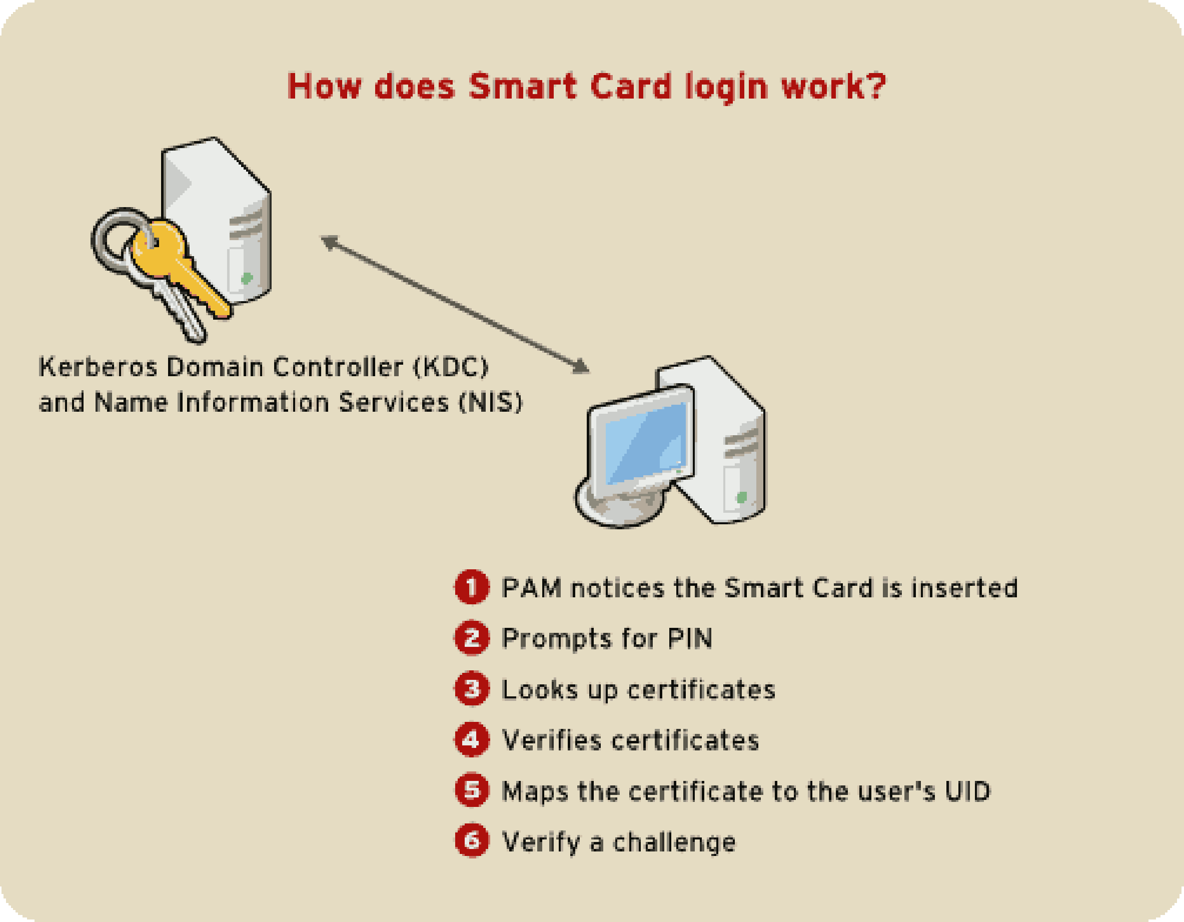
How does Smart Card login work?
Kerberos Domain Controller (KDC)
and Name Information Services (NIS)
1 PAM notices the Smart Card is inserted
2 Prompts for PIN
3 Looks up certificates
4 Verifies certificates
5 Maps the certificate to the user's UID
6 Verify a challenge

그림 46.5. How Smart Card Login Works

Note

You cannot log in with a card that has not been enrolled, even if it has been formatted. You need to log in with a formatted, enrolled card, or not using a smart card, before you can enroll a new card.

Refer to 46.6절. "Kerberos" and 46.4절. "PAM (Pluggable Authentication Modules)" for more information on Kerberos and PAM.

46.3.5. Configuring Firefox to use Kerberos for SSO

You can configure Firefox to use Kerberos for Single Sign-on. In order for this functionality to work correctly, you need to configure your web browser to send your Kerberos credentials to the appropriate KDC.The following section describes the configuration changes and other requirements to achieve this.

1. In the address bar of Firefox, type about:config to display the list of current configuration options.

2. In the Filter field, type negotiate to restrict the list of options.

3. Double-click the network.negotiate-auth.trusted-uris entry to display the Enter string value dialog box.

4. Enter the name of the domain against which you want to authenticate, for example, .example.com.

5. Repeat the above procedure for the network.negotiate-auth.delegation-uris entry, using the same domain.

> **Note**
>
> You can leave this value blank, as it allows Kerberos ticket passing, which is not required.
>
> If you do not see these two configuration options listed, your version of Firefox may be too old to support Negotiate authentication, and you should consider upgrading.

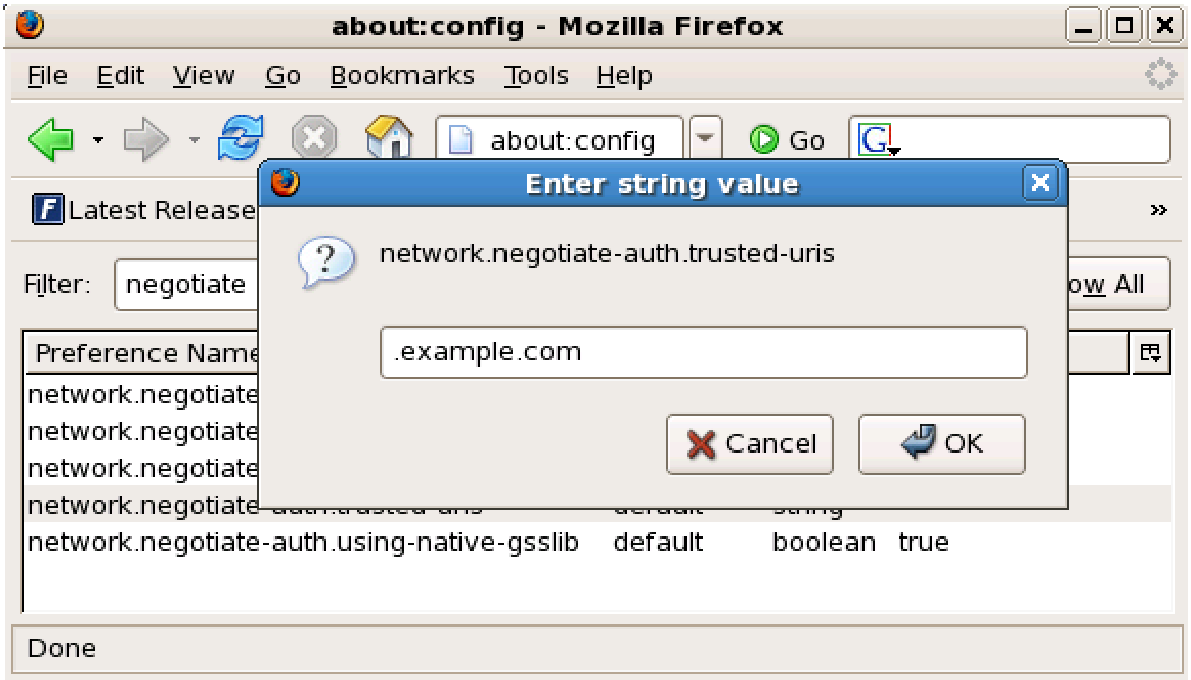

Enter string value

network.negotiate-auth.trusted-uris

.example.com

```
network.negotiate
network.negotiate
network.negotiate
network.negotiate
network.negotiate-auth.using-native-gsslib   default   boolean   true
```

그림 46.6. Configuring Firefox for SSO with Kerberos

You now need to ensure that you have Kerberos tickets. In a command shell, type kinit to retrieve Kerberos tickets. To display the list of available tickets, type klist. The following shows an example output from these commands:

```
~]$ kinit
Password for user@EXAMPLE.COM:

~]$ klist
Ticket cache: FILE:/tmp/krb5cc_10920
Default principal: user@EXAMPLE.COM

Valid starting      Expires             Service principal
10/26/06 23:47:54  10/27/06 09:47:54  krbtgt/USER.COM@USER.COM
        renew until 10/26/06 23:47:54

Kerberos 4 ticket cache: /tmp/tkt10920
klist: You have no tickets cached
```

46.3.5.1. Troubleshooting

If you have followed the configuration steps above and Negotiate authentication is not working, you can turn on verbose logging of the authentication process. This could help you find the cause of the problem. To enable verbose logging, use the following procedure:

1. Close all instances of Firefox.

2. Open a command shell, and enter the following commands:

```
export NSPR_LOG_MODULES=negotiateauth:5
export NSPR_LOG_FILE=/tmp/moz.log
```

3. Restart Firefox from that shell, and visit the website you were unable to authenticate to earlier. Information will be logged to /tmp/moz.log, and may give a clue to the problem. For example:

```
-1208550944[90039d0]: entering nsNegotiateAuth::GetNextToken()
-1208550944[90039d0]: gss_init_sec_context() failed: Miscellaneous failure
No credentials cache found
```

   This indicates that you do not have Kerberos tickets, and need to run kinit.

If you are able to run kinit successfully from your machine but you are unable to authenticate, you might see something like this in the log file:

```
-1208994096[8d683d8]: entering nsAuthGSSAPI::GetNextToken()
-1208994096[8d683d8]: gss_init_sec_context() failed: Miscellaneous failure
Server not found in Kerberos database
```

This generally indicates a Kerberos configuration problem. Make sure that you have the correct entries in the [domain_realm] section of the /etc/krb5.conf file. For example:

```
.example.com = EXAMPLE.COM
example.com = EXAMPLE.COM
```

If nothing appears in the log it is possible that you are behind a proxy, and that proxy is stripping off the HTTP headers required for Negotiate authentication. As a workaround, you can try to connect to the server using HTTPS instead, which allows the request to pass through unmodified. Then proceed to debug using the log file, as described above.

46.4. PAM (Pluggable Authentication Modules)

Programs that grant users access to a system use authentication to verify each other's identity (that is, to establish that a user is who they say they are).

일반적으로, 각각의 프로그램은 고유한 사용자 인증 방식을 갖고 있습니다. Red Hat Enterprise Linux에서 대부분의 프로그램은 PAM (Pluggable Authentication Modules)이라고 불리우는 중앙 집중 인증 메카니즘을 사용하여 설정됩니다.

PAM은 장착가능한 모듈러 아키텍쳐를 사용하여 시스템 관리자에게 시스템의 인증 정책을 설정하는 데에 있어서 유연성을 제공합니다.

In most situations, the default PAM configuration file for a PAM-aware application is sufficient. Sometimes, however, it is necessary to edit a PAM configuration file. Because misconfiguration of PAM can compromise system security, it is important to understand the structure of these files before making any modifications. Refer to 46.4.3절. "PAM 설정 파일 포멧" for more information.

46.4.1. PAM의 장점

PAM에는 다음과 같은 장점이 있습니다:

- 다양한 응용 프로그램과 함께 사용할 수 있는 일반적인 인증 설계.
- 인증에 있어서 시스템 관리자와 프로그램 개발자를 위한 유연성 및 제어 기능.
- 인증 설계를 생성하지 않고 개발자가 프로그램에 쓰는 것을 허용하는 완전 문서화된 단독 라이브러리.

46.4.2. PAM 설정 파일

/etc/pam.d/ 디렉토리에는 각각의 PAM-aware 응용 프로그램에 대한 PAM 설정 파일이 있습니다. 이전 버전의 PAM에는 /etc/pam.conf 파일이 사용되었으나, 현재 이 파일은 삭제되었으며 이는 /etc/pam.d/ 디렉토리가 존재하지 않을 경우에만 사용됩니다.

46.4.2.1. PAM 서비스 파일

각각의 PAM-aware 응용 프로그램이나 서비스는 /etc/pam.d/ 디렉토리에 파일을 갖습니다. 이 디렉토리에 있는 각각의 파일은 액세스를 제어하는 서비스와 같은 이름을 갖습니다.

PAM-aware 프로그램은 자신의 서비스명을 정의하고 /etc/pam.d/ 디렉토리에 PAM 설정 파일을 설치합니다. 예를 들어, login 프로그램은 login으로 자신의 서비스명을 정의하고 /etc/pam.d/login PAM 설정 파일을 설치합니다.

46.4.3. PAM 설정 파일 포멧

각각의 PAM 설정 파일에는 다음과 같이 포멧된 지시문 그룹이 있습니다:

```
<module interface>  <control flag>   <module name>   <module arguments>
```

이러한 각각의 요소에 대해서는 다음 부분에서 설명합니다.

46.4.3.1. 모듈 인터페이스

현재 네 가지 유형의 PAM 모듈 인터페이스가 사용 가능합니다. 각각의 유형은 권한 부여 과정의 다른 양상을 띠고 있습니다.

- auth — 이러한 모듈 인터페이스는 사용을 인증합니다. 예를 들어, 이는 암호의 유효성을 요청 및 확인합니다. 이러한 인터페이스를 갖는 모듈은 그룹 멤버쉽이나 커베로스 티켓과 같이 인증을 설정할 수 있습니다.
- account — 이러한 모듈 인터페이스는 액세스를 허용하는 지를 확인합니다. 예를 들어, 사용자 계정이 만료되었는지 또는 특정 시간에 사용자의 로그인이 허용되는지를 확인합니다.
- password — 이러한 모듈 인터페이스는 사용자 암호를 변경하는 데 사용됩니다.
- session — This module interface configures and manages user sessions. Modules with this interface can also perform additional tasks that are needed to allow access, like mounting a user's home directory and making the user's mailbox available.

주의

개별 모듈은 모든 모듈 인터페이스를 제공할 수 있습니다. 예를 들어, pam_unix.so는 네 가지 모듈 인터페이스 모두를 제공합니다.

PAM 설정 파일에서, 모듈 인터페이스는 처음으로 정의되어야 할 영역입니다. 예를 들어, 설정에서 나타나는 전형적인 줄은 다음과 같습니다:

```
auth required pam_unix.so
```

This instructs PAM to use the pam_unix.so module's auth interface.

46.4.3.1.1. 모듈 인터페이스 스택하는 중

Module interface directives can be stacked, or placed upon one another, so that multiple modules are used together for one purpose. If a module's control flag uses the "sufficient" or "requisite" value (refer to 46.4.3.2절. “제어 플래그” for more information on these flags), then the order in which the modules are listed is important to the authentication process.

스택하는 것은 사용자 인증을 허용하기 전에 존재하는 특정 사항을 관리자가 쉽게 필요로하게 합니다.예를 들어, 일반적으로 reboot 명령은 PAM 설정 파일에서 볼 수 있듯이 스택된 여러 모듈을 사용합니다.

```
~]# cat /etc/pam.d/reboot
#%PAM-1.0
auth sufficient pam_rootok.so
auth required pam_console.so
#auth include system-auth
account required pam_permit.so
```

- 첫 번째 줄은 주석으로 실행되지 않습니다.
- auth sufficient pam_rootok.so — 이 줄은 pam_rootok.so 모듈을 사용하여 사용자의 UID가 0임을 확인함으로써 사용자가 현재 루트에 있는 지를 확인합니다. 이러한 테스트가 성공적으로 이루어지면, 더이상 다른 모듈이 제시되지 않으며 명령이 실행됩니다. 테스트가 실패할 경우, 다음 모듈이 제시됩니다.
- auth required pam_console.so — 이 줄은 pam_console.so 모듈을 사용하여 사용자 인증을 확인합니다. 사용자가 이미 콘솔에 로그인되어 있을 경우, pam_console.so 파일은 /etc/security/console.apps/ 디렉토리에 서비스명과 같은 이름을 가진 파일이 있는 지를 확인합니다 (재부팅). 이러한 파일이 존재할 경우, 성공적으로 인증이 이루어지며 다음 모듈로 제어권이 넘어갑니다.
- #auth include system-auth — 이 줄은 주석으로 실행되지 않습니다.
- account required pam_permit.so — 이 줄은 pam_permit.so 모듈을 사용하여 루트 사용자나 또는 시스템에 재부팅하기 위해 콘솔에 로그인하는 사용자를 허용합니다.

46.4.3.2. 제어 플래그

모든 PAM 모듈은 호출되었을 때 성공 또는 실패 결과를 생성합니다. 제어 플래그는 PAM에게 결과에 어떻게 대응할 지를 지시합니다. 모듈은 특정한 순서로 스택될 수 있으며, 제어 플래그는 특정 모듈의 성공 또는 실패 결과가 사용자를 서비스에 인증하게 하는 전체 목적에 얼마나 중요한 지를 결정합니다.

네 개의 미리 정의된 제어 플래그입니다:

- required — 모듈 결과가 반드시 성공적으로 이루어져야 인증 작업을 계속 진행할 수 있습니다. 이 시점에서 모듈 테스트를 실패할 경우, 인터페이스를 참조하는 모든 모듈 테스트의 결과가 완료될 때 까지 사용자는 인식되지 않게 됩니다.

- requisite — 모듈 결과가 반드시 성공적으로 이루어져야 인증 작업을 계속 진행할 수 있습니다. 하지만, 이 시점에서 모듈 테스트를 실패할 경우, 사용자는 처음으로 실패한 required 또는 requisite 모듈 테스트를 반영하는 메세지와 함께 바로 인식됩니다.

- sufficient — 모듈 테스트를 실패할 경우, 모듈 결과는 무시됩니다. 하지만, sufficient로 플래그된 모듈의 결과가 성공적으로 이루어지고 required로 플래그된 이전 모듈이 실패하지 않은 경우, 다른 결과가 필요하지 않게 되며 사용자는 서비스에 인증됩니다.

- optional — 모듈 결과가 무시됩니다. 다른 모듈이 인터페이스를 참조하지 않을 때 optional로 플래그된 모듈만이 성공적인 인증 작업을 위해 필요합니다.

중요 사항

required 모듈이 호출되는 순서는 중요하지 않습니다. sufficient 및 requisite 제어 플래그만이 순서가 중요하게 됩니다.

보다 정교하게 제어할 수 있는 새로운 제어 플래그 구문이 현재 PAM에서 사용 가능합니다.

The pam.d man page, and the PAM documentation, located in the /usr/share/doc/pam-<version-number>/ directory, where <version-number> is the version number for PAM on your system, describe this newer syntax in detail.

46.4.3.3. 모듈명

모듈명은 특정 모듈 인터페이스를 포함하는 장착식 모듈 명과 함께 PAM을 제공합니다. Red Hat Enterprise Linux 버전 순서에서, 모듈로의 완전한 경로는 PAM 설정 파일에 제공되지만, /lib64/security/ 디렉토리에 있는 64비트 PAM 모듈을 저장할 수 있는 multilib 시스템이 나타나면서, 디렉토리명이 생략되었습니다. 이는 응용 프로그램이 올바른 모듈 버전을 배치할 수 있는 libpam 버전에 링크되기 때문입니다.

46.4.3.4. 모듈 인자

PAM은 몇몇 모듈에 대해 인증하는 동안 인자를 사용하여 장착식 모듈에 정보를 전달합니다.

예를 들어, pam_userdb.so 모듈은 Berkeley DB 파일에 저장된 정보를 사용하여 사용자를 인증합니다. Berkeley DB는 많은 응용 프로그램에 내장된 오픈 소스 데이터베이스 시스템입니다. 모듈이 db 인자를 갖음으로서 Berkeley DB는 어떤 데이터베이스가 요청된 서비스에 사용되는 지를 알게 됩니다.

The following is a typical pam_userdb.so line in a PAM configuration. The <path-to-file> is the full path to the Berkeley DB database file:

```
auth required pam_userdb.so db=<path-to-file>
```

잘못된 인자는 일반적으로 무시되며 PAM 모듈의 성공 또는 실패 결과에 영향을 미치지 않습니다. 하지만, 몇개의 모듈은 잘못된 인자로 인해 실패될 수 도 있습니다. 대부분의 모듈은 /var/log/secure 파일에 오류를 보고합니다.

46.4.4. PAM 설정 파일의 예

다음은 PAM 응용 프로그램 설정 파일의 예입니다:

```
#%PAM-1.0
auth required   pam_securetty.so
auth required   pam_unix.so nullok
auth required   pam_nologin.so
account required   pam_unix.so
password required   pam_cracklib.so retry=3
password required   pam_unix.so shadow nullok use_authtok
session required   pam_unix.so
```

- 첫 번째 줄은 주석으로, 줄은 해쉬 표시 (#)로 시작합니다.
- 로그인 인증을 위해 4개의 스택 3개의 모듈을 통해 두 줄을 만듭니다.

  auth required pam_securetty.so — 이 모듈은 사용자가 루트로 로그인하려 할 경우, 사용자가 로그인되어 있는 tty는 파일이 존재할 경우 /etc/securetty 파일 목록에 있게 됩니다.

  tty가 파일 목록에 없을 경우, Login incorrect 메세지와 함께 루트로 로그인을 실패하게 됩니다.

  auth required pam_unix.so nullok — 이 모듈은 사용자에게암호를 요구하며 /etc/passwd와 /etc/shadow (이 파일이 존재할 경우)에 저장된 정보를 사용하여 암호를 확인합니다.

  In the authentication phase, the pam_unix.so module automatically detects whether the user's password is in the passwd file or the shadow file. Refer to 35.6절. "Shadow Passwords" for more information.

  - nullok 인자는 공백 암호를 허용하는 pam_unix.so 모듈을 지시합니다.

- auth required pam_nologin.so — 마지막 인증 단계로 /etc/nologin 파일이 존재하는 지를 확인합니다. 파일이 존재하고 사용자가 루트에 있지 않을 경우, 인증 실패합니다.

> **주의**
>
> 이 예시에서, 첫 번째 auth 모듈이 실패했지만 세 개의 모든 auth 모듈이 확인되었습니다. 이는 사용자가 어떤 단계에서 자신의 인증이 실패되었는지를 알지 못하게 합니다. 이러한 내용이 침입자의 손에 들어가면 침입자가 시스템에 침입하는 방법을 보다 쉽게 추론하게 할 수 있습니다.

- account required pam_unix.so — 이 모듈은 필요한 계정 확인 작업을 실행합니다. 예를 들어, 쉐도우 암호가 활성화되었을 경우, pam_unix.so 모듈의 계정 인터페이스는 계정이 만료되었는지 또는 사용자가 허용된 유예 기간 안에 암호를 변경하였는지를 확인합니다.

- password required pam_cracklib.so retry=3 — 암호가 만료된 경우, pam_cracklib.so 모듈의 암호 구성 요소는 새로운 암호를 만들 것을 요구합니다. 그 후에 새로 생성된 암호에 대해 사전 기반 암호 해킹 프로그램에 의해 쉽게 결정될 수 있는 지를 테스트합니다.

  - retry=3 인수는 테스트가 처음으로 실패했는 지를 지정하고,사용자에게 강력한 암호를 생성할 수 있는 두 번의 기회가 주어집니다.

- password required pam_unix.so shadow nullok use_authtok — This line specifies that if the program changes the user's password, it should use the password interface of the pam_unix.so module to do so.

  - The argument shadow instructs the module to create shadow passwords when updating a user's password.

  - nullok 인수는 모듈을 지시하여 사용자가 공백 암호에서 자신의 암호를 변경할 수 있게 허용합니다. 그렇지 않을 경우, 공백 암호는 계정 잠금으로 다루어 집니다.

  - 이 줄에 있는 마지막 인수인 use_authtok는 PAM 모듈을 스택할 때 순서의 중요함을 보여주는 예 입니다. 이러한 인수는 모듈이 사용자에게 새로운 암호를 요구하지 않도록 지시합니다. 대신, 이전 암호 모듈에 의해 기록된 암호중 아무것이나 허용합니다. 이러한 방법에서, 모든 새로운 암호는 암호의 보안을 위해 암호를 허용하기 전pam_cracklib.so 테스트를 거쳐야 합니다.

- session required pam_unix.so — 마지막 줄은 세션을 관리하기 위한 pam_unix.so 모듈의 세션 인터페이스를 지시합니다. 이 모듈은 각 세션의 시작과 마지막에 사용자명과 서비스 유형을 /var/log/secure에 기록합니다. 이 모듈은 다른 세션 모듈을 사용하여 스택함으로서 추가 기능을 보완할 수 있습니다.

46.4.5. PAM 모듈 생성

PAM-aware 응용 프로그램을 사용하여 언제든지 새로운 PAM 모듈을 생성하거나 추가하실 수 있습니다.

예를 들어, 개발자는 일회용 암호 생성 방식을 만들어 이를 지원하기 위해 PAM 모듈을 기록할 수도 있습니다. PAM-aware 프로그램은 새로운 모듈과 암호를 재컴파일하거나 수정하지 않고 즉시 사용할 수 있습니다.

이는 개발자와 시스템 관리자에게 다른 프로그램에 대해 이를 재컴파일하지 않고 인증 방식을 테스트함은 물론 혼합하고 붙일수 있게 합니다.

Documentation on writing modules is included in the /usr/share/doc/pam-<version-number>/ directory, where <version-number> is the version number for PAM on your system.

46.4.6. PAM 및 관리자 인증 캐싱

Red Hat Enterprise Linux에 있는 그래픽 관리자 도구의 수는 pam_timestamp.so 모듈을 사용하여 최대 5분 동안 사용자에게 극도의 권한을 제공합니다. pam_timestamp.so 파일이 작동하고 있는 동안 터미널에서 떨어진 사용자가 콘솔로 물리적으로 액세스하려는 누군가에 의한 조작으로 인해 컴퓨터가 열린 채로 있을 수 있으므로 이러한 메카니즘이 어떻게 작동하는 지를 이해하는 것이 중요합니다.

PAM 타임스탬프 설계에서, 그래픽 관리자 응용프로그램은이 시작할 때 이는 사용자에게 루트 암호를 요청합니다. 사용자가 인증 확인되면, pam_timestamp.so 모듈은 타임스탬프 파일을 생성합니다. 이는 /var/run/sudo/ 디렉토리에 기본으로 생성됩니다. 타임스탬프 파일이 이미 존재할 경우, 그래픽 관리자 프로그램은 암호를 요청하지 않고 대신, pam_timestamp.so 모듈은 타임스탬프 파일을 새롭게 하고 사용자를 위해 문제가 되지 않는 관리자 액세스에 대한 5분의 여유 시간을 보유해 둡니다.

You can verify the actual state of the timestamp file by inspecting the /var/run/sudo/<user> file. For the desktop, the relevant file is unknown:root. If it is present and its timestamp is less than five minutes old, the credentials are valid.

인증 아이콘이 타임스탬프 파일의 존재를 보여주며, 이는 패널의 알림 상자에 나타납니다.

그림 46.7. 인증 아이콘

46.4.6.1. 타임스탬프 파일 삭제 중

PAM 타임스탬프가 활성화되는 곳에서 콘솔을 배제하기 전에, 타임스탬프 파일을 삭제시킬 것을 권장합니다. 그래픽 환경에서 이를 실행하기 위해, 패널 상의 인증 아이콘을 클릭하면 대화 상자가 나타납니다. 관리자 권한 해제하기 버튼을 클릭하여 활성화된 타임스탬프 파일을 삭제합니다.

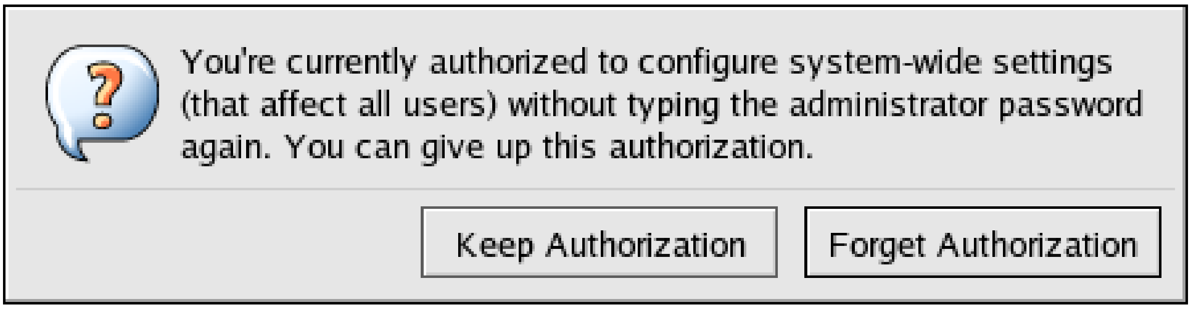

You're currently authorized to configure system-wide settings (that affect all users) without typing the administrator password again. You can give up this authorization.

그림 46.8. 인증 다이얼로그 해제

PAM 타임스탬프 파일에 대해 다음 사항을 주의하시기 바랍니다:

- ssh를 사용하여 원격으로 시스템에 로그인할 경우, /sbin/pam_timestamp_check -k root 명령을 사용하여 타임스탬프 파일을 삭제합니다.
- 권한이 있는 응용 프로그램을 시작하는 것으로 부터 같은 터미널 윈도우에서 /sbin/pam_timestamp_check -k root 명령을 실행하셔야합니다.
- /sbin/pam_timestamp_check -k 명령을 사용하시려면pam_timestamp.so 모듈을 불러내는 사용자로 로그인하셔야 합니다. 이 명령을 사용하기 위해 루트로 로그인하지 마십시오.
- 데스크탑 상에서 (아이콘에 있는 권한부여 해제하기 작업을 사용하지 않고) 인증을 종료하시려면, 다음 명령을 사용하시기 바랍니다:

  ```
  pam_timestamp_check -k root </dev/null >/dev/null 2>/dev/null
  ```

  이 명령의 사용을 실패는 명령을 실행하신 곳에 있는 pty 에서 (만일 있을 경우) 인증을 삭제하게 됩니다.

pam_timestamp_check를 사용하여 타임스탬프 파일을 삭제하는 것에 관한 보다 자세한 정보는 pam_timestamp_check 메뉴얼 페이지를 참조하시기 바랍니다.

46.4.6.2. 일반적인 pam_timestamp 지시문

pam_timestamp.so 모듈은 여러 지시문을 허용합니다. 다음은 가장 일반적으로 사용되는 두 개의 옵션입니다:

- timestamp_timeout — 타임스탬프 파일이 유효한 기간을 (초 단위) 지정합니다. 기본 값은 300초 (5분)입니다.

- timestampdir — 타임스탬프 파일이 저장된 디렉토리를 지정합니다. 기본 값은 /var/run/sudo/입니다.

Refer to 46.4.8.1절. "설치된 문서" for more information about controlling the pam_timestamp.so module.

46.4.7. PAM 및 장치 소유권

Red Hat Enterprise Linux에서, 컴퓨터의 물리적 콘솔에 로그인한 첫 번째 사용자는 특정한 장치를 조작할 수 있으며 루트 사용자을 위해 보유된 특정 작업을 실행할 수 있습니다. 이는 pam_console.so라고 불리는 PAM 모듈에 의해 제어됩니다.

46.4.7.1. 장치 소유권

사용자가 Red Hat Enterprise Linux 시스템에 로그인할 때, pam_console.so 모듈은 login 또는 그래픽 로그인 프로그램, gdm, kdm, xdm에 의해 호출됩니다. 이러한 사용자가 물리적 콘솔에 로그인하는 첫 번째 사용자일 경우 — console user로 불려짐 — 모듈은 일반적으로 루트 사용자에 의해 소유되었던 여러 장치의 사용자 소유권이 확장됩니다. 콘솔 사용자는 사용자의 마지막 로컬 세션이 끝날때 까지 이러한 장치를 소유합니다. 이러한 사용자가 로그 아웃한 후, 장치의 소유권은 루트 사용자에게로 복귀됩니다.

영향을 미치는 장치에는 사운드 카드, 디스켓 드라이브, CD-ROM 드라이브가 포함되지만, 이에 제한되지는 않습니다.

이러한 장치는 로컬 사용자가 루트 액세스 없이 이러한 장치를 조작하는 것을 허용하므로, 콘솔 사용자를 위한 일반적인 작업을 단순화시킬 수 있습니다.

다음의 파일을 편집하여 pam_console.so에 의해 제어되는 장치 목록을 수정하실 수 있습니다.

- /etc/security/console.perms

- /etc/security/console.perms.d/50-default.perms

위의 파일 목록보다는 다른 장치의 사용 권한을 변경하시거나 또는 특정 기본값으로 덮어쓰실 수 있습니다. 50-default.perms 파일을 수정하는 대신, 새로운 파일을 생성하시고 (예, xx-name.perms) 필요한 사항을 수정하시기 바랍니다. 새로운 기본 파일명은 50 이상의 숫자로 시작해야 합니다. (예, 51-default.perms) 이는 50-default.perms 파일에 있는 기본값을 덮어 쓰게 됩니다.

경고

If the gdm, kdm, or xdm display manager configuration file has been altered to allow remote users to log in and the host is configured to run at runlevel 5, it is advisable to change the <console> and <xconsole> directives in the /etc/security/console.perms to the following values:

```
<console>=tty[0-9][0-9]* vc/[0-9][0-9]* :0\.[0-9] :0
<xconsole>=:0\.[0-9] :0
```

이는 원격 사용자가 컴퓨터 상의 장치 및 제한된 응용 그로그램에 액세스하지 못하게 합니다.

If the gdm, kdm, or xdm display manager configuration file has been altered to allow remote users to log in and the host is configured to run at any multiple user runlevel other than 5, it is advisable to remove the <xconsole> directive entirely and change the <console> directive to the following value:

```
<console>=tty[0-9][0-9]* vc/[0-9][0-9]*
```

46.4.7.2. 응용 프로그램 액세스

콘솔 사용자는 /etc/security/console.apps/ 디렉토리에서 사용하도록 설정된 특정 프로그램에도 액세스합니다.

이 디렉토리에는 콘솔 사용자가 /sbin and /usr/sbin에 있는 특정 응용 프로그램을 실행할 수 있는 설정 파일이 포함되어 있습니다.

이 설정 파일에는 콘솔 사용자가 설정한 응용 프로그램과 같은 이름이 있습니다.

콘솔 사용자가 액세스한 주목할만한 응용 프로그램 그룹 중 하나는 시스템을 종료하거나 재부팅하는 세가지 프로그램입니다:

- /sbin/halt
- /sbin/reboot
- /sbin/poweroff

이는 PAM-aware 응용 프로그램이므로, 사용을 위한 기본 조건으로서 pam_console.so 모듈을 호출합니다.

Refer to 46.4.8.1절. “설치된 문서” for more information.

46.4.8. 추가 자료

다음의 자료는 PAM을 사용하고 설정하는 방식에 대해 보다 자세하게 설명합니다. 이러한 자료에 더하여, 시스템에 있는 PAM 설정 파일을 읽어보시면 PAM의 구성 방식에 대해 보다 명확하게 이해하실 수 있습니다.

46.4.8.1. 설치된 문서

- PAM과 관련된 메뉴얼 페이지 — PAM과 관련된 다양한 응용프로그램 및 설정 파일에 대한 여러 메뉴얼 페이지가 있습니다. 다음은 보다 중요한 메뉴얼 페이지의 목록입니다.

  설정 파일

  - pam — PAM 설정 파일에 대한 구조 및 목적을 포함하는 PAM에 관한 개요.

    이 메뉴얼 페이지에서는 /etc/pam.d/ 디렉토리에 있는 /etc/pam.conf 및 개별 설정 파일에 관해 설명하고 있음을 유의하시기 바랍니다. 기본적으로 Red Hat Enterprise Linux에서는 /etc/pam.d/ 디렉토리에 있는 개별 설정 파일을 사용하며, /etc/pam.conf 파일이 있을지라도 이를 무시합니다.

  - pam_console — pam_console.so 모듈의 목적에 대해 설명하며, PAM 설정 파일 안에 있는 항목에 대한 올바른 구문에 대해서도 설명합니다.

  - console.apps — PAM에 의해 할당된 콘솔 사용자가 액세스할 수 있는 응용 프로그램에는 어떤 것이 있는 지를 정의하는 /etc/security/console.apps 설정 파일에서 사용 가능한 포맷 및 옵션에 대해 설명합니다.

  - console.perms — PAM에 의해 할당된 콘솔 사용자 허용을 지정하는 /etc/security/console.perms 설정 파일에서 사용 가능한 포맷 및 옵션에 대해 설명합니다.

  - pam_timestamp — pam_timestamp.so 모듈에 대해 설명합니다.

- /usr/share/doc/pam-<version-number> — Contains a System Administrators' Guide, a Module Writers' Manual, and the Application Developers' Manual, as well as a copy of the PAM standard, DCE-RFC 86.0, where <version-number> is the version number of PAM.

- /usr/share/doc/pam-<version-number>/txts/README.pam_timestamp — Contains information about the pam_timestamp.so PAM module, where <version-number> is the version number of PAM.

46.4.8.2. 유용한 웹사이트

- http://www.kernel.org/pub/linux/libs/pam/ — Linux-PAM 프로젝트에 대한 주요 웹사이트로 이에는 다양한 PAM 모듈에 관한 정보와 FAQ 그리고 추가적 PAM 문서가 포함되어 있습니다.

> 주의
>
> 위의 웹사이트에 있는 문서는 마지막으로 출시된 새로운 PAM 버전으로 Red Hat Enterprise Linux에 포함된 PAM 버전과 100% 일치하지 않을 수 도 있습니다.

46.5. TCP Wrappers and xinetd

Controlling access to network services is one of the most important security tasks facing a server administrator. Red Hat Enterprise Linux provides several tools for this purpose. For example, an iptables-based firewall filters out unwelcome network packets within the kernel's network stack. For network services that utilize it, TCP Wrappers add an additional layer of protection by defining which hosts are or are not allowed to connect to "wrapped" network services. One such wrapped network service is the xinetd super server. This service is called a super server because it controls connections to a subset of network services and further refines access control.

그림 46.9. “Access Control to Network Services” is a basic illustration of how these tools work together to protect network services.

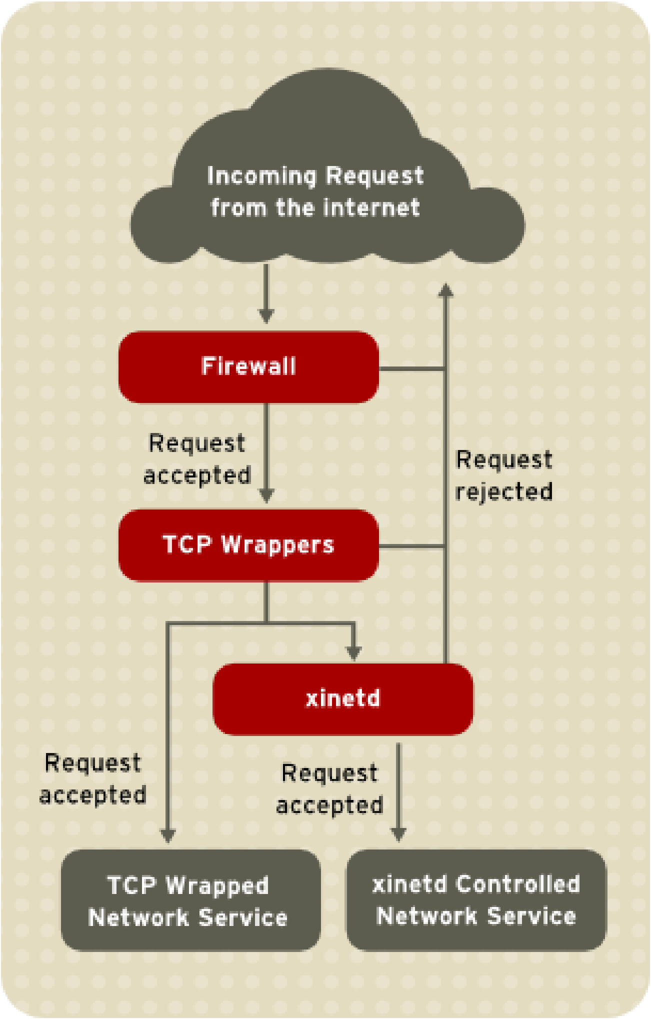
Incoming Request from the internet
Firewall
Request accepted
Request rejected
TCP Wrappers
xinetd
Request accepted
Request accepted
TCP Wrapped Network Service
xinetd Controlled Network Service

그림 46.9. Access Control to Network Services

This chapter focuses on the role of TCP Wrappers and xinetd in controlling access to network services and reviews how these tools can be used to enhance both logging and utilization management. Refer to 46.9절. “IPTables” for information about using firewalls with iptables.

46.5.1. TCP Wrappers

The TCP Wrappers package (tcp_wrappers) is installed by default and provides host-based access control to network services. The most important component within the package is the /usr/lib/libwrap.a library. In general terms, a TCP-wrapped service is one that has been compiled against the libwrap.a library.

When a connection attempt is made to a TCP-wrapped service, the service first references the host's access files (/etc/hosts.allow and /etc/hosts.deny) to determine whether or not the client is allowed to connect. In most cases, it then uses the syslog daemon (syslogd) to write the name of the requesting client and the requested service to /var/log/secure or /var/log/messages.

If a client is allowed to connect, TCP Wrappers release control of the connection to the requested service and take no further part in the communication between the client and the server.

In addition to access control and logging, TCP Wrappers can execute commands to interact with the client before denying or releasing control of the connection to the requested network service.

Because TCP Wrappers are a valuable addition to any server administrator's arsenal of security tools, most network services within Red Hat Enterprise Linux are linked to the libwrap.a library. Some such applications include /usr/sbin/sshd, /usr/sbin/sendmail, and /usr/sbin/xinetd.

Note

To determine if a network service binary is linked to libwrap.a, type the following command as the root user:

```
ldd <binary-name> | grep libwrap
```

Replace <binary-name> with the name of the network service binary.

If the command returns straight to the prompt with no output, then the network service is not linked to libwrap.a.

The following example indicates that /usr/sbin/sshd is linked to libwrap.a:

```
~]# ldd /usr/sbin/sshd | grep libwrap
        libwrap.so.0 => /usr/lib/libwrap.so.0 (0x00655000)
~]#
```

46.5.1.1. Advantages of TCP Wrappers

TCP Wrappers provide the following advantages over other network service control techniques:

- Transparency to both the client and the wrapped network service — Both the connecting client and the wrapped network service are unaware that TCP Wrappers are in use. Legitimate users are logged and connected to the requested service while connections from banned clients fail.

- Centralized management of multiple protocols — TCP Wrappers operate separately from the network services they protect, allowing many server applications to share a common set of access control configuration files, making for simpler management.

46.5.2. TCP Wrappers Configuration Files

To determine if a client is allowed to connect to a service, TCP Wrappers reference the following two files, which are commonly referred to as hosts access files:

- /etc/hosts.allow

- /etc/hosts.deny

When a TCP-wrapped service receives a client request, it performs the following steps:

1. It references /etc/hosts.allow. — The TCP-wrapped service sequentially parses the /etc/hosts.allow file and applies the first rule specified for that service. If it finds a matching rule, it allows the connection. If not, it moves on to the next step.

2. It references /etc/hosts.deny. — The TCP-wrapped service sequentially parses the /etc/hosts.deny file. If it finds a matching rule, it denies the connection. If not, it grants access to the service.

The following are important points to consider when using TCP Wrappers to protect network services:

- Because access rules in hosts.allow are applied first, they take precedence over rules specified in hosts.deny. Therefore, if access to a service is allowed in hosts.allow, a rule denying access to that same service in hosts.deny is ignored.
- The rules in each file are read from the top down and the first matching rule for a given service is the only one applied. The order of the rules is extremely important.
- If no rules for the service are found in either file, or if neither file exists, access to the service is granted.
- TCP-wrapped services do not cache the rules from the hosts access files, so any changes to hosts.allow or hosts.deny take effect immediately, without restarting network services.

Warning

If the last line of a hosts access file is not a newline character (created by pressing the Enter key), the last rule in the file fails and an error is logged to either /var/log/messages or /var/log/secure. This is also the case for a rule that spans multiple lines without using the backslash character. The following example illustrates the relevant portion of a log message for a rule failure due to either of these circumstances:

```
warning: /etc/hosts.allow, line 20: missing newline or line too long
```

46.5.2.1. Formatting Access Rules

The format for both /etc/hosts.allow and /etc/hosts.deny is identical. Each rule must be on its own line. Blank lines or lines that start with a hash (#) are ignored.

Each rule uses the following basic format to control access to network services:

```
<daemon list>: <client list> [: <option>: <option>: ...]
```

- <daemon list> — A comma-separated list of process names (not service names) or the ALL wildcard. The daemon list also accepts operators (refer to 46.5.2.1.4절. "Operators") to allow greater flexibility.
- <client list> — A comma-separated list of hostnames, host IP addresses, special patterns, or wildcards which identify the hosts affected by the rule. The client list also accepts operators listed in 46.5.2.1.4절. "Operators" to allow greater flexibility.
- <option> — An optional action or colon-separated list of actions performed when the rule is triggered. Option fields support expansions, launch shell commands, allow or deny access, and alter logging behavior.

Note

More information on the specialist terms above can be found elsewhere in this Guide:

- 46.5.2.1.1절. "Wildcards"
- 46.5.2.1.2절. "Patterns"
- 46.5.2.2.4절. "Expansions"
- 46.5.2.2절. "Option Fields"

The following is a basic sample hosts access rule:

```
vsftpd : .example.com
```

This rule instructs TCP Wrappers to watch for connections to the FTP daemon (vsftpd) from any host in the example.com domain. If this rule appears in hosts.allow, the connection is accepted. If this rule appears in hosts.deny, the connection is rejected.

The next sample hosts access rule is more complex and uses two option fields:

```
sshd : .example.com  \ : spawn /bin/echo `/bin/date` access denied>>/var/log/sshd.log \ : deny
```

Note that each option field is preceded by the backslash (\). Use of the backslash prevents failure of the rule due to length.

This sample rule states that if a connection to the SSH daemon (sshd) is attempted from a host in the example.com domain, execute the echo command to append the attempt to a special log file, and deny the connection. Because the optional deny directive is used, this line denies access even if it appears in the hosts.allow file. Refer to 46.5.2.2절. "Option Fields" for a more detailed look at available options.

46.5.2.1.1. Wildcards

Wildcards allow TCP Wrappers to more easily match groups of daemons or hosts. They are used most frequently in the client list field of access rules.

The following wildcards are available:

- ALL — Matches everything. It can be used for both the daemon list and the client list.
- LOCAL — Matches any host that does not contain a period (.), such as localhost.
- KNOWN — Matches any host where the hostname and host address are known or where the user is known.
- UNKNOWN — Matches any host where the hostname or host address are unknown or where the user is unknown.
- PARANOID — Matches any host where the hostname does not match the host address.

Caution

The KNOWN, UNKNOWN, and PARANOID wildcards should be used with care, because they rely on functioning DNS server for correct operation. Any disruption to name resolution may prevent legitimate users from gaining access to a service.

46.5.2.1.2. Patterns

Patterns can be used in the client field of access rules to more precisely specify groups of client hosts.

The following is a list of common patterns for entries in the client field:

- Hostname beginning with a period (.) — Placing a period at the beginning of a hostname matches all hosts sharing the listed components of the name. The following example applies to any host within the example.com domain:

  ```
  ALL : .example.com
  ```

- IP address ending with a period (.) — Placing a period at the end of an IP address matches all hosts sharing the initial numeric groups of an IP address. The following example applies to any host within the 192.168.x.x network:

  ```
  ALL : 192.168.
  ```

- IP address/netmask pair — Netmask expressions can also be used as a pattern to control access to a particular group of IP addresses. The following example applies to any host with an address range of 192.168.0.0 through 192.168.1.255:

  ```
  ALL : 192.168.0.0/255.255.254.0
  ```

  Important

  When working in the IPv4 address space, the address/prefix length (prefixlen) pair declarations (CIDR notation) are not supported. Only IPv6 rules can use this format.

- [IPv6 address]/prefixlen pair — [net]/prefixlen pairs can also be used as a pattern to control access to a particular group of IPv6 addresses. The following example would apply to any host with an address range of 3ffe:505:2:1:: through 3ffe:505:2:1:ffff:ffff:ffff:ffff:

  ```
  ALL : [3ffe:505:2:1::]/64
  ```

- The asterisk (*) — Asterisks can be used to match entire groups of hostnames or IP addresses, as long as they are not mixed in a client list containing other types of patterns. The following example would apply to any host within the example.com domain:

```
ALL : *.example.com
```

- The slash (/) — If a client list begins with a slash, it is treated as a file name. This is useful if rules specifying large numbers of hosts are necessary. The following example refers TCP Wrappers to the /etc/telnet.hosts file for all Telnet connections:

```
in.telnetd : /etc/telnet.hosts
```

Other, lesser used, patterns are also accepted by TCP Wrappers. Refer to the hosts_access man 5 page for more information.

Warning

Be very careful when using hostnames and domain names. Attackers can use a variety of tricks to circumvent accurate name resolution. In addition, disruption to DNS service prevents even authorized users from using network services. It is, therefore, best to use IP addresses whenever possible.

46.5.2.1.3. Portmap and TCP Wrappers

Portmap's implementation of TCP Wrappers does not support host look-ups, which means portmap can not use hostnames to identify hosts. Consequently, access control rules for portmap in hosts.allow or hosts.deny must use IP addresses, or the keyword ALL, for specifying hosts.

Changes to portmap access control rules may not take effect immediately. You may need to restart the portmap service.

Widely used services, such as NIS and NFS, depend on portmap to operate, so be aware of these limitations.

46.5.2.1.4. Operators

At present, access control rules accept one operator, EXCEPT. It can be used in both the daemon list and the client list of a rule.

The EXCEPT operator allows specific exceptions to broader matches within the same rule.

In the following example from a hosts.allow file, all example.com hosts are allowed to connect to all services except cracker.example.com:

```
ALL: .example.com EXCEPT cracker.example.com
```

In another example from a hosts.allow file, clients from the 192.168.0.x network can use all services except for FTP:

```
ALL EXCEPT vsftpd: 192.168.0.
```

Note

Organizationally, it is often easier to avoid using EXCEPT operators. This allows other administrators to quickly scan the appropriate files to see what hosts are allowed or denied access to services, without having to sort through EXCEPT operators.

46.5.2.2. Option Fields

In addition to basic rules that allow and deny access, the Red Hat Enterprise Linux implementation of TCP Wrappers supports extensions to the access control language through option fields. By using option fields in hosts access rules, administrators can accomplish a variety of tasks such as altering log behavior, consolidating access control, and launching shell commands.

46.5.2.2.1. Logging

Option fields let administrators easily change the log facility and priority level for a rule by using the severity directive.

In the following example, connections to the SSH daemon from any host in the example.com domain are logged to the default authpriv syslog facility (because no facility value is specified) with a priority of emerg:

```
sshd : .example.com : severity emerg
```

It is also possible to specify a facility using the severity option. The following example logs any SSH connection attempts by hosts from the example.com domain to the local0 facility with a priority of alert:

```
sshd : .example.com : severity local0.alert
```

Note

In practice, this example does not work until the syslog daemon (syslogd) is configured to log to the local0 facility. Refer to the syslog.conf man page for information about configuring custom log facilities.

46.5.2.2.2. Access Control

Option fields also allow administrators to explicitly allow or deny hosts in a single rule by adding the allow or deny directive as the final option.

For example, the following two rules allow SSH connections from client-1.example.com, but deny connections from client-2.example.com:

```
sshd : client-1.example.com : allow
sshd : client-2.example.com : deny
```

By allowing access control on a per-rule basis, the option field allows administrators to consolidate all access rules into a single file: either hosts.allow or hosts.deny. Some administrators consider this an easier way of organizing access rules.

46.5.2.2.3. Shell Commands

Option fields allow access rules to launch shell commands through the following two directives:

- spawn — Launches a shell command as a child process. This directive can perform tasks like using /usr/sbin/safe_finger to get more information about the requesting client or create special log files using the echo command.

  In the following example, clients attempting to access Telnet services from the example.com domain are quietly logged to a special file:

  ```
  in.telnetd : .example.com \
   : spawn /bin/echo `/bin/date` from %h>>/var/log/telnet.log \
   : allow
  ```

- twist — Replaces the requested service with the specified command. This directive is often used to set up traps for intruders (also called "honey pots"). It can also be used to send messages to connecting clients. The twist directive must occur at the end of the rule line.

  In the following example, clients attempting to access FTP services from the example.com domain are sent a message using the echo command:

  ```
  vsftpd : .example.com \
   : twist /bin/echo "421 This domain has been black-listed. Access denied!"
  ```

For more information about shell command options, refer to the hosts_options man page.

46.5.2.2.4. Expansions

Expansions, when used in conjunction with the spawn and twist directives, provide information about the client, server, and processes involved.

The following is a list of supported expansions:

- %a — Returns the client's IP address.

- %A — Returns the server's IP address.

- %c — Returns a variety of client information, such as the username and hostname, or the username and IP address.

- %d — Returns the daemon process name.

- %h — Returns the client's hostname (or IP address, if the hostname is unavailable).

- %H — Returns the server's hostname (or IP address, if the hostname is unavailable).

- %n — Returns the client's hostname. If unavailable, unknown is printed. If the client's hostname and host address do not match, paranoid is printed.

- %N — Returns the server's hostname. If unavailable, unknown is printed. If the server's hostname and host address do not match, paranoid is printed.
- %p — Returns the daemon's process ID.
- %s —Returns various types of server information, such as the daemon process and the host or IP address of the server.
- %u — Returns the client's username. If unavailable, unknown is printed.

The following sample rule uses an expansion in conjunction with the spawn command to identify the client host in a customized log file.

When connections to the SSH daemon (sshd) are attempted from a host in the example.com domain, execute the echo command to log the attempt, including the client hostname (by using the %h expansion), to a special file:

```
sshd : .example.com  \
 : spawn /bin/echo `/bin/date` access denied to %h>>/var/log/sshd.log \
 : deny
```

Similarly, expansions can be used to personalize messages back to the client. In the following example, clients attempting to access FTP services from the example.com domain are informed that they have been banned from the server:

```
vsftpd : .example.com \
: twist /bin/echo "421 %h has been banned from this server!"
```

For a full explanation of available expansions, as well as additional access control options, refer to section 5 of the man pages for hosts_access (man 5 hosts_access) and the man page for hosts_options.

Refer to 46.5.5절. "Additional Resources" for more information about TCP Wrappers.

46.5.3. xinetd

The xinetd daemon is a TCP-wrapped super service which controls access to a subset of popular network services, including FTP, IMAP, and Telnet. It also provides service-specific configuration options for access control, enhanced logging, binding, redirection, and resource utilization control.

When a client attempts to connect to a network service controlled by xinetd, the super service receives the request and checks for any TCP Wrappers access control rules.

If access is allowed, xinetd verifies that the connection is allowed under its own access rules for that service. It also checks that the service can have more resources allotted to it and that it is not in breach of any defined rules.

If all these conditions are met (that is, access is allowed to the service; the service has not reached its resource limit; and the service is not in breach of any defined rule), xinetd then starts an instance of the requested service and passes control of the connection to it. After the connection has been established, xinetd takes no further part in the communication between the client and the server.

46.5.4. xinetd Configuration Files

The configuration files for xinetd are as follows:

- /etc/xinetd.conf — The global xinetd configuration file.
- /etc/xinetd.d/ — The directory containing all service-specific files.

46.5.4.1. The /etc/xinetd.conf File

The /etc/xinetd.conf file contains general configuration settings which affect every service under xinetd's control. It is read when the xinetd service is first started, so for configuration changes to take effect, you need to restart the xinetd service. The following is a sample /etc/xinetd.conf file:

```
defaults
{
         instances               = 60
  log_type                = SYSLOG authpriv
  log_on_success          = HOST PID
  log_on_failure          = HOST
  cps                     = 25 30
}
includedir /etc/xinetd.d
```

These lines control the following aspects of xinetd:

- instances — Specifies the maximum number of simultaneous requests that xinetd can process.
- log_type — Configures xinetd to use the authpriv log facility, which writes log entries to the /var/log/secure file. Adding a directive such as FILE /var/log/xinetdlog would create a custom log file called xinetdlog in the /var/log/ directory.
- log_on_success — Configures xinetd to log successful connection attempts. By default, the remote host's IP address and the process ID of the server processing the request are recorded.
- log_on_failure — Configures xinetd to log failed connection attempts or if the connection was denied.
- cps — Configures xinetd to allow no more than 25 connections per second to any given service. If this limit is exceeded, the service is retired for 30 seconds.
- includedir /etc/xinetd.d/ — Includes options declared in the service-specific configuration files located in the /etc/xinetd.d/ directory. Refer to 46.5.4.2절. "The /etc/xinetd.d/ Directory" for more information.

Note

Often, both the log_on_success and log_on_failure settings in /etc/xinetd.conf are further modified in the service-specific configuration files. More information may therefore appear in a given service's log file than the /etc/xinetd.conf file may indicate. Refer to 46.5.4.3.1절. "Logging Options" for further information.

46.5.4.2. The /etc/xinetd.d/ Directory

The /etc/xinetd.d/ directory contains the configuration files for each service managed by xinetd and the names of the files correlate to the service. As with xinetd.conf, this directory is read only when the xinetd service is started. For any changes to take effect, the administrator must restart the xinetd service.

The format of files in the /etc/xinetd.d/ directory use the same conventions as /etc/xinetd.conf. The primary reason the configuration for each service is stored in a separate file is to make customization easier and less likely to affect other services.

To gain an understanding of how these files are structured, consider the /etc/xinetd.d/krb5-telnet file:

```
service telnet
{
         flags           = REUSE
  socket_type      = stream
  wait             = no
  user             = root
  server           = /usr/kerberos/sbin/telnetd
  log_on_failure  += USERID
  disable          = yes
}
```

These lines control various aspects of the telnet service:

- service — Specifies the service name, usually one of those listed in the /etc/services file.
- flags — Sets any of a number of attributes for the connection. REUSE instructs xinetd to reuse the socket for a Telnet connection.

  Note

  The REUSE flag is deprecated. All services now implicitly use the REUSE flag.

- socket_type — Sets the network socket type to stream.
- wait — Specifies whether the service is single-threaded (yes) or multi-threaded (no).
- user — Specifies which user ID the process runs under.
- server — Specifies which binary executable to launch.
- log_on_failure — Specifies logging parameters for log_on_failure in addition to those already defined in xinetd.conf.
- disable — Specifies whether the service is disabled (yes) or enabled (no).

Refer to the xinetd.conf man page for more information about these options and their usage.

46.5.4.3. Altering xinetd Configuration Files

A range of directives is available for services protected by xinetd. This section highlights some of the more commonly used options.

46.5.4.3.1. Logging Options

The following logging options are available for both /etc/xinetd.conf and the service-specific configuration files within the /etc/xinetd.d/ directory.

The following is a list of some of the more commonly used logging options:

- ATTEMPT — Logs the fact that a failed attempt was made (log_on_failure).
- DURATION — Logs the length of time the service is used by a remote system (log_on_success).
- EXIT — Logs the exit status or termination signal of the service (log_on_success).
- HOST — Logs the remote host's IP address (log_on_failure and log_on_success).
- PID — Logs the process ID of the server receiving the request (log_on_success).
- USERID — Logs the remote user using the method defined in RFC 1413 for all multi-threaded stream services (log_on_failure andlog_on_success).

For a complete list of logging options, refer to the xinetd.conf man page.

46.5.4.3.2. Access Control Options

Users of xinetd services can choose to use the TCP Wrappers hosts access rules, provide access control via the xinetd configuration files, or a mixture of both. Refer to 46.5.2절. "TCP Wrappers Configuration Files" for more information about TCP Wrappers hosts access control files.

This section discusses using xinetd to control access to services.

> **Note**
>
> Unlike TCP Wrappers, changes to access control only take effect if the xinetd administrator restarts the xinetd service.
>
> Also, unlike TCP Wrappers, access control through xinetd only affects services controlled by xinetd.

The xinetd hosts access control differs from the method used by TCP Wrappers. While TCP Wrappers places all of the access configuration within two files, /etc/hosts.allow and /etc/hosts.deny, xinetd's access control is found in each service's configuration file in the /etc/xinetd.d/ directory.

The following hosts access options are supported by xinetd:

- only_from — Allows only the specified hosts to use the service.
- no_access — Blocks listed hosts from using the service.
- access_times — Specifies the time range when a particular service may be used. The time range must be stated in 24-hour format notation, HH:MM-HH:MM.

The only_from and no_access options can use a list of IP addresses or host names, or can specify an entire network. Like TCP Wrappers, combining xinetd access control with the enhanced logging configuration can increase security by blocking requests from banned hosts while verbosely recording each connection attempt.

For example, the following /etc/xinetd.d/telnet file can be used to block Telnet access from a particular network group and restrict the overall time range that even allowed users can log in:

```
service telnet
{
         disable          = no
  flags              = REUSE
  socket_type        = stream
  wait               = no
  user               = root
  server             = /usr/kerberos/sbin/telnetd
  log_on_failure    += USERID
  no_access          = 172.16.45.0/24
  log_on_success    += PID HOST EXIT
  access_times       = 09:45-16:15
}
```

In this example, when a client system from the 10.0.1.0/24 network, such as 10.0.1.2, tries to access the Telnet service, it receives the following message:

```
Connection closed by foreign host.
```

In addition, their login attempts are logged in /var/log/messages as follows:

```
Sep  7 14:58:33 localhost xinetd[5285]: FAIL: telnet address from=172.16.45.107
Sep  7 14:58:33 localhost xinetd[5283]: START: telnet pid=5285 from=172.16.45.107
Sep  7 14:58:33 localhost xinetd[5283]: EXIT: telnet status=0 pid=5285 duration=0(sec)
```

When using TCP Wrappers in conjunction with xinetd access controls, it is important to understand the relationship between the two access control mechanisms.

The following is the sequence of events followed by xinetd when a client requests a connection:

1. The xinetd daemon accesses the TCP Wrappers hosts access rules using a libwrap.a library call. If a deny rule matches the client, the connection is dropped. If an allow rule matches the client, the connection is passed to xinetd.

2. The xinetd daemon checks its own access control rules both for the xinetd service and the requested service. If a deny rule matches the client, the connection is dropped. Otherwise, xinetd starts an instance of the requested service and passes control of the connection to that service.

Important

Care should be taken when using TCP Wrappers access controls in conjunction with xinetd access controls. Misconfiguration can cause undesirable effects.

46.5.4.3.3. Binding and Redirection Options

The service configuration files for xinetd support binding the service to an IP address and redirecting incoming requests for that service to another IP address, hostname, or port.

Binding is controlled with the bind option in the service-specific configuration files and links the service to one IP address on the system. When this is configured, the bind option only allows requests to the correct IP address to access the service. You can use this method to bind different services to different network interfaces based on requirements.

This is particularly useful for systems with multiple network adapters or with multiple IP addresses. On such a system, insecure services (for example, Telnet), can be configured to listen only on the interface connected to a private network and not to the interface connected to the Internet.

The redirect option accepts an IP address or hostname followed by a port number. It configures the service to redirect any requests for this service to the specified host and port number. This feature can be used to point to another port number on the same system, redirect the request to a different IP address on the same machine, shift the request to a totally different system and port number, or any combination of these options. A user connecting to a certain service on a system may therefore be rerouted to another system without disruption.

The xinetd daemon is able to accomplish this redirection by spawning a process that stays alive for the duration of the connection between the requesting client machine and the host actually providing the service, transferring data between the two systems.

The advantages of the bind and redirect options are most clearly evident when they are used together. By binding a service to a particular IP address on a system and then redirecting requests for this service to a second machine that only the first machine can see, an internal system can be used to provide services for a totally different network. Alternatively, these options can be used to limit the exposure of a particular service on a multi-homed machine to a known IP address, as well as redirect any requests for that service to another machine especially configured for that purpose.

For example, consider a system that is used as a firewall with this setting for its Telnet service:

```
service telnet
{
         socket_type   = stream
  wait    = no
  server    = /usr/kerberos/sbin/telnetd
  log_on_success   += DURATION USERID
  log_on_failure   += USERID
  bind                         = 123.123.123.123
  redirect                     = 10.0.1.13 23
}
```

The bind and redirect options in this file ensure that the Telnet service on the machine is bound to the external IP address (123.123.123.123), the one facing the Internet. In addition, any requests for Telnet service sent to 123.123.123.123 are redirected via a second network adapter to an internal IP address (10.0.1.13) that only the firewall and internal systems can access. The firewall then sends the communication between the two systems, and the connecting system thinks it is connected to 123.123.123.123 when it is actually connected to a different machine.

This feature is particularly useful for users with broadband connections and only one fixed IP address. When using Network Address Translation (NAT), the systems behind the gateway machine, which are using internal-only IP addresses, are not available from outside the gateway system. However, when certain services controlled by xinetd are configured with the bind and redirect options, the gateway machine can act as a proxy between outside systems and a particular internal machine configured to provide the service. In addition, the various xinetd access control and logging options are also available for additional protection.

46.5.4.3.4. Resource Management Options

The xinetd daemon can add a basic level of protection from Denial of Service (DoS) attacks. The following is a list of directives which can aid in limiting the effectiveness of such attacks:

- per_source — Defines the maximum number of instances for a service per source IP address. It accepts only integers as an argument and can be used in both xinetd.conf and in the service-specific configuration files in the xinetd.d/ directory.
- cps — Defines the maximum number of connections per second. This directive takes two integer arguments separated by white space. The first argument is the maximum number of connections allowed to the service per second. The second argument is the number of seconds that xinetd must wait before re-enabling the service. It accepts only integers as arguments and can be used in either the xinetd.conf file or the service-specific configuration files in the xinetd.d/ directory.
- max_load — Defines the CPU usage or load average threshold for a service. It accepts a floating point number argument.

  The load average is a rough measure of how many processes are active at a given time. See the uptime, who, and procinfo commands for more information about load average.

There are more resource management options available for xinetd. Refer to the xinetd.conf man page for more information.

46.5.5. Additional Resources

More information about TCP Wrappers and xinetd is available from system documentation and on the Internet.

46.5.5.1. Installed Documentation

The documentation on your system is a good place to start looking for additional configuration options for TCP Wrappers, xinetd, and access control.

- /usr/share/doc/tcp_wrappers-<version>/ — This directory contains a README file that discusses how TCP Wrappers work and the various hostname and host address spoofing risks that exist.
- /usr/share/doc/xinetd-<version>/ — This directory contains a README file that discusses aspects of access control and a sample.conf file with various ideas for modifying service-specific configuration files in the /etc/xinetd.d/ directory.
- TCP Wrappers and xinetd-related man pages — A number of man pages exist for the various applications and configuration files involved with TCP Wrappers and xinetd. The following are some of the more important man pages:

  Server Applications
  - man xinetd — The man page for xinetd.

  Configuration Files
  - man 5 hosts_access — The man page for the TCP Wrappers hosts access control files.
  - man hosts_options — The man page for the TCP Wrappers options fields.
  - man xinetd.conf — The man page listing xinetd configuration options.

46.5.5.2. Useful Websites

- http://www.xinetd.org/[4] — The home of xinetd, containing sample configuration files, a full listing of features, and an informative FAQ.
- http://www.macsecurity.org/resources/xinetd/tutorial.shtml — A thorough tutorial that discusses many different ways to optimize default xinetd configuration files to meet specific security goals.

46.5.5.3. Related Books

- Hacking Linux Exposed by Brian Hatch, James Lee, and George Kurtz; Osbourne/McGraw-Hill — An excellent security resource with information about TCP Wrappers and xinetd.

46.6. Kerberos

네트워크에서 시스템의 보안을 유지하는 것은 힘이 듭니다. 시스템 관리자는 네트워크에서 어떤 서비스가 실행되고 있는 지와 이러한 서비스가 사용하는 방식을 지속적으로 추적하기 위해 시간을 소요할 수 있습니다.

Further, authenticating users to network services can prove dangerous when the method used by the protocol is inherently insecure, as evidenced by the transfer of unencrypted passwords over a network using the traditional FTP and Telnet protocols.

커베로스는 안전하지 않은 인증 방식을 허용하여 모든 네트워크 보안을 증강시키는데 필요한 프로토콜을 삭제하기 위한 방법입니다.

46.6.1. 커베로스란?

커베로스는 미국 MIT에서 개발한 네트워크 인증 프로토콜로서 대칭키 암호화 (symmetric key cryptography)[5]를 사용하여 네트워크 상에서 암호를 보낼 필요가 없이 네트워크 서비스 사용자를 인증하는 방법입니다.

커베로스를 사용하여 네트워크 서비스 사용자를 인증함으로서, 다른 권한이 없는 사용자가 네트워크 소통량을 모니터하여 암호에 대한 정보가 누설되는 위험 부담이 줄어듭니다.

46.6.1.1. 커베로스의 장점

대부분의 전통적인 네트워크 시스템은 암호 기반 인증 방식을 사용합니다. 이러한 방식은 사용자가 네트워크 서버에 접속하기 위하여 사용자명과 암호를 입력해야 합니다. 그러나 불행히도 여러 서비스에서 이러한 인증 정보가 암호화되지 않은 채 전송됩니다. 따라서 이러한 인증 방식의 보안을 증강시키기 위해서는 외부 사용자의 네트워크 접속을 방지하고 신용할 수 있는 컴퓨터와 사용자만 네트워크에 접속할 수 있도록 해야 합니다.

외부 사용자가 차단되고 네트워크 상 모든 컴퓨터와 사용자를 신용할 수 있다고 해도, 일단 네트워크가 인터넷에 연결된 후에는 더 이상 네트워크가 안전하다고 할 수 없습니다. 인터넷을 통하여 네트워크에 접속한 해커가 패킷 스니퍼 (packet sniffer)라고도 알려진 단순 패킷 분석기 (packet analyzer)를 사용하여 암호화되지 않은 채 전송되는 사용자명과 암호를 중간에서 가로챈 후 사용자 계정을 통하여 전체 보안 시스템을 위협할 가능성이 있습니다.

[4] http://www.xinetd.org
[5] 클라이언트와 서버가 모두 네트워크 통신을 암호화하고 암호 해독하는데 공유키를 사용하는 시스템

커베로스의 주된 목적은 네트워크 내에서 암호화되지 않은 암호가 전송되는 것을 막는 것입니다. 커베로스가 적절히 사용된다면, 네트워크 상에 패킷 스니퍼 공격을 효율적으로 방지할 수 있습니다.

46.6.1.2. 커베로스의 단점

비록 커베로스를 사용하여 가장 흔하고 심각한 보안 위협을 제거할 수는 있지만, 다음과 같은 다양한 이유 때문에 커베로스를 구현하는 것이 쉽지 않습니다:

- /etc/passwd 또는 /etc/shadow와 같은 표준 UNIX 암호 데이터베이스에서 사용자 암호를 커베로스 암호 데이터베이스로 옮기는 작업은 아직 자동화되지 않아서 매우 느리고 지루한 작업이 될 수 있습니다. 보다 많은 정보를 원하신다면, 다음 온라인 커베로스 FAQ에서 질문 2.23 번을 참조하시기 바랍니다:

  http://www.nrl.navy.mil/CCS/people/kenh/kerberos-faq.html[6]

- Kerberos has only partial compatibility with the Pluggable Authentication Modules (PAM) system used by most Red Hat Enterprise Linux servers. Refer to 46.6.4절. "커베로스와 PAM" for more information about this issue.

- 커베로스는 신뢰할 수 있는 개별 사용자가 신뢰할 수 없는 네트워크 상에서 신뢰할 수 없는 호스트를 사용하고 있다고 간주합니다. 커베로스의 주된 목적은 네트워크 상에서 암호가 전송되는 것을 방지하는 것입니다. 그러나 만일 적절한 권한이 없는 사용자가 사용자 인증에 사용되는 티켓을 배포하는 호스트 — 키 배포 센터 (KDC) —에 접근하게 된다면, 전체 커베로스 인증 시스템의 보안이 위험해질 수 있습니다.

- 응용 프로그램이 커베로스를 사용하도록 설정하기 위해서는, 커베로스 라이브러리를 적절히 호출할 수 있도록 해당 응용 프로그램의 소스를 수정해야 합니다. 이렇게 수정된 응용 프로그램을 커베로스-인식(Kerberos-aware), 또는 커베로스화(kerberized)되었다고 간주합니다. 그러나 일부 응용 프로그램의 경우 그 응용 프로그램의 크기나 디자인 때문에 이러한 설정이 매우 힘듭니다. 다른 호환되지 않는 응용 프로그램의 경우, 서버와 클라이언트 측면에서 서로 통신할 수 있는 방식으로 소스를 수정해야 하며, 이러한 수정은 방대한 양의 프로그래밍을 요구합니다. 종종 커베로스를 기본적으로 지원하지 않는 소스를 공개하지 않는 (closed-source) 응용 프로그램이 가장 큰 문제가 됩니다.

- 커베로스는 모든 응용 프로그램에서 사용되지 않고 부분적으로 사용되면 아무런 역할을 하지 않습니다. 네트워크 상에서 커베로스를 사용하는 경우 암호를 암호화되지 않은 상태에서 커베로스를 사용하지 않는 서비스로 전송하게되면 암호가 누출될 위험이 있습니다. 따라서 여러분의 네트워크가 커베로스를 사용한다고 해도 인증에 아무런 효과를 볼 수 없게 됩니다. 커베로스를 사용하여 네트워크 보안을 강화하기 위해서는 암호화되지 않은 암호를 보내는 모든 클라이언트/서버 응용 프로그램은 커베로스를 사용해야 합니다. 그렇지 않으면 차라리 아무런 클라이언트/서버 응용 프로그램도 사용하지 않는 것이 낫습니다.

46.6.2. 키베로스 용어

커베로스는 다양한 서비스를 정의하기 위하여 독자적인 용어를 사용합니다. 커베로스가 어떻게 작용하는지 배우기 전에, 다음과 같은 용어를 알아두시는 것이 좋습니다.

[6] http://www.nrl.navy.mil/CCS/people/kenh/kerberos-faq.html#pwconvert

인증 서버 (AS)

사용자가 서비스를 요청시 티켓을 발행한 후 사용자가 서비스를 사용할 수 있도록 돌려주는 서버. AS는 서비스를 요청시 증명(credentials)이 없는 또는 증명을 보내지 않는 클라이언트의 요청에 응답합니다. AS는 일반적으로 티켓 부여 티켓 (TGT)를 발행하여 티켓 부여 서버 (TGS)에 접속할때 사용됩니다. AS는 보통 키 배포 센터 (KDC)와 동일한 호스트 상에서 실행됩니다.

암호문 (ciphertext)

암호화된 데이터.

클라이언트

커베로스에서 티켓을 부여받을 수 있는 네트워크 상의 개체 (사용자, 호스트 또는 응용 프로그램).

증명 (credentials)

서버에 특정 서비스를 요구하는 클라이언트를 제한된 시간 동안 인증하는 전자 증명으로서 티켓으로도 불리웁니다.

증명 캐시 혹은 티켓 파일

사용자와 다양한 네트워크 서비스 간에 주고 받는 통신을 암호화하는데 사용되는 키들을 포함한 파일. 커베로스 5는 다른 캐시 유형에 사용되는 공유 메모리와 같은 프레임워크를 지원하지만, 파일은 보다 완전하게 지원됩니다.

암호 해시

사용자 인증을 위해 사용되는 단방향 해시 (one way hash). 암호화되지 않은 데이터를 사용하는 것보다 안전하지만, 여전히 숙련된 해커라면 쉽게 암호를 해독할 수 있습니다.

GSS-API

IETF(Internet Engineering Task Force)에서 발표한 RFC-2743에 정의된 GSS-API (Generic Security Service Application Program Interface)는 보안 서비스를 제공하는 함수 집합입니다. 이 API는 클라이언트와 서비스 간에 기본적인 메커니즘에 대한 특별한 지식 없이도 각 프로그램이 서로 인증할 수 있게 해줍니다. 만일 cyrus-IMAP와 같은 네트워크 서비스가 GSS-API를 사용한다면, 커베로스를 사용한 인증을 수행 가능합니다.

해시

해시 값(hash value)이라고도 알려짐. 해시 기능(hash function)을 통해 문자열의 전달에 의해 생성된 값. 이러한 값은 일반적으로 전송된 데이터가 변경되었는 지를 확인하는 데 사용됩니다.

해시 기능

A way of generating a digital "fingerprint" from input data. These functions rearrange, transpose or otherwise alter data to produce a hash value.

key

다른 데이터를 암호화하고 암호 해독하는데 사용되는 데이터. 암호화된 데이터는 해커가 아무리 추측하더라도 적절한 키가 없이는 해독하기 힘듭니다.

키 배포 센터 (KDC)

일반적으로 티켓 부여 서버 (TGS)와 같은 호스트 상에서 실행되는 커베로스 티켓 발행 서비스.

키탭(keytab) (또는 키 테이블)

주 멤버 (principals)와 그들의 키가 담긴 암호화되지 않은 목록을 포함한 파일. 서버는 필요한 키를 kinit을 사용하는 대신 keytab 파일에서 가져옵니다. 기본 keytab 파일은 /etc/krb5.keytab

입니다. KDC 관리 서버인 /usr/kerberos/sbin/kadmind만 유일하게 다른 파일인 /var/kerberos/krb5kdc/kadm5.keytab 파일을 사용합니다.

kinit

kinit 명령은 이미 로그인한 주 멤버가 초기 티켓 부여 티켓 (TGT)을 받아서 캐시할 수 있도록 해줍니다. kinit 명령 사용에 대한 보다 많은 정보를 원하신다면 메뉴얼 페이지를 참조하시기 바랍니다.

주 멤버 (또는 주 멤버명)

The principal is the unique name of a user or service allowed to authenticate using Kerberos. A principal follows the form root[/instance]@REALM. For a typical user, the root is the same as their login ID. The instance is optional. If the principal has an instance, it is separated from the root with a forward slash ("/"). An empty string ("") is considered a valid instance (which differs from the default NULL instance), but using it can be confusing. All principals in a realm have their own key, which for users is derived from a password or is randomly set for services.

커베로스 영역 (realm)

KDC라고 부르는 한 개 이상의 서버와 다수의 클라이언트로 이루어진 커베로스를 사용하는 네트워크.

서비스

네트워크 상에서 사용되는 프로그램.

티켓

서버에 특정 서비스를 요구하는 클라이언트를 제한된 시간 동안 인증하는 전자 증명 티켓.

티켓 부여 서비스 (TGS)

사용자가 서비스를 요청할 때 서비스에 보낼 수 있는 티켓을 돌려주는 서버. TGS는 일반적으로 KDC와 동일한 호스트 상에서 서비스를 수행합니다.

티켓 부여 티켓 (TGT)

클라이언트가 KDC에서 티켓을 요청할 필요가 없이 추가 티켓을 받을 수 있도록 해주는 특별한 티켓.

암호화되지 않은 암호

평문으로서 판독 가능한 암호.

46.6.3. 커베로스 작업 방식

Kerberos differs from username/password authentication methods. Instead of authenticating each user to each network service, Kerberos uses symmetric encryption and a trusted third party (a KDC), to authenticate users to a suite of network services. When a user authenticates to the KDC, the KDC sends a ticket specific to that session back to the user's machine, and any Kerberos-aware services look for the ticket on the user's machine rather than requiring the user to authenticate using a password.

커베로스를 사용하는 네트워크 상에서 사용자가 자신의 워크스테이션에 로그인하면, 그 사용자의 주 멤버 정보가 인증 서버(AS)로 부터 티켓 부여 티켓 (TGT)을 받기 위하여 KDC로 전달됩니다. 이러한 요청은 사용자가 알 수 있도록 로그인 프로그램을 통해 전달될 수 도 있고 또는 사용자가 로그인한 후 kinit 프로그램에 의해 전달될 수 도 있습니다.

The KDC then checks for the principal in its database. If the principal is found, the KDC creates a TGT, which is encrypted using the user's key and returned to that user.

The login or kinit program on the client then decrypts the TGT using the user's key, which it computes from the user's password. The user's key is used only on the client machine and is not transmitted over the network.

The TGT is set to expire after a certain period of time (usually ten to twenty-four hours) and is stored in the client machine's credentials cache. An expiration time is set so that a compromised TGT is of use to an attacker for only a short period of time. After the TGT has been issued, the user does not have to re-enter their password until the TGT expires or until they log out and log in again.

사용자가 네트워크 서비스를 필요로 할 때마다, 클라이언트 소프트웨어는 TGT를 사용하여 특정 서비스에 대한 새로운 티켓을 발급받도록 TGS에게 요청합니다. 그 후 서비스 티켓은 해당 서비스로 사용자를 인증하는데 사용됩니다.

경고

네트워크 상에서 사용자가 평문으로된 암호를 커베로스를 사용하지 않는 서비스에 보내어 인증한다면 커베로스 시스템은 언제든지 깨어질 위험이 있습니다. 따라서 커베로스를 사용하지 않는 서비스는 사용하지 않는 것이 좋습니다. 이러한 서비스에는 Telnet과 FTP가 있습니다. SSH 또는 SSL 보안 서비스와 같은 암호화된 프로토콜을 사용하는 것은 가능하지만, 그리 완벽하지는 않습니다.

This is only a broad overview of how Kerberos authentication works. Refer to 46.6.10절. “추가 자료” for links to more in-depth information.

주목

커베로스는 다음의 네트워크 서비스에 따라 올바르게 작동합니다.

- 네트워크 상의 컴퓨터 간의 대략적인 시간을 동기화합니다.

  A clock synchronization program should be set up for the network, such as ntpd. Refer to /usr/share/doc/ntp-<version-number>/index.html for details on setting up Network Time Protocol servers (where <version-number> is the version number of the ntp package installed on your system).

- DNS (Domain Name Service).

  You should ensure that the DNS entries and hosts on the network are all properly configured. Refer to the Kerberos V5 System Administrator's Guide in /usr/share/doc/krb5-server-<version-number> for more information (where <version-number> is the version number of the krb5-server package installed on your system).

46.6.4. 커베로스와 PAM

현재 커베로스 서비스는 PAM (Pluggable Authentication Modules)을 사용하지 않습니다 — 커베로스를 사용하는 서버는 PAM을 완전히 무시합니다. 그러나 PAM을 사용하는 응용 프로그램은 (pam_krb5 패키지에 포함된) pam_krb5 모듈이 설치된 경우 커베로스를 사용하여 인증 작업을 수행할 수 있습니다. pam_krb5 패키지에는 사용자를 인증하고 사용자의 암호를 사용하여 초기 인증 정보를 가져오는 login과 gdm과 같은 서비스를 가능하게 해주는 샘플 설정 파일이 포함되어 있습

니다. 만일 커베로스를 사용하는 서비스나 IMAP과 같은 GSS-API를 사용하는 서비스만 사용하여 네트워크 서버를 사용한다면, 네트워크가 상당히 안전하다고 간주됩니다.

Tip

관리자는 주의해서 사용자가 커베로스 암호를 사용하여 대부분의 네트워크 서비스로 자신을 인증하는 것을 막아야 합니다. 이러한 서비스는 네트워크로 암호를 암호화하지 않은채 전송하는 많은 프로토콜을 사용하기 때문에, 커베로스 시스템을 사용하는 목적이 상실됩니다. 예를 들면, 사용자가 telnet 상에서 커베로스 암호를 사용하여 인증하는 것을 허용하시면 안됩니다.

46.6.5. 커베로스 5 서버 설정하기

When setting up Kerberos, install the KDC first. If it is necessary to set up slave servers, install the master first.

To configure the first Kerberos KDC, follow these steps:

1. 커베로스를 설정하시기 전에 네트워크 상 컴퓨터 간에 시간이 동기화되었는지와 서버 상 DNS가 작동하고 있는지를 확인해 주시기 바랍니다. 커베로스 서버와 다양한 클라이언트 간에 시간이 동기화 되었는지에 특히 주의하여 살펴보시기 바랍니다. 만일 서버와 클라이언트의 시간이 5분 (이 기본 값은 커베로스 5에서 설정 가능합니다) 이상 차이가 난다면, 커베로스 클라이언트는 서버에 인증할 수 없게 됩니다. 이 시간 동기화 작업은 해커가 이전 커베로스 티켓을 사용하여 유효한 사용자로 위장하는 것을 방지하는데 중요한 역할을 합니다.

   It is advisable to set up a Network Time Protocol (NTP) compatible client/server network even if Kerberos is not being used. Red Hat Enterprise Linux includes the ntp package for this purpose. Refer to /usr/share/doc/ntp-<version-number>/index.html (where <version-number> is the version number of the ntp package installed on your system) for details about how to set up Network Time Protocol servers, and http://www.ntp.org for more information about NTP.

2. KDC를 실행할 기계에 krb5-libs, krb5-server, ="none">krb5-workstation 패키지를 설치하시기 바랍니다. 이 기계는 매우 안전해야 합니다 — 만일 가능하다면, KDC가 아닌 다른 서비스는 실행하지 않는 것이 좋습니다.

3. 커베로스 영역명과 도메인과 영역명을 묶는 매핑에 대한 정보를 담도록 /etc/krb5.conf 설정 파일과 /var/kerberos/krb5kdc/kdc.conf 설정 파일을 편집하십시오. EXAMPLE.COM과 example.com — 반드시 대문자와 소문자를 올바른 형식으로 유지하셔야 합니다 — 을 도메인명으로 대체하고 KDC를 kerberos.example.com에서 케베로스 서버 이름으로 변경하여 쉽게 단순 영역명을 생성하실 수 있습니다. 일반적으로 모든 영역명은 대문자이며 모든 DNS 호스트명과 도메인 이름은 소문자입니다. 이러한 파일 형식에 대한 보다 자세한 정보는 각 메뉴얼 페이지를 참조하시기 바랍니다.

4. 쉘 프롬프트에서 kdb5_util 유틸리티를 사용하여 데이터베이스를 생성하십시오:

   ```
   /usr/kerberos/sbin/kdb5_util create -s
   ```

   create 명령은 여러분의 커베로스 영역에 사용되는 키를 저장할 데이터베이스를 생성합니다. -s 스위치 옵션은 마스터 서버 키 파일이 저장될 숨은 (stash) 파일을 생성하게 합니다. 키를 읽어올 숨은 파일이 존재하지 않는다면, 커베로스 서버 (krb5kdc)는 매번 시작할 때마다 사용자에게 (키를 재생성하는데 사용되는) 마스터 서버 암호를 입력하도록 요구할 것입니다.

5. /var/kerberos/krb5kdc/kadm5.acl 파일을 수정하시기 바랍니다. kadmind 명령은 이 파일을 사용하여 어느 주 멤버가 커베로스 데이터베이스에 관리 권한을 가지고 있는지와 그 권한 수준을 알아냅니다. 대부분의 경우 다음 한 줄만 입력하시면 됩니다:

   ```
   */admin@EXAMPLE.COM   *
   ```

   Most users are represented in the database by a single principal (with a NULL, or empty, instance, such as joe@EXAMPLE.COM). In this configuration, users with a second principal with an instance of admin (for example, joe/admin@EXAMPLE.COM) are able to wield full power over the realm's Kerberos database.

   일단 kadmind가 서버에서 시작되면, 어느 사용자든 커베로스 영역 내 어느 클라이언트나 서버상에서 kadmin을 실행하여 서비스를 사용할 수 있습니다. 그러나 kadm5.acl 파일 목록에 포함된 사용자만이 자신의 암호를 변경하는 작업을 제외하고는 어떠한 방식으로든 데이터베이스를 수정할 수 있습니다.

   주목

   kadmin 유틸리티는 네트워크 상에서 kadmind 서버와 통신을 주고 받으며, 커베로스를 사용하여 인증을 처리합니다. 따라서 네트워크 상에서 서버를 관리하기 위하여 서버에 연결하기 전에 첫번째 주 멤버를 생성하셔야 합니다. kadmin.local 명령을 사용하여 첫번째 주 멤버를 생성하시기 바랍니다. 이 명령은 KDC가 실행되는 동일한 호스트에서만 사용되도록 고안되었으며 인증을 위해 커베로스를 사용하지 않습니다.

   KDC 터미널에서 다음 kadmin.local 명령을 입력하여 첫번째 주 멤버를 생성하시기 바랍니다:

   ```
   /usr/kerberos/sbin/kadmin.local -q "addprinc username/admin"
   ```

6. 다음 명령을 사용하여 커베로스를 시작하십시오:

   ```
   service krb5kdc start
   service kadmin start
   service krb524 start
   ```

7. kadmin 명령과 함께 addprinc 명령을 사용하여 사용자를 위한 주 멤버를 추가하시기 바랍니다. kadmin과 kadmin.local는 KDC에 대한 명령행 인터페이스입니다. 따라서 kadmin 프로그램을 시작 후 많은 명령어 — 예: addprinc —를 사용 가능합니다. 보다 많은 정보를 원하신다면, kadmin 메뉴얼 페이지를 참조하시기 바랍니다.

8. KDC가 티켓을 발행하는지 확인해주십시오. 우선 kinit 명령을 실행하여 티켓을 받아 증명 캐시 파일에 보관합니다. 다음으로 캐시 파일에서 인증 목록을 살펴 보시려면 klist 명령을 사용하시고, 캐시와 그 캐시가 포함한 인증 정보를 삭제하시려면 kdestroy 명령을 사용하시기 바랍니다.

주목

By default, kinit attempts to authenticate using the same system login username (not the Kerberos server). If that username does not correspond to a principal in the Kerberos database, kinit issues an error message. If that happens, supply kinit with the name of the correct principal as an argument on the command line (kinit <principal>).

앞서 설명된 단계를 마치셨다면, 커베로스 서버가 작동 시작할 것입니다.

46.6.6. 커베로스 5 클라이언트 설정하기

Setting up a Kerberos 5 client is less involved than setting up a server. At a minimum, install the client packages and provide each client with a valid krb5.conf configuration file. While ssh and slogin are the preferred method of remotely logging in to client systems, Kerberized versions of rsh and rlogin are still available, though deploying them requires that a few more configuration changes be made.

1. Be sure that time synchronization is in place between the Kerberos client and the KDC. Refer to 46.6.5절. "커베로스 5 서버 설정하기" for more information. In addition, verify that DNS is working properly on the Kerberos client before configuring the Kerberos client programs.

2. 모든 클라이언트 기계에서 krb5-libs와krb5-workstation 패키지를 설치해 주십시오. 각 클라이언트마다 유효한 /etc/krb5.conf 파일을 입력해 주셔야 합니다 (일반적으로 이 파일은 KDC에 의해 사용되는 krb5.conf 파일과 동일합니다).

3. Before a workstation in the realm can use Kerberos to authenticate users who connect using ssh or Kerberized rsh or rlogin, it must have its own host principal in the Kerberos database. The sshd, kshd, and klogind server programs all need access to the keys for the host service's principal. Additionally, in order to use the kerberized rsh and rlogin services, that workstation must have the xinetd package installed.

   Using kadmin, add a host principal for the workstation on the KDC. The instance in this case is the hostname of the workstation. Use the -randkey option for the kadmin's addprinc command to create the principal and assign it a random key:

   ```
   addprinc -randkey host/blah.example.com
   ```

   이제 주 멤버를 생성을 마치셨으니, 해당 워크스테이션에서 kadmin 명령에 ktadd 명령을 함께 사용하여 워크스테이션에 사용될 키를 가져올 수 있습니다:

   ```
   ktadd -k /etc/krb5.keytab host/blah.example.com
   ```

4. 다른 커베로스 네트워크 서비스를 사용하시려면, 그 서비스들을 시작시켜야 합니다. 다음은 자주 사용되는 커베로스 서비스 목록과 이 서비스들을 활성화하는데 필요한 지시 사항입니다:

   - ssh — OpenSSH uses GSS-API to authenticate users to servers if the client's and server's configuration both have GSSAPIAuthentication enabled. If the client also has

GSSAPIDelegateCredentials enabled, the user's credentials are made available on the remote system.

- rsh와 rlogin — 커베로스를 사용하는 버전의 rsh와 rlogin을 사용하기 위해서는, klogin, eklogin, kshell을 활성화하셔야 합니다.

- Telnet — 커베로스를 사용한 telnet을 사용하시려면, krb5-telnet을 활성화하십시오.

- FTP — FTP를 사용할 수 있게 하시려면, ftp의 루트로서 주 멤버의 키를 생성하신 후 가져오셔야 합니다. FTP 서버의 완전한 호스트명 (fully qualified hostname)에 인스턴스를 잊지말고 설정하신 후 gssftp를 활성화하시기 바랍니다.

- IMAP — 커베로화된 IMAP 서버를 사용하기 위해, cyrus-sasl-gssapi 패키지기 설치되어 있을 경우 cyrus-imap 패키지는 Kerberos 5를 사용합니다. cyrus-sasl-gssapi 패키지에는 GSS-API 인증을 지원하는 Cyrus SASL이 포함되어 있습니다. Cyrus IMAP는 cyrus 사용자가 /etc/krb5.keytab에서 올바른 키를 찾을 수 있고, 주 멤버의 루트(root)가 (kadmin와 함께 생성된) imap에 설정되어 있는 한 커베로스와 함께 올바르게 작동해야 합니다.

  An alternative to cyrus-imap can be found in the dovecot package, which is also included in Red Hat Enterprise Linux. This package contains an IMAP server but does not, to date, support GSS-API and Kerberos.

- CVS — CVS에서 커베로스를 사용하는 gserver는 cvs를 루트로 가진 주 멤버를 사용합니다. 이 점만 제외하고는 CVS pserver와 동일합니다.

Refer to 17장. 서비스로의 접근 통제 for details about how to enable services.

46.6.7. Domain-to-Realm Mapping

When a client attempts to access a service running on a particular server, it knows the name of the service (host) and the name of the server (foo.example.com), but because more than one realm may be deployed on your network, it must guess at the name of the realm in which the service resides.

By default, the name of the realm is taken to be the DNS domain name of the server, upper-cased.

```
foo.example.org → EXAMPLE.ORG
  foo.example.com → EXAMPLE.COM
  foo.hq.example.com → HQ.EXAMPLE.COM
```

In some configurations, this will be sufficient, but in others, the realm name which is derived will be the name of a non-existent realm. In these cases, the mapping from the server's DNS domain name to the name of its realm must be specified in the domain_realm section of the client system's krb5.conf. For example:

```
[domain_realm]
.example.com = EXAMPLE.COM
example.com = EXAMPLE.COM
```

The above configuration specifies two mappings. The first mapping specifies that any system in the "example.com" DNS domain belongs to the EXAMPLE.COM realm. The second specifies that a system with the exact name "example.com" is also in the realm. (The distinction between a domain and a specific host is marked by the presence or lack of an initial ".".) The mapping can also be stored directly in DNS.

46.6.8. Setting Up Secondary KDCs

For a number of reasons, you may choose to run multiple KDCs for a given realm. In this scenario, one KDC (the master KDC) keeps a writable copy of the realm database and runs kadmind (it is also your realm's admin server), and one or more KDCs (slave KDCs) keep read-only copies of the database and run kpropd.

The master-slave propagation procedure entails the master KDC dumping its database to a temporary dump file and then transmitting that file to each of its slaves, which then overwrite their previously-received read-only copies of the database with the contents of the dump file.

To set up a slave KDC, first ensure that the master KDC's krb5.conf and kdc.conf files are copied to the slave KDC.

Start kadmin.local from a root shell on the master KDC and use its add_principal command to create a new entry for the master KDC's host service, and then use its ktadd command to simultaneously set a random key for the service and store the random key in the master's default keytab file. This key will be used by the kprop command to authenticate to the slave servers. You will only need to do this once, regardless of how many slave servers you install.

```
~]# kadmin.local -r EXAMPLE.COM
Authenticating as principal root/admin@EXAMPLE.COM with password.
kadmin: add_principal -randkey host/masterkdc.example.com
Principal "host/host/masterkdc.example.com@EXAMPLE.COM" created.
kadmin: ktadd host/masterkdc.example.com
Entry for principal host/masterkdc.example.com with kvno 3, encryption type Triple DES cbc mode with \
 HMAC/sha1 added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/masterkdc.example.com with kvno 3, encryption type ArcFour with HMAC/md5 \
 added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/masterkdc.example.com with kvno 3, encryption type DES with HMAC/sha1 added \
 to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/masterkdc.example.com with kvno 3, encryption type DES cbc mode with RSA-MD5 \
 added to keytab WRFILE:/etc/krb5.keytab.
kadmin: quit
```

Start kadmin from a root shell on the slave KDC and use its add_principal command to create a new entry for the slave KDC's host service, and then use kadmin's ktadd command to simultaneously set a random key for the service and store the random key in the slave's default keytab file. This key is used by the kpropd service when authenticating clients.

```
~]# kadmin -p jimbo/admin@EXAMPLE.COM -r EXAMPLE.COM
Authenticating as principal jimbo/admin@EXAMPLE.COM with password.
Password for jimbo/admin@EXAMPLE.COM:
kadmin: add_principal -randkey host/slavekdc.example.com
Principal "host/slavekdc.example.com@EXAMPLE.COM" created.
kadmin: ktadd host/slavekdc.example.com@EXAMPLE.COM
Entry for principal host/slavekdc.example.com with kvno 3, encryption type Triple DES cbc mode with \
 HMAC/sha1 added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/slavekdc.example.com with kvno 3, encryption type ArcFour with HMAC/md5 added \
 to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/slavekdc.example.com with kvno 3, encryption type DES with HMAC/sha1 added \
 to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/slavekdc.example.com with kvno 3, encryption type DES cbc mode with RSA-MD5 added \
 to keytab WRFILE:/etc/krb5.keytab.
kadmin: quit
```

With its service key, the slave KDC could authenticate any client which would connect to it. Obviously, not all of them should be allowed to provide the slave's kprop service with a new realm database. To restrict access, the kprop service on the slave KDC will only accept updates from clients whose principal names are listed in /var/kerberos/krb5kdc/kpropd.acl. Add the master KDC's host service's name to that file.

```
~]# echo host/masterkdc.example.com@EXAMPLE.COM > /var/kerberos/krb5kdc/kpropd.acl
```

Once the slave KDC has obtained a copy of the database, it will also need the master key which was used to encrypt it. If your KDC database's master key is stored in a stash file on the master KDC (typically named /var/kerberos/krb5kdc/.k5.REALM, either copy it to the slave KDC using any available secure method, or create a dummy database and identical stash file on the slave KDC by running kdb5_util create -s (the dummy database will be overwritten by the first successful database propagation) and supplying the same password.

Ensure that the slave KDC's firewall allows the master KDC to contact it using TCP on port 754 (krb5_prop), and start the kprop service. Then, double-check that the kadmin service is disabled.

Now perform a manual database propagation test by dumping the realm database, on the master KDC, to the default data file which the kprop command will read (/var/kerberos/krb5kdc/slave_datatrans), and then use the kprop command to transmit its contents to the slave KDC.

```
~]# /usr/kerberos/sbin/kdb5_util dump /var/kerberos/krb5kdc/slave_datatrans
~]# kprop slavekdc.example.com
```

Using kinit, verify that a client system whose krb5.conf lists only the slave KDC in its list of KDCs for your realm is now correctly able to obtain initial credentials from the slave KDC.

That done, simply create a script which dumps the realm database and runs the kprop command to transmit the database to each slave KDC in turn, and configure the cron service to run the script periodically.

46.6.9. Setting Up Cross Realm Authentication

Cross-realm authentication is the term which is used to describe situations in which clients (typically users) of one realm use Kerberos to authenticate to services (typically server processes running on a particular server system) which belong to a realm other than their own.

For the simplest case, in order for a client of a realm named A.EXAMPLE.COM to access a service in the B.EXAMPLE.COM realm, both realms must share a key for a principal named krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM, and both keys must have the same key version number associated with them.

To accomplish this, select a very strong password or passphrase, and create an entry for the principal in both realms using kadmin.

```
~]# kadmin -r A.EXAMPLE.COM
kadmin: add_principal krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM
Enter password for principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM":
Re-enter password for principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM":
Principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM" created.
kadmin: quit
~]# kadmin -r B.EXAMPLE.COM
kadmin: add_principal krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM
Enter password for principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM":
Re-enter password for principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM":
Principal "krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM" created.
kadmin: quit
```

Use the get_principal command to verify that both entries have matching key version numbers (kvno values) and encryption types.

Dumping the Database Doesn't Do It

Security-conscious administrators may attempt to use the add_principal command's -randkey option to assign a random key instead of a password, dump the new entry from the database of the first realm, and import it into the second. This will not work unless the master keys for the realm databases are identical, as the keys contained in a database dump are themselves encrypted using the master key.

Clients in the A.EXAMPLE.COM realm are now able to authenticate to services in the B.EXAMPLE.COM realm. Put another way, the B.EXAMPLE.COM realm now trusts the A.EXAMPLE.COM realm, or phrased even more simply, B.EXAMPLE.COM now trusts A.EXAMPLE.COM.

This brings us to an important point: cross-realm trust is unidirectional by default. The KDC for the B.EXAMPLE.COM realm may trust clients from the A.EXAMPLE.COM to authenticate to services in the B.EXAMPLE.COM realm, but the fact that it does has no effect on whether or not clients in the B.EXAMPLE.COM realm are trusted to authenticate to services in the A.EXAMPLE.COM realm. To establish trust in the other direction, both realms would need to share keys for the krbtgt/A.EXAMPLE.COM@B.EXAMPLE.COM service (take note of the reversed in order of the two realms compared to the example above).

If direct trust relationships were the only method for providing trust between realms, networks which contain multiple realms would be very difficult to set up. Luckily, cross-realm trust is transitive. If clients from A.EXAMPLE.COM can authenticate to services in B.EXAMPLE.COM, and clients from B.EXAMPLE.COM can authenticate to services in C.EXAMPLE.COM, then clients in A.EXAMPLE.COM can also authenticate to services in C.EXAMPLE.COM, even if C.EXAMPLE.COM doesn't directly trust A.EXAMPLE.COM. This means that, on a network with multiple realms which all need to trust each other, making good choices about which trust relationships to set up can greatly reduce the amount of effort required.

Now you face the more conventional problems: the client's system must be configured so that it can properly deduce the realm to which a particular service belongs, and it must be able to determine how to obtain credentials for services in that realm.

First things first: the principal name for a service provided from a specific server system in a given realm typically looks like this:

```
service/server.example.com@EXAMPLE.COM
```

In this example, service is typically either the name of the protocol in use (other common values include ldap, imap, cvs, and HTTP) or host, server.example.com is the fully-qualified domain name of the system which runs the service, and EXAMPLE.COM is the name of the realm.

To deduce the realm to which the service belongs, clients will most often consult DNS or the domain_realm section of /etc/krb5.conf to map either a hostname (server.example.com) or a DNS domain name (.example.com) to the name of a realm (EXAMPLE.COM).

Having determined which to which realm a service belongs, a client then has to determine the set of realms which it needs to contact, and in which order it must contact them, to obtain credentials for use in authenticating to the service.

This can be done in one of two ways.

The default method, which requires no explicit configuration, is to give the realms names within a shared hierarchy. For an example, assume realms named A.EXAMPLE.COM, B.EXAMPLE.COM, and EXAMPLE.COM. When a client in the A.EXAMPLE.COM realm attempts to authenticate to a service in B.EXAMPLE.COM, it will, by default, first attempt to get credentials for the EXAMPLE.COM realm, and then to use those credentials to obtain credentials for use in the B.EXAMPLE.COM realm.

The client in this scenario treats the realm name as one might treat a DNS name. It repeatedly strips off the components of its own realm's name to generate the names of realms which are "above" it in the hierarchy until it reaches a point which is also "above" the service's realm. At that point it begins prepending components of the service's realm name until it reaches the service's realm. Each realm which is involved in the process is another "hop".

For example, using credentials in A.EXAMPLE.COM, authenticating to a service in B.EXAMPLE.COM:

```
A.EXAMPLE.COM → EXAMPLE.COM → B.EXAMPLE.COM
```

- A.EXAMPLE.COM and EXAMPLE.COM share a key for krbtgt/EXAMPLE.COM@A.EXAMPLE.COM

- EXAMPLE.COM and B.EXAMPLE.COM share a key for krbtgt/B.EXAMPLE.COM@EXAMPLE.COM

Another example, using credentials in SITE1.SALES.EXAMPLE.COM, authenticating to a service in EVERYWHERE.EXAMPLE.COM:

```
SITE1.SALES.EXAMPLE.COM → SALES.EXAMPLE.COM → EXAMPLE.COM → EVERYWHERE.EXAMPLE.COM
```

- SITE1.SALES.EXAMPLE.COM and SALES.EXAMPLE.COM share a key for krbtgt/SALES.EXAMPLE.COM@SITE1.SALES.EXAMPLE.COM

- SALES.EXAMPLE.COM and EXAMPLE.COM share a key for krbtgt/EXAMPLE.COM@SALES.EXAMPLE.COM

- EXAMPLE.COM and EVERYWHERE.EXAMPLE.COM share a key for krbtgt/EVERYWHERE.EXAMPLE.COM@EXAMPLE.COM

Another example, this time using realm names whose names share no common suffix (DEVEL.EXAMPLE.COM and PROD.EXAMPLE.ORG):

```
DEVEL.EXAMPLE.COM → EXAMPLE.COM → COM → ORG → EXAMPLE.ORG → PROD.EXAMPLE.ORG
```

- DEVEL.EXAMPLE.COM and EXAMPLE.COM share a key for krbtgt/EXAMPLE.COM@DEVEL.EXAMPLE.COM

- EXAMPLE.COM and COM share a key for krbtgt/COM@EXAMPLE.COM

- COM and ORG share a key for krbtgt/ORG@COM

- ORG and EXAMPLE.ORG share a key for krbtgt/EXAMPLE.ORG@ORG

- EXAMPLE.ORG and PROD.EXAMPLE.ORG share a key for krbtgt/PROD.EXAMPLE.ORG@EXAMPLE.ORG

The more complicated, but also more flexible, method involves configuring the capaths section of /etc/krb5.conf, so that clients which have credentials for one realm will be able to look up which realm is next in the chain which will eventually lead to the being able to authenticate to servers.

The format of the capaths section is relatively straightforward: each entry in the section is named after a realm in which a client might exist. Inside of that subsection, the set of intermediate realms from which the client must obtain credentials is listed as values of the key which corresponds to the realm in which a service might reside. If there are no intermediate realms, the value "." is used.

Here's an example:

```
[capaths]
A.EXAMPLE.COM = {
 B.EXAMPLE.COM = .
 C.EXAMPLE.COM = B.EXAMPLE.COM
 D.EXAMPLE.COM = B.EXAMPLE.COM
 D.EXAMPLE.COM = C.EXAMPLE.COM
}
```

In this example, clients in the A.EXAMPLE.COM realm can obtain cross-realm credentials for B.EXAMPLE.COM directly from the A.EXAMPLE.COM KDC.

If those clients wish to contact a service in theC.EXAMPLE.COM realm, they will first need to obtain necessary credentials from the B.EXAMPLE.COM realm (this requires that krbtgt/B.EXAMPLE.COM@A.EXAMPLE.COM exist), and then use those credentials to obtain credentials for use in the C.EXAMPLE.COM realm (using krbtgt/C.EXAMPLE.COM@B.EXAMPLE.COM).

If those clients wish to contact a service in the D.EXAMPLE.COM realm, they will first need to obtain necessary credentials from the B.EXAMPLE.COM realm, and then credentials from the C.EXAMPLE.COM realm, before finally obtaining credentials for use with the D.EXAMPLE.COM realm.

주목

Without a capath entry indicating otherwise, Kerberos assumes that cross-realm trust relationships form a hierarchy.

Clients in the A.EXAMPLE.COM realm can obtain cross-realm credentials from B.EXAMPLE.COM realm directly. Without the "." indicating this, the client would instead attempt to use a hierarchical path, in this case:

```
A.EXAMPLE.COM → EXAMPLE.COM → B.EXAMPLE.COM
```

46.6.10. 추가 자료

커베로스에 대한 보다 많은 정보를 원하신다면 다음 자료들을 참조하시기 바랍니다.

46.6.10.1. 설치된 문서 자료

- The Kerberos V5 Installation Guide and the Kerberos V5 System Administrator's Guide in PostScript and HTML formats. These can be found in the /usr/share/doc/krb5-server-<version-number>/

directory (where <version-number> is the version number of the krb5-server package installed on your system).

- The Kerberos V5 UNIX User's Guide in PostScript and HTML formats. These can be found in the /usr/share/doc/krb5-workstation-<version-number>/ directory (where <version-number> is the version number of the krb5-workstation package installed on your system).

- 커베로스 메뉴얼 페이지 — 커베로스 실행에 사용되는 다양한 응용 프로그램과 설정 파일에 대한 메뉴얼 페이지가 존재합니다. 다음은 중요한 메뉴얼 페이지의 목록입니다.

  클라이언트 응용 프로그램

  - man kerberos — 커베로스 시스템에 대한 소개. 증명(credentials)이 작동하는 방법 및 커베로스 티켓을 획득하고 제거하는 방법에 대하여 설명합니다. 메뉴얼 페이지 마지막 부분을 보시면 다른 관련 메뉴얼 페이지가 나와 있습니다.

  - man kinit — 이 명령을 사용하여 티켓 부여 티켓(TGT)을 획득하고 캐시하는 방법을 설명합니다.

  - man kdestroy — 이 명령을 사용하여 커베로스 인증을 제거하는 방법을 설명합니다.

  - man klist — 이 명령을 사용하여 캐시 저장된 커베로스 인증 목록을 보는 방법을 설명합니다.

  관리 응용 프로그램

  - man kadmin — 이 명령을 사용하여 커베로스 V5 데이터베이스를 관리하는 방법을 설명합니다.

  - man kdb5_util — 이 명령을 사용하여 커베로스 V5 데이터베이스에 저수준 관리 기능을 생성하고 실행하는 방법을 설명합니다.

  서버 응용 프로그램

  - man krb5kdc — 커베로스 V5 키 배포 센터에 사용되는 명령행 옵션을 보여줍니다.

  - man kadmind — 커베로스 V5 관리 서버에 사용되는 명령행 옵션을 보여줍니다.

  설정 파일

  - man krb5.conf — 커베로스 V5 라이브러리에 사용되는 설정 파일 내에서 사용 가능한 형식과 옵션을 설명합니다.

  - man krb5.conf — 커베로스 V5 AS 및 KDC에 사용되는 설정 파일 내에서 사용 가능한 형식과 옵션을 설명합니다.

46.6.10.2. 유용한 웹사이트

- http://web.mit.edu/kerberos/www/ — MIT의 Kerberos: The Network Authentication Protocol 웹페이지.

- http://www.nrl.navy.mil/CCS/people/kenh/kerberos-faq.html — 커베로스에 대한 자주 문의되는 질문과 답변 (FAQ).

- ftp://athena-dist.mit.edu/pub/kerberos/doc/usenix.PS — PostScript 버전의 Kerberos: An Authentication Service for Open Network Systems. 저자 Jennifer G. Steiner, Clifford Neuman, 그리고 Jeffrey I. Schiller. 이 문서는 커베로스를 설명하는 원저 (original paper)입니다.

- http://web.mit.edu/kerberos/www/dialogue.html — Designing an Authentication System: a Dialogue in Four Scenes originally by Bill Bryant in 1988, modified by Theodore Ts'o in 1997. This document is a conversation between two developers who are thinking through the creation of a Kerberos-style authentication system. The conversational style of the discussion make this a good starting place for people who are completely unfamiliar with Kerberos.

- http://www.ornl.gov/~jar/HowToKerb.html — How to Kerberize your site는 커베로스를 사용하는 네트워크에 대한 좋은 참조 자료입니다.

- http://www.networkcomputing.com/netdesign/kerb1.html — Kerberos Network Design Manual는 커베로스 시스템에 대한 전반적인 소개입니다.

46.7. 가상 사설 통신망 (Virtual Private Networks)

여러 사무실을 두고 있는 기업체에서는 전용선으로 각 사무실을 연결하여 중요한 정보를 보다 효율적으로 또한 안전하게 전송합니다. 예를 들면, 많은 기업체에서는 end-to-end 네트워킹 솔루션으로서 프레임 중계(frame relay) 서비스나 ATM (Asynchronous Transfer Mode) 서비스를 사용하여 각 사무실 간에 업무용 데이터를 전송합니다. 이러한 방법은 값비싼 엔터프라이즈 수준의 전용 디지털 회로에 드는 비용을 절감하면서 회사를 확장하기를 바라는 기업체에게는, 특히 중소 기업체(SMB)의 경우에는 비용이 많이 드는 솔루션입니다.

값비싼 엔터프라이즈 전용 회로에 대해 비용면에서 보다 저렴한 대안책으로서 가상 사설 통신망(Virtual Private Networks) (VPN)이 개발되었습니다. VPN은 전용 회로와 동일한 기능을 제공하면서 두 위치 간에 보안 디지털 통신을 성립하여 기존 LAN (Local Area Networks)으로부터 WAN (Wide Area Network)을 구성하며, 프레임 중계나 ATM과는 사용하는 전송 매체에서 차이가 있습니다. VPN은 전송 계층(transport layer) 데이터그램을 사용하여 IP 상으로 전송하기 때문에 인터넷을 통하여 원하는 목적지까지 안전한 선로를 통하게 됩니다. 대부분의 자유 소프트웨어 VPN은 전송 중인 자료의 보다 안전한 보안을 위해 공개 표준, 공개 소스 암호화 방식을 사용합니다.

일부 기관에서는 보안을 강화시키기 위해 하드웨어 VPN 솔루션을 사용하는 반면, 다른 기업체에서는 소프트웨어나 프로토콜 기반 VPN을 사용합니다. 하드웨어 VPN 솔루션을 제공하는 업체에는 Cisco, Nortel, IBM 및 Checkpoint와 같은 기업이 있습니다. FreeS/Wan이라는 리눅스 용 자유 소프트웨어 기반 VPN 솔루션은 표준 IPsec (Internet Protocol Security) 시스템을 사용합니다. 이러한 VPN 솔루션은 하드웨어나 소프트웨어 기반에 상관없이 한 사무실에서 다른 사무실 간 중간에서 IP를 연결하는 특수 라우터로 작동합니다.

46.7.1. VPN은 어떻게 작동합니까?

클라이언트에서 패킷을 전송시, 패킷은 VPN 라우터나 게이트웨이를 통하여 보내진 후, 인증 헤더 (Authentication Header) (AH)라는 라우팅과 인증에 대한 헤더 정보가 추가됩니다. 데이터가 암호화되면 마지막으로, ESP (Encapsulating Security Payload)가 삽입됩니다. 이는 나중에 암호 해독과 지시 사항을 처리하게 됩니다.

수신 VPN 라우터는 헤더 정보를 떼내어 읽어본 후, 데이터를 해독하여 목적지 (워크스테이션이나 네트워크 상 컴퓨터)로 라우팅합니다. 네트워크 간 연결을 사용하여 지역 네트워크 상 수신 컴퓨터는 암호 해독되어 처리될 준비가 된 패킷을 전송받습니다. 네트워크 간 VPN 연결에서 암호화/암호 해독 과정은 지역 컴퓨터에서 투명하게 진행됩니다.

이러한 강화된 보안 덕분에, 크래커는 패킷을 중간에서 가로챈 경우 그 패킷을 암호 해독해야 내용을 볼 수 있습니다. 또한 침입자가 man-in-the-middle 공격 방식을 이용하여 서버와 클라이언트 사이의 통신을 엿듣기 시도할 경우 인증 세션에 사용되는 최소 한개의 비밀키가 있어야 합니다. 이렇듯 VPN은 여러 계층의 인증과 암호 방식을 사용하기 때문에, 다른 장소에 위치한 원격 시스템을 하나의 인트라넷으로 안전하고 효율적으로 연결할 수 있습니다.

46.7.2. VPN 및 Red Hat Enterprise Linux

Red Hat Enterprise Linux는 사용자에게 WAN을 보안 연결하기 위해 다양한 소프트웨어 솔루션을 제공해 드립니다. IPsec (Internet Protocol Security)은 Red Hat Enterprise Linux가 지원하는 VPN으로서, 본사와 지역 사무실 및 원거리 근무자를 하나로 묶는 인트라넷을 안전하고 효율적으로 구현 가능합니다.

46.7.3. IPsec

Red Hat Enterprise Linux는 인터넷과 같은 공중 네트워크 상에서 보안 터널을 사용하여 원격 호스트와 네트워크를 연결해주는 IPsec을 지원합니다. IPsec은 호스트 간 (한 컴퓨터 워크스테이션에서 다른 워크스테이션으로) 또는 네트워크 간 (한 LAN/WAN에서 다른 LAN/WAN으로) 연결을 사용하여 구현됩니다.

Red Hat Enterprise Linux에서는 시스템을 안전하게 연결하고 상호 인증하기 위해 IETF (Internet Engineering Task Force)가 구현한 프로토콜인 IKE (Internet Key Exchange:인터넷 키 교환)을 사용하여 IPsec을 구현합니다.

46.7.4. IPsec 연결 생성하기

An IPsec connection is split into two logical phases. In phase 1, an IPsec node initializes the connection with the remote node or network. The remote node or network checks the requesting node's credentials and both parties negotiate the authentication method for the connection.

Red Hat Enterprise Linux 시스템에서는 IPsec 연결을 위해 pre-shared key(사전 공유 키) 방식의 IPsec 컴퓨터 인증을 사용합니다. 사전 공유 키 방식을 사용하는 IPsec 연결에서는 양 컴퓨터가 동일한 키를 가지고 공유하고 있어야 IPsec 연결 2 단계로 넘어갈 수 있습니다.

IPsec 연결 2 단계에서는 IPsec 컴퓨터 간에 SA (Security Association)가 설정됩니다. 암호화 방식, 비밀 세션키 교환 변수 등과 같은 설정 정보를 담은 SA 데이터베이스가 만들어지며, 원격 컴퓨터와 네트워크 사이에서 실제 IPsec 연결을 관리하는 단계입니다.

Red Hat Enterprise Linux에서 IPsec을 구현하는데 IKE를 사용하여 인터넷 상에서 호스트 간에 키를 공유합니다. racoon 키 관리 데몬이 IKE 키를 배포하고 교환합니다. 이러한 데몬에 관한 보다 자세한 정보는 racoon 메뉴얼 페이지를 참조하시기 바랍니다.

46.7.5. IPsec 설치

IPsec을 구현하기 위해서는 모든 IPsec 호스트(호스트 간 설정을 사용하는 경우) 또는 라우터(네트워크 간 설정을 사용하는 경우)에 ipsec-tools RPM 패키지가 설치되어 있어야 합니다. 이 RPM 패키지에는 IPsec 연결을 설정하는데 필요한 라이브러리, 데몬 및 설정 파일이 포함되어 있습니다. 패키지 내용물은 다음과 같습니다:

- /sbin/setkey — 커널에서 키 관리와 IPsec의 보안 속성을 조정합니다. 이 실행 파일은 racoon 키 관리 데몬에 의해 조정됩니다. setkey에 대한 보다 자세한 정보는 setkey(8) 메뉴얼 페이지를 참조하시기 바랍니다.

- /usr/sbin/racoon — the IKE key management daemon, used to manage and control security associations and key sharing between IPsec-connected systems.

- /etc/racoon/racoon.conf — racoon 데몬 설정 파일으로서 연결에 사용된 인증 방법 및 암호화 알고리즘을 포함한 다양한 측면의 IPsec 연결을 설정하는데 사용됩니다. 사용 가능한 모든 지시자 목록을 보시려면 racoon.conf(5) 메뉴얼 페이지를 참조하시기 바랍니다.

Red Hat Enterprise Linux에서 IPsec을 설정하기 위해, Network Administration Tool: 네트워크 관리 도구를 사용하시거나, 수동으로 네트워킹과 IPsec 설정 파일을 수정하실 수 있습니다.

- To connect two network-connected hosts via IPsec, refer to 46.7.6절. "IPsec 호스트 간 설정".
- To connect one LAN/WAN to another via IPsec, refer to 46.7.7절. "IPsec 네트워크 간 설정".

46.7.6. IPsec 호스트 간 설정

IPsec을 호스트 간 연결을 통하여 데스크탑이나 워크스테이션들을 연결하도록 설정 가능합니다. 이러한 유형의 연결은 각 호스트가 연결된 네트워크를 사용하여 양 호스트 사이에 보안 터널을 생성합니다. 호스트 간 연결에 필요한 요건은 각 호스트에 IPsec만 설정하면 됩니다. 호스트에서 IPsec 연결을 생성하기 위해서는 공중 네트워크 (예, 인터넷)와 Red Hat Enterprise Linux에 연결할 전용선만 있으면 됩니다.

46.7.6.1. 호스트 간 설정

호스트간 IPsec 연결은 두 시스템 사이에서 암호화된 연결이어야 하며, 같은 인증키를 사용하고 IPsec을 실행하여야 합니다. IPsec 연결이 활성화되었을 경우, 두 호스트사이의 모든 네트워크 소통은 암호화됩니다.

호스트간 IPsec 연결을 설정하기 위해,각각의 호스트에 사용하는 단계는 다음과 같습니다:

> **주의**
>
> 설정하시고 있는 실제 컴퓨터에 다음과 같은 과정을 실행하셔야 합니다. 원격으로 IPsec 연결 설정을 시도하지 마십시오.

1. 명령 쉘에서, system-config-network를 입력하여 Network Administration Tool: 네트워크 관리 도구를 시작합니다.

2. IPsec 탭에서, 새로 시작 버튼을 클릭하여 IPsec 설정 마법사를 시작합니다.

3. 다음 버튼을 클릭하여 호스트 간 IPsec 연결을 설정합니다.

4. Enter a unique name for the connection, for example, ipsec0. If required, select the check box to automatically activate the connection when the computer starts. Click Forward to continue.

5. 연결 유형으로 호스트 간 암호화를 선택한 후, 다음 버튼을 클릭합니다.

6. 수동 또는 자동으로 사용할 암호화 유형을 선택합니다.

   수동으로 암호화할 것을 선택하신 경우, 암호화 키는 과정의 나중에 제공되어야 합니다. 자동으로 암호화할 것을 선택하신 경우, racoon 데몬이 암호화키를 관리합니다. 자동 암호화를 사용하실 경우, ipsec-tools 패키지가 설치되어 있어야 합니다.

   다음 버튼을 클릭하여 계속 진행합니다.

7. 원격 호스트의 IP 주소를 입력합니다.

   원격 호스트의 IP 주소를 설정하기 위해, 원격 호스트 상에서 다음 명령을 사용합니다:

```
ifconfig <device>
```

where <device> is the Ethernet device that you want to use for the VPN connection.

시스템에 하나의 이더넷 카드만이 있을 경우, 일반적으로 장치명은 eth0가 됩니다. 다음은 이러한 명령에 관련된 예입니다 (이는 출력 결과의 예에 불과함에 유의합니다):

```
eth0      Link encap:Ethernet  HWaddr 00:0C:6E:E8:98:1D
          inet addr:172.16.44.192  Bcast:172.16.45.255  Mask:255.255.254.0
```

IP 주소는 inet addr: 레이블 다음에 오는 숫자입니다.

주의

호스트간 연결의 경우 양쪽 호스트는 공개의 라우트 가능한 주소여야 합니다. 또는 양쪽 호스트가 같은 LAN을 사용할 경우 개인의 라우트 불가능한 주소일 수 있습니다 (예, 10.x.x.x 또는 192.168.x.x 범위에서)

If the hosts are on different LANs, or one has a public address while the other has a private address, refer to 46.7.7절. "IPsec 네트워크 간 설정".

다음 버튼을 클릭하여 계속 진행합니다.

8. If manual encryption was selected in step 6, specify the encryption key to use, or click Generate to create one.

   a. 인증키를 지정하거나 또는 생성 버튼을 클릭하여 새로 생성합니다. 인증키는 숫자와 문자의 조합으로 만들어 질 수 있습니다.

   b. 다음 버튼을 클릭하여 계속 진행합니다.

9. IPsec — 요약 페이지에서 내용을 확인한 후, 적용 버튼을 클릭합니다.

10. Click File > Save to save the configuration.

    변경 사항을 적용하기 위해 네트워크를 다시 시작하셔야 합니다. 다음 명령을 사용하여 네트워크를 다시 시작합니다:

    ```
    service network restart
    ```

11. 목록에서 IPsec 연결을 선택한 후 활성화 버튼을 클릭합니다.

12. Repeat the entire procedure for the other host. It is essential that the same keys from step 8 be used on the other hosts. Otherwise, IPsec will not work.

After configuring the IPsec connection, it appears in the IPsec list as shown in 그림 46.10. "IPsec Connection".

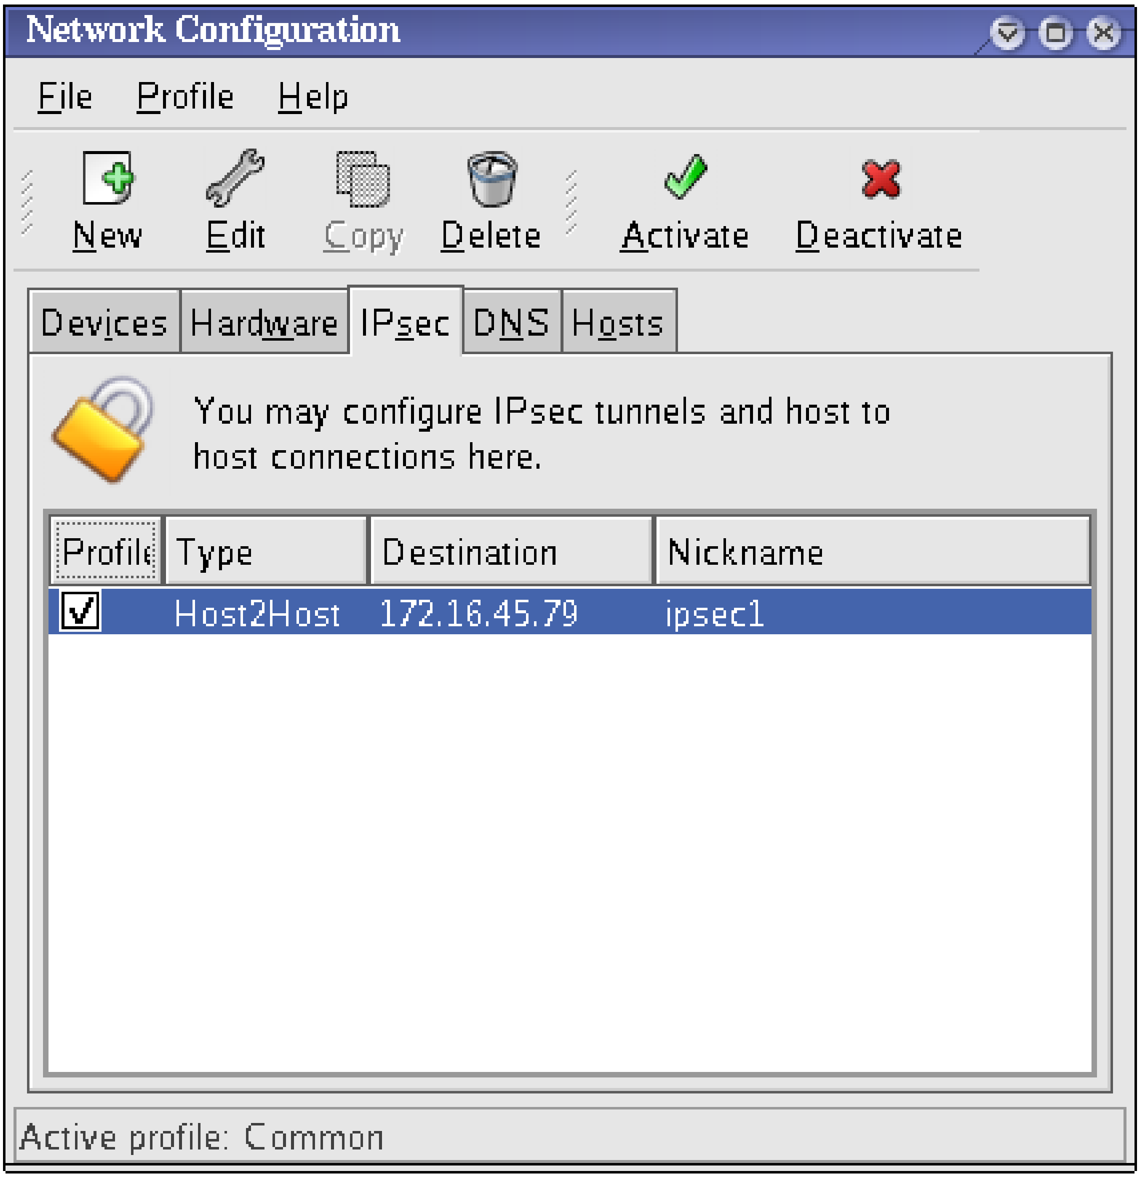

You may configure IPsec tunnels and host to host connections here.

Profile	Type	Destination	Nickname
☑	Host2Host	172.16.45.79	ipsec1

Active profile: Common

그림 46.10. IPsec Connection

IPsec 연결이 설정되면 다음의 파일이 생성됩니다:

- /etc/sysconfig/network-scripts/ifcfg-<nickname>
- /etc/sysconfig/network-scripts/keys-<nickname>
- /etc/racoon/<remote-ip>.conf
- /etc/racoon/psk.txt

자동 암호화가 선택되면, /etc/racoon/racoon.conf 역시 생성됩니다.

When the interface is up, /etc/racoon/racoon.conf is modified to include <remote-ip>.conf.

46.7.6.2. 수동으로 IPsec 호스트 간 설정

연결을 생성하는 첫번째 단계는 각 워크스테이션의 시스템 정보와 네트워크 정보를 모으는 것입니다. 호스트 간 연결을 위해서는 다음과 같은 정보를 수집하셔야 합니다:

- 각 호스트의 IP 주소

- 예를 들어, ipsec1과 같은 고유 이름. 이는 IPsec 연결을 식별하고 다른 장치나 연결에서 이를 구별하기 위해 사용됩니다.

- 고정 암호키 또는 racoon에 의해 자동으로 생성된 암호키

- 연결을 초기화하고 세션 중 암호키를 교환하는데 사용되는 미리 공유된 인증키.

예를 들어 워크스테이션 A와 워크스테이션 B가 IPsec터널을 통하여 연결하고자 한다고 가정합니다. Key_Value01의 값을 이미 공유된 키로 사용하여 연결하고자 하며, 양 사용자가 racoon 데몬이 자동으로 인증키를 생성하여 각 호스트 간에 공유하는 것에 동의하여 이 연결을 ipsec1으로 이름 붙였다고 가정합니다.

주의

대소문자, 숫자, 구두점을 혼합하여 PSK를 선택하셔야 합니다. 쉽게 추측할 수 있는 PSK는 보안 위험을 초래합니다.

각 호스트에 대해 같은 연결명을 사용하실 필요가 없습니다. 설치에 편하고 의미있는 연결명을 선택하시면 됩니다.

다음은 워크스테이션 B와 호스트 간 IPsec을 연결하기 위해 사용된 워크스테이션 A의 IPsec 설정 파일입니다. 이 예시에서 이 연결을 식별하기 위해 사용된 고유 이름은 ipsec1이므로 결과적으로 파일 이름은 /etc/sysconfig/network-scripts/ifcfg-ipsec1이 됩니다.

```
DST=X.X.X.X
TYPE=IPSEC
ONBOOT=no
IKE_METHOD=PSK
```

워크스테이션 A는 X.X.X.X 부분을 워크스테이션 B의 IP 주소로 대체하고, 워크스테이션 B는 X.X.X.X를 워크스테이션 A의 IP 주소로 대체합니다. 이 연결은 부팅시 시작되도록 (ONBOOT=no) 설정되지 않았으며 이미 공유된 키 인증 방식 (IKE_METHOD=PSK)을 사용합니다.

다음은 이미 공유된 키 파일 (/etc/sysconfig/network-scripts/keys-ipsec1)의 내용입니다. 이 파일은 양 워크스테이션이 상대방을 인증하는데 필요합니다. 이 파일의 내용은 양 컴퓨터에서 동일해야 하며 루트 사용자만이 이 파일을 읽거나 수정할 수 있습니다.

```
IKE_PSK=Key_Value01
```

중요

keys-ipsec1 파일을 루트 사용자만 읽고 수정할 수 있도록 하시려면, 파일을 만드신 후 다음 명령을 입력하십시오:

```
chmod 600 /etc/sysconfig/network-scripts/keys-ipsec1
```

언제든지 인증키를 변경하시려면 양 워크스테이션에서 keys-ipsec1 파일을 수정하시면 됩니다. 양 키가 동일해야만 제대로 연결할 수 있습니다.

다음은 원격 호스트로 1 단계 연결하는데 사용되는 설정 파일의 예입니다. 이 파일의 이름은 X.X.X.X.conf입니다 (여기서 X.X.X.X는 원격 IPsec 호스트의 IP 주소로 대체하십시오). 이 파일은 IPsec 터널이 활성화되면 자동으로 생성되기 때문에 직접 수정하시면 안됩니다.

```
remote X.X.X.X
{
         exchange_mode aggressive, main;
  my_identifier address;
  proposal {
   encryption_algorithm 3des;
  hash_algorithm sha1;
  authentication_method pre_shared_key;
  dh_group 2 ;
 }
}
```

IPsec 연결이 초기화될때 생성된 기본 단계 1 설정 파일에는 Red Hat Enterprise Linux가 IPsec를 구현하는데 사용하는 다음과 같은 문구가 포함됩니다:

remote X.X.X.X

이 줄은 이 설정 파일에서 다음 부분을 IP 주소가 X.X.X.X인 원격 컴퓨터에만 적용하도록 지정합니다.

exchange_mode aggressive

윗 줄은 Red Hat Enterprise Linux의 기본 IPsec 설정으로서 여러 호스트 간에 IPsec 연결을 설정하는 동시에 연결 작업 부하를 낮춰주는 적극적 (aggressive) 인증 모드를 사용합니다.

my_identifier address

컴퓨터 인증시 사용할 식별 방식을 지정합니다. Red Hat Enterprise Linux는 IP 주소를 사용하여 컴퓨터를 식별합니다.

encryption_algorithm 3des

인증시 사용할 암호화 방식을 지정합니다. 3DES (Triple Data Encryption Standard)가 기본으로 사용됩니다.

hash_algorithm sha1;

컴퓨터 간에 1 단계 협상 단계에서 사용할 해시 알고리즘을 지정합니다. 기본값으로 SHA (Secure Hash Algorithm) 버전 1이 사용됩니다.

authentication_method pre_shared_key

컴퓨터 간 협상시 사용할 인증 방식을 지정합니다. Red Hat Enterprise Linux는 인증을 위해 기존 공유 키(pre-shared keys)를 기본으로 사용합니다.

dh_group 2

동적으로 생성된 세션키를 분배하는데 사용할 Diffie-Hellman 그룹 번호를 지정합니다. modp1024 (그룹 2)가 기본값입니다.

46.7.6.2.1. Racoon 설정 파일

The /etc/racoon/racoon.conf files should be identical on all IPsec nodes except for the include "/etc/racoon/X.X.X.X.conf" statement. This statement (and the file it references) is generated when the IPsec tunnel is activated. For Workstation A, the X.X.X.X in the include statement is Workstation B's IP address. The opposite is true of Workstation B. The following shows a typical racoon.conf file when the IPsec connection is activated.

```
# Racoon IKE daemon configuration file.
# See 'man racoon.conf' for a description of the format and entries.
```

```
path include "/etc/racoon";
path pre_shared_key "/etc/racoon/psk.txt";
path certificate "/etc/racoon/certs";

sainfo anonymous
{
        pfs_group 2;
        lifetime time 1 hour ;
        encryption_algorithm 3des, blowfish 448, rijndael ;
        authentication_algorithm hmac_sha1, hmac_md5 ;
        compression_algorithm deflate ;
}
include "/etc/racoon/X.X.X.X.conf";
```

이 기본 racoon.conf 파일에는 IPsec 설정, 기존 공유 키 파일 및 인증서의 지정된 경로가 포함됩니다. sainfo anonymous 항목은 IPsec 컴퓨터 간 2 단계 SA — IPsec 연결 속성 (사용된 암호화 알고리즘 포함) 및 키 교환 방식을 설명합니다. 다음은 2 단계 항목을 정의하는 목록입니다:

sainfo anonymous
: IPsec 인증서만 일치한다면 SA가 어느 컴퓨터와도 익명으로 연결을 초기화할 수 있다는 것을 의미합니다.

pfs_group 2
: Diffie-Hellman 키 교환 프로토콜을 지정합니다. 이 프로토콜은 IPsec 연결 2 단계에서 IPsec 시스템이 상호 임시 세션키를 분배하는 방식을 결정합니다. 기본적으로 Red Hat Enterprise Linux는 Diffie-Hellman 암호키 교환 그룹의 그룹 2 (modp1024)를 사용하여 IPsec을 구현합니다. 그룹 2는 1024 비트 방식을 사용하기 때문에 침입자가 비밀키를 알아낸 경우에도 이전 IPsec 전송 내용을 암호 해독 불가능합니다.

lifetime time 1 hour
: 이 변수는 SA의 수명을 시간이나 데이터 바이트 수 단위로 지정합니다. Red Hat Enterprise Linux에서 IPsec 구현시 수명은 한 시간이 기본입니다.

encryption_algorithm 3des, blowfish 448, rijndael
: 2 단계에서 사용할 암호화 방식을 지정합니다. Red Hat Enterprise Linux는 3DES, 448-bit Blowfish, 및 Rijndael(Advanced Encryption Standard 또는 AES에 사용된 암호방식)을 지원합니다.

authentication_algorithm hmac_sha1, hmac_md5
: 인증에 사용할 해시 알고리즘을 열거합니다. 지원 가능한 모드는 sha1와 md5 해시 메시지 인증 코드 (HMAC) 입니다.

compression_algorithm deflate
: IPCOMP (IP Payload Compression)을 위해 Deflate 압축 알고리즘을 사용하도록 지정합니다. 이 알고리즘을 사용하면 IP 데이터그램을 보다 빠르게 압축 전송할 수 있습니다.

연결을 시작하려면 각 호스트에서 다음 명령을 실행하시면 됩니다:

```
ifup <nickname>
```

where <nickname> is the name you specified for the IPsec connection.

To test the IPsec connection, run the tcpdump utility to view the network packets being transferred between the hosts and verify that they are encrypted via IPsec. The packet should include an AH header and should be shown as ESP packets. ESP means it is encrypted. For example:

```
~]# tcpdump -n -i eth0 host <targetSystem>
```

```
IP 172.16.45.107 > 172.16.44.192: AH(spi=0x0954ccb6,seq=0xbb): ESP(spi=0x0c9f2164,seq=0xbb)
```

46.7.7. IPsec 네트워크 간 설정

IPsec can also be configured to connect an entire network (such as a LAN or WAN) to a remote network using a network-to-network connection. A network-to-network connection requires the setup of IPsec routers on each side of the connecting networks to transparently process and route information from one node on a LAN to a node on a remote LAN. 그림 46.11. "A network-to-network IPsec tunneled connection" shows a network-to-network IPsec tunneled connection.

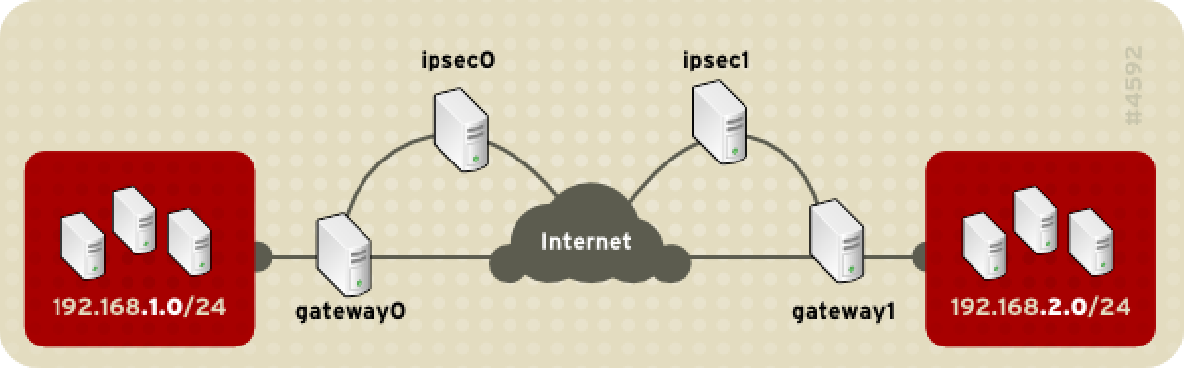

그림 46.11. A network-to-network IPsec tunneled connection

도표에서는 인터넷으로 분리된 별개의 두 LAN을 보여줍니다. 이 LAN은 인터넷을 통해 보안 터널을 사용하여 연결을 인증하고 초기화하는데 IPsec 라우터를 사용합니다. 이러한 두 LAN 사이에서 전송 중인 패킷을 가로채어 암호를 해독하기 위해서는 brute-force (암호 알고리즘에 사용할 수 있는 모든 키를 전부 조사하는 방법)을 사용해야 합니다. IPsec 패킷의 처리, 암호화/암호 해독 및 라우팅 과정은 전적으로 IPsec 라우터에 의해 처리되므로 192.168.1.0/24 IP 상 컴퓨터에서 192.168.2.0/24 컴퓨터 간에 주고받는 통신 과정은 그 컴퓨터에서는 전혀 알 수 없습니다.

네트워크 간 연결에 필요한 정보는 다음과 같습니다:

- 외부에서 접근 가능한 전용 IPsec 라우터의 IP 주소
- IPsec 라우터에서 처리하는 LAN/WAN의 네트워크 주소 범위 (예, 192.168.1.0/24 또는 10.0.1.0/24)
- 네트워크로 연결된 시스템에서 인터넷으로 데이터를 라우팅하는 게이트웨이 장치의 IP 주소
- 예를 들어, ipsec1과 같은 고유 이름. 이는 IPsec 연결을 식별하고 다른 장치나 연결에서 이를 구별하기 위해 사용됩니다.
- 고정 암호키 또는 racoon에 의해 자동으로 생성된 암호키
- 연결을 초기화하고 세션 중 암호키를 교환하는데 사용되는 미리 공유된 인증키.

46.7.7.1. 네트워크 간 (VPN) 설정

네트워크 간 IPsec 연결에서는 사설 서브넷에 대한 네트워크 소통을 라우트하여 각각의 네트워크에 하나씩 두개의 IPsec 라우터를 사용합니다.

For example, as shown in 그림 46.12. "Network-to-Network IPsec", if the 192.168.1.0/24 private network sends network traffic to the 192.168.2.0/24 private network, the packets go through gateway0, to ipsec0, through the Internet, to ipsec1, to gateway1, and to the 192.168.2.0/24 subnet.

IPsec 라우터에는 공개적으로 주소를 지정할 수 있는 IP 주소와 사설 네트워크와 연결된 다른 이더넷 장치가 있어야 합니다. 암호화된 연결이 있는 다른 IPsec 라우터를 통해 소통하려할 경우, IPsec 라우터만을 통해 소통이 이루어집니다.

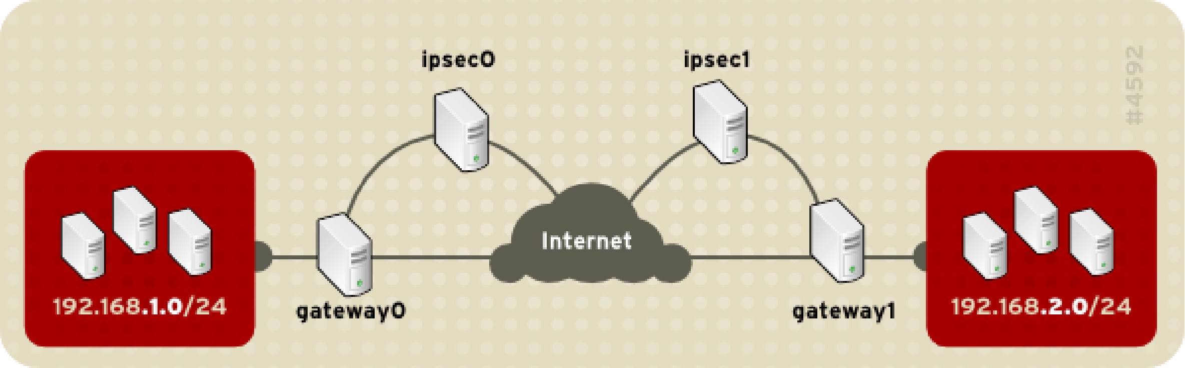

그림 46.12. Network-to-Network IPsec

각 IP 라우터와 인터넷 간의 방화벽 및 각 IPsec 라우터와 서브넷 게이트웨이 간의 인트라넷 방화벽을 포함하는 네트워크 설정 대체 옵션. 서브넷의 IPsec 라우터와 게이트웨이는 두개의 이더넷 장치를 가진 하나의 시스템이 될 수 있습니다: 하나는 공개 IP 주소를 가지고 IPsec 라우터로 작동하며; 다른 하나는 사설 IP 주소를 가지고 사설 서브넷에 대한 게이트웨이로 작동합니다. 각각의 IPsec 라우터는 사설 네트워크 또는 공개 게이트웨이에 대해 게이트웨이를 사용하여 패킷을 다른 IPsec 라우터로 전송합니다.

다음과 같은 과정을 실행하여 네트워크 간 IPsec 연결을 설정합니다:

1. 명령 쉘에서, system-config-network를 입력하여 Network Administration Tool: 네트워크 관리 도구를 시작합니다.

2. IPsec 탭에서, 새로 시작 버튼을 클릭하여 IPsec 설정 마법사를 시작합니다.

3. 다음 버튼을 클릭하여 네트워크 간 IPsec연결을 설정합니다.

4. Enter a unique nickname for the connection, for example, ipsec0. If required, select the check box to automatically activate the connection when the computer starts. Click Forward to continue.

5. 연결 유형으로 네트워크 간 암호화 (VPN)를 선택하고, 다음 버튼을 클릭합니다.

6. 수동 또는 자동으로 사용할 암호화 유형을 선택합니다.

   수동으로 암호화할 것을 선택하신 경우, 암호화 키는 과정의 나중에 제공되어야 합니다. 자동으로 암호화할 것을 선택하신 경우, racoon 데몬이 암호화키를 관리합니다. 자동 암호화를 사용하실 경우, ipsec-tools 패키지가 설치되어 있어야 합니다.

   다음 버튼을 클릭하여 계속 진행합니다.

7. 지역 네트워크 페이지에서, 다음의 정보를 입력합니다:

   - 지역 네트워크 주소 — 사설 네트워크에 연결된 IPsec 라우터에 있는 장치의 IP 주소.
   - 지역 서브넷 마스크 — 지역 네트워크 IP 주소의 서브넷 마스크.
   - 지역 네트워크 게이트웨이 — 사설 서브넷의 게이트웨이.

   다음 버튼을 클릭하여 계속 진행합니다.

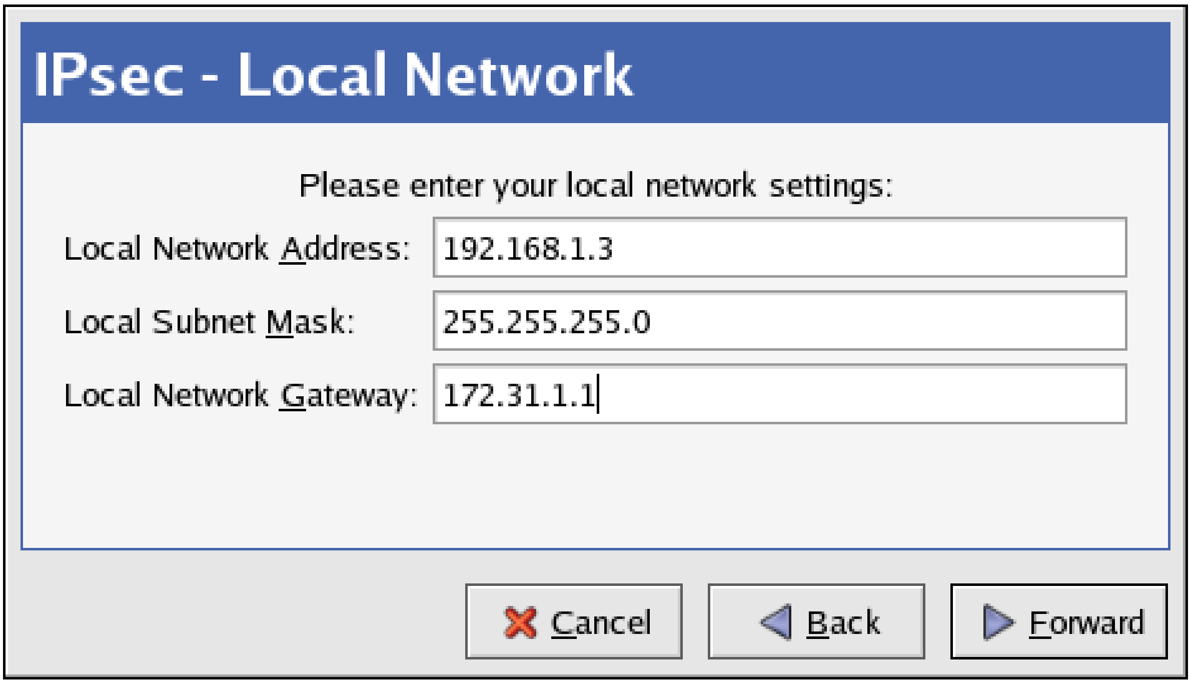

IPsec - Local Network

Please enter your local network settings:

Local Network Address: 192.168.1.3

Local Subnet Mask: 255.255.255.0

Local Network Gateway: 172.31.1.1

그림 46.13. Local Network Information

8. 원격 네트워크 페이지에서, 다음의 정보를 입력합니다:

   - 원격 IP 주소 — 다른 사설 네트워크에 대한 IPsec 라우터에 공개적으로 주소 지정이 가능한 IP 주소. 이 예시에서, ipsec0는 ipsec1의 공개적으로 주소 지정이 가능한 IP 주소를 입력하며, ipsec1는 반대로 적용합니다.

   - 원격 네트워크 주소 — 다른 IPsec 라우터 뒤에 있는 사설 서브넷의 네트워크 주소. 예를 들어, ipsec1을 설정하는 경우 192.168.1.0을 입력하고, ipsec0 설정하는 경우 192.168.2.0을 입력합니다.

   - 원격 서브넷 마스크 — 원격 IP 주소의 서브넷 마스크.

   - 원격 네트워크 게이트웨이 — 원격 네트워크 주소에 대한 게이트웨이의 IP 주소.

   - If manual encryption was selected in step 6, specify the encryption key to use or click Generate to create one.

     인증키를 지정하거나 또는 생성 버튼을 클릭하여 키를 생성합니다. 이러한 키는 숫자와 문자의 조합으로 만들어 질 수 있습니다.

   다음 버튼을 클릭하여 계속 진행합니다.

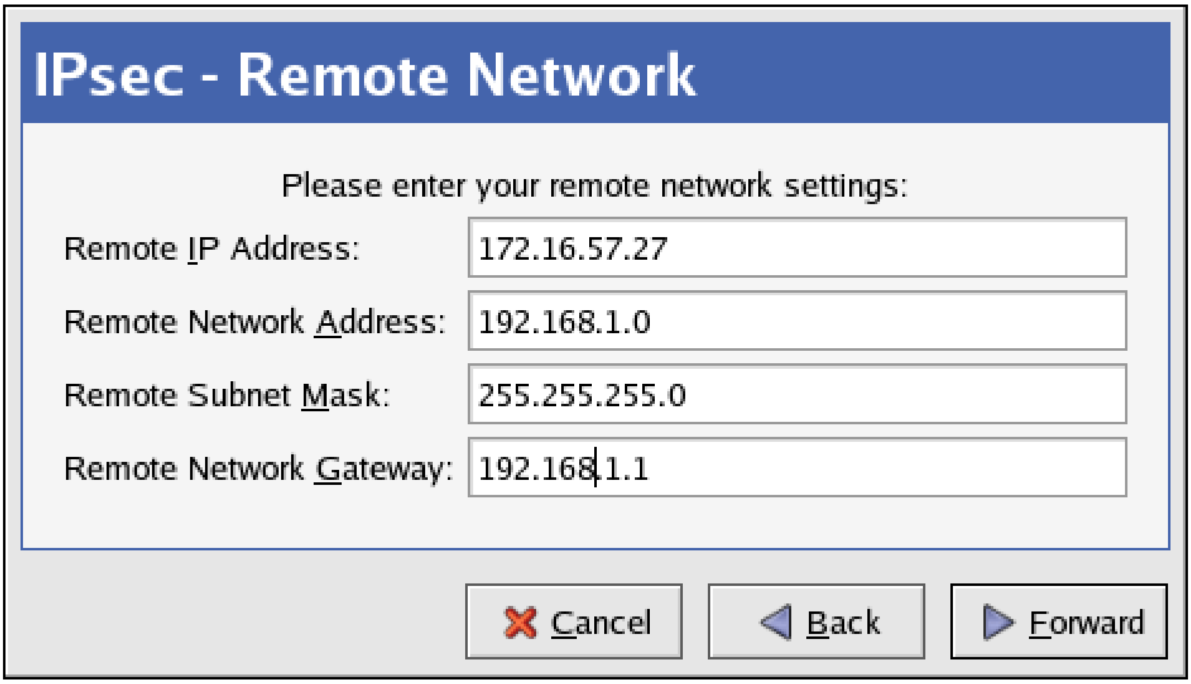

IPsec - Remote Network

Please enter your remote network settings:

Remote IP Address: 172.16.57.27

Remote Network Address: 192.168.1.0

Remote Subnet Mask: 255.255.255.0

Remote Network Gateway: 192.168.1.1

그림 46.14. Remote Network Information

9. IPsec — 요약 페이지에서 내용을 확인한 후, 적용 버튼을 클릭합니다.

10. Select File > Save to save the configuration.

11. 목록에서 IPsec 연결을 선택한 후, 활성화 버튼을 클릭하여 연결을 활성화함.

12. IP forwarding 활성화:

    a. Edit /etc/sysctl.conf and set net.ipv4.ip_forward to 1.

    b. 변경 사항이 적용되도록 다음 명령을 실행합니다:

    ```
    sysctl -p /etc/sysctl.conf
    ```

IPsec 연결을 활성화하기 위한 네트워크 스크립트는 자동으로 네트워크 라우터를 생성하여 필요한 경우 IPsec 라우터를 통해 패킷을 전송합니다.

46.7.7.2. 수동으로 IPsec 네트워크 간 설정

예를 들어 LANA (lana.example.com)와 LAN B (lanb.example.com)가 IPsec 터널을 통하여 서로 연결하고자 한다고 가정합니다. LAN A의 네트워크 주소는 192.168.1.0/24 범위에 속하고 LAN B의 네트워크 주소는 192.168.2.0/24 범위를 사용한다고 합시다. LAN A의 게이트웨이 IP 주소는 192.168.1.254 이며 LAN B의 게이트위에 IP 주소는 192.168.2.254 입니다. IPsec 라우터는 각각의 LAN에서 분리되어 두개의 네트워크 장치를 사용합니다: eth0는 인터넷에 접속하는 외부에서 접근 가능한 정적 IP 주소에 할당되었으며, 반면 eth1은 한 네트워크 노드에서 원격 네트워크 노드로 LAN 패킷을 처리하고 전송하는 라우팅 지점으로 사용됩니다.

각 네트워크 간 IPsec 연결은 r3dh4tl1nux을 이미 공유된 키로 사용하며 A와 B의 관리자가 racoon 데몬이 자동으로 인증키를 생성하고 각 IPsec 라우터 간에 인증키를 공유하도록 동의했다고 가정합니다. LAN A의 관리자는 IPsec 연결을 ipsec0이라 이름 붙인 반면 LANB의 관리자는 IPsec 연결을 ipsec1이라고 이름 붙였습니다.

다음 예시는 LAN A에서 네트워크 간 IPsec 연결에 사용된 ifcfg 파일의 내용을 보여줍니다. 이 예시에서 이 연결을 식별하기 위해 사용된 고유 이름은 ipsec0입니다, 따라서 결과적으로 파일 이름은 /etc/sysconfig/network-scripts/ifcfg-ipsec0이 됩니다.

```
TYPE=IPSEC
ONBOOT=yes
IKE_METHOD=PSK
SRCGW=192.168.1.254
DSTGW=192.168.2.254
SRCNET=192.168.1.0/24
DSTNET=192.168.2.0/24
DST=X.X.X.X
```

다음의 목록은 이 파일의 내용을 설명합니다:

TYPE=IPSEC
: 연결 유형을 지정합니다.

ONBOOT=yes
: 부팅시 초기화해야 하는 연결을 지정합니다.

IKE_METHOD=PSK
: 기존 공유 키 인증 방식을 사용하는 연결을 지정합니다.

SRCGW=192.168.1.254
: 소스 게이트웨이의 IP 주소. LAN A는 LAN A 게이트웨이를 LAN B는 LAN B 게이트웨이를 지정합니다.

DSTGW=192.168.2.254
: 수신지 게이트웨이의 IP 주소. LAN A는 LAN B 게이트웨이로 LAN B는 LAN A 게이트웨이로 지정합니다.

SRCNET=192.168.1.0/24
: IPsec 연결에 대한 소스 네트워크를 지정합니다, 이 예시에서는 네트워크 범위가 LAN A로 되어 있습니다.

DSTNET=192.168.2.0/24
: IPsec 연결에 대한 수신지 네트워크를 지정합니다, 이 예시에서는 네트워크 범위가 LAN B로 되어 있습니다.

DST=X.X.X.X
: 외부에서 접근 가능한 LAN B의 IP 주소.

다음은 /etc/sysconfig/network-scripts/keys-ipsecX 기존 공유 키 파일 (여기서 X는 LAN A의 경우 0이며 LAN B는 1 입니다)의 내용으로서 이 파일은 양 네트워크가 상대방을 인증하는데 사용됩니다. 이 파일의 내용은 두 네트워크에서 동일해야 하며 루트 사용자만이 이 파일을 읽거나 작성할 수 있습니다.

```
IKE_PSK=r3dh4tl1nux
```

중요

keys-ipsecX 파일을 루트 사용자만 읽고 수정할 수 있도록 설정하시려면, 파일을 생성 후 다음 명령을 입력합니다:

```
chmod 600 /etc/sysconfig/network-scripts/keys-ipsec1
```

언제든지 인증키를 변경하려면 양 IPsec 라우터에서 keys-ipsecX 파일을 수정합니다. 반드시 양 키가 같아야 제대로 연결됩니다.

다음은 IPsec 연결에 사용된 /etc/racoon/racoon.conf 설정 파일입니다. 파일 마지막 부분의 include 부분은 자동으로 생성되며 IPsec 터널이 실행된 경우에만 나타납니다.

```
# Racoon IKE daemon configuration file.
# See 'man racoon.conf' for a description of the format and entries.
path include "/etc/racoon";
path pre_shared_key "/etc/racoon/psk.txt";
path certificate "/etc/racoon/certs";

sainfo anonymous
{
 pfs_group 2;
 lifetime time 1 hour ;
 encryption_algorithm 3des, blowfish 448, rijndael ;
 authentication_algorithm hmac_sha1, hmac_md5 ;
 compression_algorithm deflate ;
}
include "/etc/racoon/X.X.X.X.conf"
```

다음은 원격 네트워크로 연결하는데 사용되는 설정 파일입니다. 파일의 이름은 X.X.X.X.conf이며 여기서 X.X.X.X를 원격 IPsec 라우터의 IP 주소로 대체합니다. 이 파일은 IPsec 터널이 활성화되면 자동으로 생성되기 때문에 직접 수정하시면 안됩니다.

```
remote X.X.X.X
{
        exchange_mode aggressive, main;
 my_identifier address;
 proposal {
  encryption_algorithm 3des;
  hash_algorithm sha1;
  authentication_method pre_shared_key;
  dh_group 2 ;
 }
}
```

IPsec 연결을 시작하기 전에 커널에서 IP fowarding 기능을 활성화합니다. IP forwarding 기능을 활성화하려면 다음 명령을 실행합니다:

1. Edit /etc/sysctl.conf and set net.ipv4.ip_forward to 1.

2. 변경 사항이 적용되도록 다음 명령을 실행합니다:

```
sysctl -p /etc/sysctl.conf
```

IPsec 연결을 시작하려면 각 라우터에서 다음 명령을 실행합니다:

```
ifup ipsec0
```

LAN A와 LAN B가 서로 통신할 수 있도록 연결이 활성화됩니다. IPsec 연결에서 ifup을 실행하여 호출된 초기화 스크립트에 의해 라우트가 자동으로 생성됩니다. 네트워크 상 라우트 목록을 보시려면, 다음 명령을 실행합니다:

```
ip route list
```

To test the IPsec connection, run the tcpdump utility on the externally-routable device (eth0 in this example) to view the network packets being transferred between the hosts (or networks), and verify that they are encrypted via IPsec. For example, to check the IPsec connectivity of LAN A, use the following command:

```
tcpdump -n -i eth0 host lana.example.com
```

패킷은 AH 헤더를 포함해야 하며 ESP 패킷으로 보여져야 하니다. ESP는 패킷이 암호화되었다는 것을 의미합니다. 예를 들면 (백슬래쉬는 한 줄이 계속 된다는 것을 의미합니다):

```
12:24:26.155529 lanb.example.com > lana.example.com: AH(spi=0x021c9834,seq=0x358): \
 lanb.example.com > lana.example.com: ESP(spi=0x00c887ad,seq=0x358) (DF) \
 (ipip-proto-4)
```

46.7.8. IPsec 연결 시작하기 및 중지하기

부팅시 IPsec 연결이 활성화되도록 설정되지 않은 경우, 명령행에서 이를 제어하실 수 있습니다.

연결을 시작하려면 호스트 간 IPsec의 각각의 호스트에서나 또는 네트워크 간 IPsec의 각각의 IPsec 라운터에서 다음 명령을 실행합니다:

```
ifup <nickname>
```

where <nickname> is the nickname configured earlier, such as ipsec0.

다음 명령을 사용하여 명령을 중지합니다:

```
ifdown <nickname>
```

46.8. Firewalls

Information security is commonly thought of as a process and not a product. However, standard security implementations usually employ some form of dedicated mechanism to control access privileges and restrict network resources to users who are authorized, identifiable, and traceable. Red Hat Enterprise Linux includes several tools to assist administrators and security engineers with network-level access control issues.

Firewalls are one of the core components of a network security implementation. Several vendors market firewall solutions catering to all levels of the marketplace: from home users protecting one PC to data center solutions safeguarding vital enterprise information. Firewalls can be stand-alone hardware solutions, such as firewall appliances by Cisco, Nokia, and Sonicwall. Vendors such as

Checkpoint, McAfee, and Symantec have also developed proprietary software firewall solutions for home and business markets.

Apart from the differences between hardware and software firewalls, there are also differences in the way firewalls function that separate one solution from another. 표 46.5. "방화벽 유형" details three common types of firewalls and how they function:

표 46.5. 방화벽 유형

방법	설명	장점	단점
NAT	Network Address Translation (NAT) places private IP subnetworks behind one or a small pool of public IP addresses, masquerading all requests to one source rather than several. The Linux kernel has built-in NAT functionality through the Netfilter kernel subsystem.	· Can be configured transparently to machines on a LAN · Protection of many machines and services behind one or more external IP addresses simplifies administration duties · Restriction of user access to and from the LAN can be configured by opening and closing ports on the NAT firewall/gateway	· Cannot prevent malicious activity once users connect to a service outside of the firewall
패킷 필터	A packet filtering firewall reads each data packet that passes through a LAN. It can read and process packets by header information and filters the packet based on sets of programmable rules implemented by the firewall administrator. The Linux kernel has built-in packet filtering functionality through the Netfilter kernel subsystem.	· Customizable through the iptables front-end utility · Does not require any customization on the client side, as all network activity is filtered at the router level rather than the application level · Since packets are not transmitted through a proxy, network performance is faster due to direct connection from client to remote host	· Cannot filter packets for content like proxy firewalls · Processes packets at the protocol layer, but cannot filter packets at an application layer · Complex network architectures can make establishing packet filtering rules difficult, especially if coupled with IP masquerading or local subnets and DMZ networks
프록시	프록시 방화벽은 LAN 클라이언트로부터 프록시 기계로 들어오는 모든 요청에서 특정 프로토콜이나 유형을 걸러낸 후 이러한 유형을 로컬 클라이언트를 대신하여 인터넷에 보냅니다. 프록시 기계는 악의를 가진 원격 사용자와 내부 네트워크 클라이언트 기계 사이에서 버퍼로 작동합니다.	· Gives administrators control over what applications and protocols function outside of the LAN · Some proxy servers can cache frequently-accessed data locally rather than having to use the Internet connection to request it. This helps to reduce bandwidth consumption · Proxy services can be logged and monitored closely, allowing tighter	· Proxies are often application-specific (HTTP, Telnet, etc.), or protocol-restricted (most proxies work with TCP-connected services only) · Application services cannot run behind a proxy, so your application servers must use a separate form of network security · Proxies can become a network bottleneck, as all requests and transmissions are passed through one

방법	설명	장점	단점
		control over resource utilization on the network	source rather than directly from a client to a remote service

46.8.1. Netfilter and IPTables

The Linux kernel features a powerful networking subsystem called Netfilter. The Netfilter subsystem provides stateful or stateless packet filtering as well as NAT and IP masquerading services. Netfilter also has the ability to mangle IP header information for advanced routing and connection state management. Netfilter is controlled using the iptables tool.

46.8.1.1. IPTables Overview

The power and flexibility of Netfilter is implemented using the iptables administration tool, a command line tool similar in syntax to its predecessor, ipchains.

A similar syntax does not mean similar implementation, however. ipchains requires intricate rule sets for: filtering source paths; filtering destination paths; and filtering both source and destination connection ports.

By contrast, iptables uses the Netfilter subsystem to enhance network connection, inspection, and processing. iptables features advanced logging, pre- and post-routing actions, network address translation, and port forwarding, all in one command line interface.

This section provides an overview of iptables. For more detailed information, refer to 46.9절. “IPTables” .

46.8.2. Basic Firewall Configuration

Just as a firewall in a building attempts to prevent a fire from spreading, a computer firewall attempts to prevent malicious software from spreading to your computer. It also helps to prevent unauthorized users from accessing your computer.

In a default Red Hat Enterprise Linux installation, a firewall exists between your computer or network and any untrusted networks, for example the Internet. It determines which services on your computer remote users can access. A properly configured firewall can greatly increase the security of your system. It is recommended that you configure a firewall for any Red Hat Enterprise Linux system with an Internet connection.

46.8.2.1. Security Level Configuration Tool

During the Firewall Configuration screen of the Red Hat Enterprise Linux installation, you were given the option to enable a basic firewall as well as to allow specific devices, incoming services, and ports.

After installation, you can change this preference by using the Security Level Configuration Tool.

To start this application, use the following command:

```
system-config-securitylevel
```

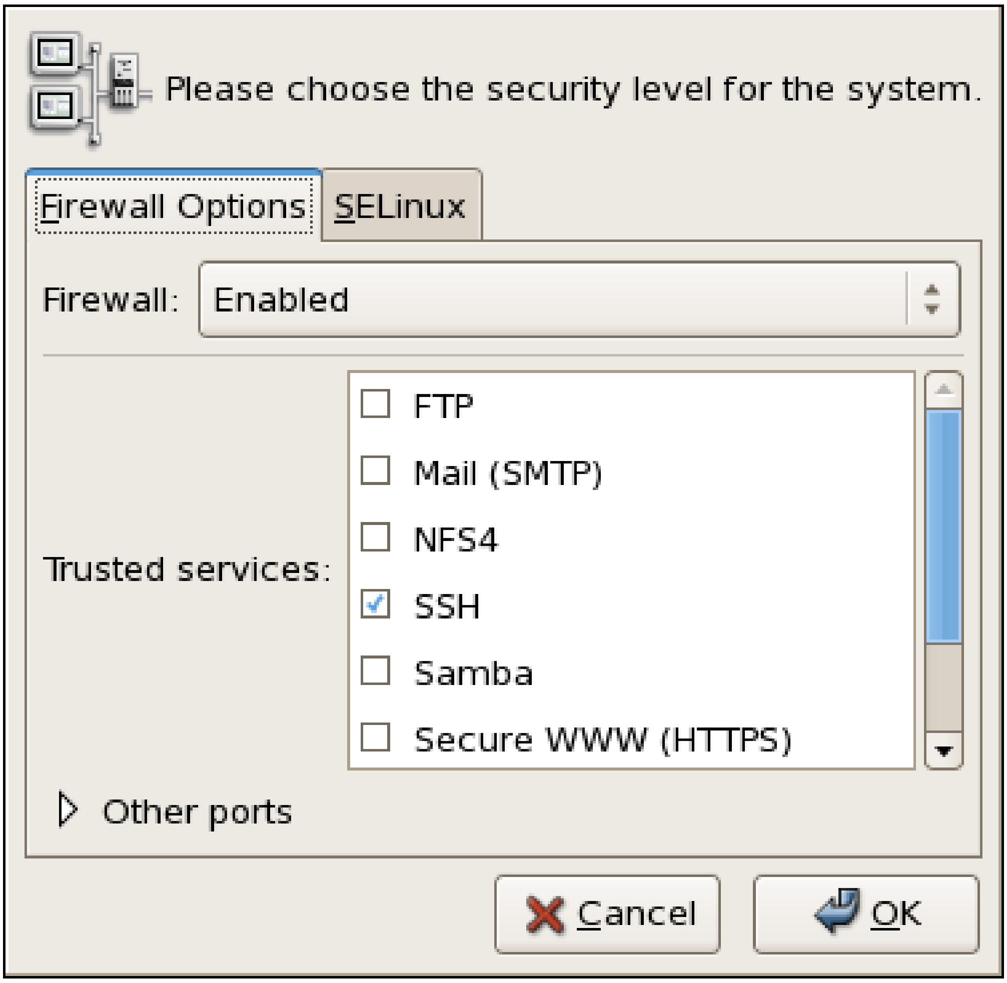

Please choose the security level for the system.

Firewall: Enabled

Trusted services:

□ FTP
□ Mail (SMTP)
□ NFS4
☑ SSH
□ Samba
□ Secure WWW (HTTPS)

그림 46.15. Security Level Configuration Tool

알림

The Security Level Configuration Tool only configures a basic firewall. If the system needs more complex rules, refer to 46.9절. "IPTables" for details on configuring specific iptables rules.

46.8.2.2. Enabling and Disabling the Firewall

Select one of the following options for the firewall:

- Disabled — Disabling the firewall provides complete access to your system and does no security checking. This should only be selected if you are running on a trusted network (not the Internet) or need to configure a custom firewall using the iptables command line tool.

경고

Firewall configurations and any customized firewall rules are stored in the /etc/sysconfig/iptables file. If you choose Disabled and click OK, these configurations and firewall rules will be lost.

- Enabled — This option configures the system to reject incoming connections that are not in response to outbound requests, such as DNS replies or DHCP requests. If access to services running on this machine is needed, you can choose to allow specific services through the firewall.

  If you are connecting your system to the Internet, but do not plan to run a server, this is the safest choice.

46.8.2.3. Trusted Services

Enabling options in the Trusted services list allows the specified service to pass through the firewall.

WWW (HTTP)

The HTTP protocol is used by Apache (and by other Web servers) to serve web pages. If you plan on making your Web server publicly available, select this check box. This option is not required for viewing pages locally or for developing web pages. This service requires that the httpd package be installed.

Enabling WWW (HTTP) will not open a port for HTTPS, the SSL version of HTTP. If this service is required, select the Secure WWW (HTTPS) check box.

FTP

The FTP protocol is used to transfer files between machines on a network. If you plan on making your FTP server publicly available, select this check box. This service requires that the vsftpd package be installed.

SSH

Secure Shell (SSH) is a suite of tools for logging into and executing commands on a remote machine. To allow remote access to the machine via ssh, select this check box. This service requires that the openssh-server package be installed.

Telnet

Telnet is a protocol for logging into remote machines. Telnet communications are unencrypted and provide no security from network snooping. Allowing incoming Telnet access is not recommended. To allow remote access to the machine via telnet, select this check box. This service requires that the telnet-server package be installed.

Mail (SMTP)

SMTP is a protocol that allows remote hosts to connect directly to your machine to deliver mail. You do not need to enable this service if you collect your mail from your ISP's server using POP3 or IMAP, or if you use a tool such as fetchmail. To allow delivery of mail to your machine, select this check box. Note that an improperly configured SMTP server can allow remote machines to use your server to send spam.

NFS4

The Network File System (NFS) is a file sharing protocol commonly used on *NIX systems. Version 4 of this protocol is more secure than its predecessors. If you want to share files or directories on your system with other network users, select this check box.

Samba

Samba is an implementation of Microsoft's proprietary SMB networking protocol. If you need to share files, directories, or locally-connected printers with Microsoft Windows machines, select this check box.

46.8.2.4. Other Ports

The Security Level Configuration Tool includes an Other ports section for specifying custom IP ports as being trusted by iptables. For example, to allow IRC and Internet printing protocol (IPP) to pass through the firewall, add the following to the Other ports section:

194:tcp,631:tcp

46.8.2.5. Saving the Settings

Click OK to save the changes and enable or disable the firewall. If Enable firewall was selected, the options selected are translated to iptables commands and written to the /etc/sysconfig/iptables file. The iptables service is also started so that the firewall is activated immediately after saving the selected options. If Disable firewall was selected, the /etc/sysconfig/iptables file is removed and the iptables service is stopped immediately.

The selected options are also written to the /etc/sysconfig/system-config-securitylevel file so that the settings can be restored the next time the application is started. Do not edit this file by hand.

Even though the firewall is activated immediately, the iptables service is not configured to start automatically at boot time. Refer to 46.8.2.6절. “Activating the IPTables Service” for more information.

46.8.2.6. Activating the IPTables Service

The firewall rules are only active if the iptables service is running. To manually start the service, use the following command:

```
service iptables restart
```

To ensure that iptables starts when the system is booted, use the following command:

```
chkconfig --level 345 iptables on
```

The ipchains service is not included in Red Hat Enterprise Linux. However, if ipchains is installed (for example, an upgrade was performed and the system had ipchains previously installed), the ipchains and iptables services should not be activated simultaneously. To make sure the ipchains service is disabled and configured not to start at boot time, use the following two commands:

```
service ipchains stop
```

```
chkconfig --level 345 ipchains off
```

46.8.3. Using IPTables

The first step in using iptables is to start the iptables service. Use the following command to start the iptables service:

```
service iptables start
```

> **알림**
>
> The ip6tables service can be turned off if you intend to use the iptables service only. If you deactivate the ip6tables service, remember to deactivate the IPv6 network also. Never leave a network device active without the matching firewall.

To force iptables to start by default when the system is booted, use the following command:

```
chkconfig --level 345 iptables on
```

This forces iptables to start whenever the system is booted into runlevel 3, 4, or 5.

46.8.3.1. IPTables Command Syntax

The following sample iptables command illustrates the basic command syntax:

```
iptables -A <chain> -j <target>
```

The -A option specifies that the rule be appended to <chain>. Each chain is comprised of one or more rules, and is therefore also known as a ruleset.

The three built-in chains are INPUT, OUTPUT, and FORWARD. These chains are permanent and cannot be deleted. The chain specifies the point at which a packet is manipulated.

The -j <target> option specifies the target of the rule; i.e., what to do if the packet matches the rule. Examples of built-in targets are ACCEPT, DROP, and REJECT.

Refer to the iptables man page for more information on the available chains, options, and targets.

46.8.3.2. 기본 방화벽 정책

Establishing basic firewall policies creates a foundation for building more detailed, user-defined rules.

Each iptables chain is comprised of a default policy, and zero or more rules which work in concert with the default policy to define the overall ruleset for the firewall.

The default policy for a chain can be either DROP or ACCEPT. Security-minded administrators typically implement a default policy of DROP, and only allow specific packets on a case-by-case basis. For example, the following policies block all incoming and outgoing packets on a network gateway:

```
iptables -P INPUT DROP
iptables -P OUTPUT DROP
```

It is also recommended that any forwarded packets — network traffic that is to be routed from the firewall to its destination node — be denied as well, to restrict internal clients from inadvertent exposure to the Internet. To do this, use the following rule:

```
iptables -P FORWARD DROP
```

When you have established the default policies for each chain, you can create and save further rules for your particular network and security requirements.

The following sections describe how to save iptables rules and outline some of the rules you might implement in the course of building your iptables firewall.

46.8.3.3. Saving and Restoring IPTables Rules

Changes to iptables are transitory; if the system is rebooted or if the iptables service is restarted, the rules are automatically flushed and reset. To save the rules so that they are loaded when the iptables service is started, use the following command:

```
service iptables save
```

The rules are stored in the file /etc/sysconfig/iptables and are applied whenever the service is started or the machine is rebooted.

46.8.4. Common IPTables Filtering

Preventing remote attackers from accessing a LAN is one of the most important aspects of network security. The integrity of a LAN should be protected from malicious remote users through the use of stringent firewall rules.

However, with a default policy set to block all incoming, outgoing, and forwarded packets, it is impossible for the firewall/gateway and internal LAN users to communicate with each other or with external resources.

To allow users to perform network-related functions and to use networking applications, administrators must open certain ports for communication.

예를 들어 방화벽에서 포트 80로의 액세스를 허용하시려면 다음 규칙을 추가하십시오:

```
iptables -A INPUT -p tcp -m tcp --dport 80 -j ACCEPT
```

This allows users to browse websites that communicate using the standard port 80. To allow access to secure websites (for example, https://www.example.com/), you also need to provide access to port 443, as follows:

```
iptables -A INPUT -p tcp -m tcp --dport 443 -j ACCEPT
```

중요

When creating an iptables ruleset, order is important.

If a rule specifies that any packets from the 192.168.100.0/24 subnet be dropped, and this is followed by a rule that allows packets from 192.168.100.13 (which is within the dropped subnet), then the second rule is ignored.

The rule to allow packets from 192.168.100.13 must precede the rule that drops the remainder of the subnet.

To insert a rule in a specific location in an existing chain, use the -I option. For example:

```
iptables -I INPUT 1 -i lo -p all -j ACCEPT
```

This rule is inserted as the first rule in the INPUT chain to allow local loopback device traffic.

There may be times when you require remote access to the LAN. Secure services, for example SSH, can be used for encrypted remote connection to LAN services.

Administrators with PPP-based resources (such as modem banks or bulk ISP accounts), dial-up access can be used to securely circumvent firewall barriers. Because they are direct connections, modem connections are typically behind a firewall/gateway.

For remote users with broadband connections, however, special cases can be made. You can configure iptables to accept connections from remote SSH clients. For example, the following rules allow remote SSH access:

```
iptables -A INPUT -p tcp --dport 22 -j ACCEPT
iptables -A OUTPUT -p tcp --sport 22 -j ACCEPT
```

These rules allow incoming and outbound access for an individual system, such as a single PC directly connected to the Internet or a firewall/gateway. However, they do not allow nodes behind the firewall/gateway to access these services. To allow LAN access to these services, you can use Network Address Translation (NAT) with iptables filtering rules.

46.8.5. FORWARD and NAT Rules

Most ISPs provide only a limited number of publicly routable IP addresses to the organizations they serve.

Administrators must, therefore, find alternative ways to share access to Internet services without giving public IP addresses to every node on the LAN. Using private IP addresses is the most common way of allowing all nodes on a LAN to properly access internal and external network services.

Edge routers (such as firewalls) can receive incoming transmissions from the Internet and route the packets to the intended LAN node. At the same time, firewalls/gateways can also route outgoing requests from a LAN node to the remote Internet service.

This forwarding of network traffic can become dangerous at times, especially with the availability of modern cracking tools that can spoof internal IP addresses and make the remote attacker's machine act as a node on your LAN.

To prevent this, iptables provides routing and forwarding policies that can be implemented to prevent abnormal usage of network resources.

The FORWARD chain allows an administrator to control where packets can be routed within a LAN. For example, to allow forwarding for the entire LAN (assuming the firewall/gateway is assigned an internal IP address on eth1), use the following rules:

```
iptables -A FORWARD -i eth1 -j ACCEPT
iptables -A FORWARD -o eth1 -j ACCEPT
```

This rule gives systems behind the firewall/gateway access to the internal network. The gateway routes packets from one LAN node to its intended destination node, passing all packets through its eth1 device.

알림

By default, the IPv4 policy in Red Hat Enterprise Linux kernels disables support for IP forwarding. This prevents machines that run Red Hat Enterprise Linux from functioning as dedicated edge routers. To enable IP forwarding, use the following command:

```
sysctl -w net.ipv4.ip_forward=1
```

This configuration change is only valid for the current session; it does not persist beyond a reboot or network service restart. To permanently set IP forwarding, edit the /etc/sysctl.conf file as follows:

Locate the following line:

```
net.ipv4.ip_forward = 0
```

Edit it to read as follows:

```
net.ipv4.ip_forward = 1
```

Use the following command to enable the change to the sysctl.conf file:

```
sysctl -p /etc/sysctl.conf
```

46.8.5.1. Postrouting and IP Masquerading

Accepting forwarded packets via the firewall's internal IP device allows LAN nodes to communicate with each other; however they still cannot communicate externally to the Internet.

To allow LAN nodes with private IP addresses to communicate with external public networks, configure the firewall for IP masquerading, which masks requests from LAN nodes with the IP address of the firewall's external device (in this case, eth0):

```
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

This rule uses the NAT packet matching table (-t nat) and specifies the built-in POSTROUTING chain for NAT (-A POSTROUTING) on the firewall's external networking device (-o eth0).

POSTROUTING allows packets to be altered as they are leaving the firewall's external device.

The -j MASQUERADE target is specified to mask the private IP address of a node with the external IP address of the firewall/gateway.

46.8.5.2. Prerouting

If you have a server on your internal network that you want make available externally, you can use the -j DNAT target of the PREROUTING chain in NAT to specify a destination IP address and port where incoming packets requesting a connection to your internal service can be forwarded.

For example, if you want to forward incoming HTTP requests to your dedicated Apache HTTP Server at 172.31.0.23, use the following command:

```
iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j DNAT --to 172.31.0.23:80
```

This rule specifies that the nat table use the built-in PREROUTING chain to forward incoming HTTP requests exclusively to the listed destination IP address of 172.31.0.23.

알림

If you have a default policy of DROP in your FORWARD chain, you must append a rule to forward all incoming HTTP requests so that destination NAT routing is possible. To do this, use the following command:

```
iptables -A FORWARD -i eth0 -p tcp --dport 80 -d 172.31.0.23 -j ACCEPT
```

This rule forwards all incoming HTTP requests from the firewall to the intended destination; the Apache HTTP Server behind the firewall.

46.8.5.3. DMZs and IPTables

You can create iptables rules to route traffic to certain machines, such as a dedicated HTTP or FTP server, in a demilitarized zone (DMZ). A DMZ is a special local subnetwork dedicated to providing services on a public carrier, such as the Internet.

For example, to set a rule for routing incoming HTTP requests to a dedicated HTTP server at 10.0.4.2 (outside of the 192.168.1.0/24 range of the LAN), NAT uses the PREROUTING table to forward the packets to the appropriate destination:

```
iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j DNAT --to-destination 10.0.4.2:80
```

With this command, all HTTP connections to port 80 from outside of the LAN are routed to the HTTP server on a network separate from the rest of the internal network. This form of network segmentation can prove safer than allowing HTTP connections to a machine on the network.

If the HTTP server is configured to accept secure connections, then port 443 must be forwarded as well.

46.8.6. Malicious Software and Spoofed IP Addresses

More elaborate rules can be created that control access to specific subnets, or even specific nodes, within a LAN. You can also restrict certain dubious applications or programs such as Trojans, worms, and other client/server viruses from contacting their server.

For example, some Trojans scan networks for services on ports from 31337 to 31340 (called the elite ports in cracking terminology).

Since there are no legitimate services that communicate via these non-standard ports, blocking them can effectively diminish the chances that potentially infected nodes on your network independently communicate with their remote master servers.

The following rules drop all TCP traffic that attempts to use port 31337:

```
iptables -A OUTPUT -o eth0 -p tcp --dport 31337 --sport 31337 -j DROP
iptables -A FORWARD -o eth0 -p tcp --dport 31337 --sport 31337 -j DROP
```

You can also block outside connections that attempt to spoof private IP address ranges to infiltrate your LAN.

For example, if your LAN uses the 192.168.1.0/24 range, you can design a rule that instructs the Internet-facing network device (for example, eth0) to drop any packets to that device with an address in your LAN IP range.

Because it is recommended to reject forwarded packets as a default policy, any other spoofed IP address to the external-facing device (eth0) is rejected automatically.

```
iptables -A FORWARD -s 192.168.1.0/24 -i eth0 -j DROP
```

알림

There is a distinction between the DROP and REJECT targets when dealing with appended rules.

The REJECT target denies access and returns a connection refused error to users who attempt to connect to the service. The DROP target, as the name implies, drops the packet without any warning.

Administrators can use their own discretion when using these targets. However, to avoid user confusion and attempts to continue connecting, the REJECT target is recommended.

46.8.7. IPTables and Connection Tracking

You can inspect and restrict connections to services based on their connection state. A module within iptables uses a method called connection tracking to store information about incoming connections. You can allow or deny access based on the following connection states:

- NEW — 새로운 연결을 요청하는 패킷, 예, HTTP 요청
- ESTABLISHED — 기존 연결의 일부인 패킷
- RELATED — A packet that is requesting a new connection but is part of an existing connection. For example, FTP uses port 21 to establish a connection, but data is transferred on a different port (typically port 20).
- INVALID — 연결 추적표에서 어디 연결에도 속하지 않은 패킷.

You can use the stateful functionality of iptables connection tracking with any network protocol, even if the protocol itself is stateless (such as UDP). The following example shows a rule that uses connection tracking to forward only the packets that are associated with an established connection:

```
iptables -A FORWARD -m state --state ESTABLISHED,RELATED -j ACCEPT
```

46.8.8. IPv6

The introduction of the next-generation Internet Protocol, called IPv6, expands beyond the 32-bit address limit of IPv4 (or IP). IPv6 supports 128-bit addresses, and carrier networks that are IPv6 aware are therefore able to address a larger number of routable addresses than IPv4.

Red Hat Enterprise Linux supports IPv6 firewall rules using the Netfilter 6 subsystem and the ip6tables command. In Red Hat Enterprise Linux 5, both IPv4 and IPv6 services are enabled by default.

The ip6tables command syntax is identical to iptables in every aspect except that it supports 128-bit addresses. For example, use the following command to enable SSH connections on an IPv6-aware network server:

```
ip6tables -A INPUT -i eth0 -p tcp -s 3ffe:ffff:100::1/128 --dport 22 -j ACCEPT
```

IPv6 네트워킹에 대한 자세한 정보를 원하신다면 http://www.ipv6.org/에서 IPv6 정보 페이지를 참조하시기 바랍니다.

46.8.9. 추가 자료

이 메뉴얼에서는 방화벽과 리눅스 넷필터 서브시스템에 대하여 깊게 다양한 측면을 다루지 못했습니다. 보다 자세한 정보를 보시려면 다음에 나온 자료를 참조하시기 바랍니다.

46.8.9.1. 설치된 문서 자료

- Refer to 46.9절. "IPTables" for more detailed information on the iptables command, including definitions for many command options.
- The iptables man page contains a brief summary of the various options.

46.8.9.2. 유용한 웹사이트

- http://www.netfilter.org/ — The official homepage of the Netfilter and iptables project.
- http://www.tldp.org/ — 리눅스 문서화 프로젝트 사이트에서 방화벽 생성과 관리와 관련된 유용한 가이드를 찾으실 수 있습니다.

- http://www.iana.org/assignments/port-numbers — IANA (인터넷 할당 번호 관리 기관)에 의해 할당된 등록된 일반 서비스 포트 번호의 공식 목록.

46.8.9.3. 관련 서적

- Red Hat Linux Firewalls, by Bill McCarty; Red Hat Press — a comprehensive reference to building network and server firewalls using open source packet filtering technology such as Netfilter and iptables. It includes topics that cover analyzing firewall logs, developing firewall rules, and customizing your firewall using various graphical tools.

- Linux Firewalls, by Robert Ziegler; New Riders Press — contains a wealth of information on building firewalls using both 2.2 kernel ipchains as well as Netfilter and iptables. Additional security topics such as remote access issues and intrusion detection systems are also covered.

46.9. IPTables

Red Hat Enterprise Linux와 함께 네트워크 패킷 필터링 (packet filtering)에 대한 고급 도구 모음이 포함되어 있습니다 — 커널안에서 패킷을 입력, 이동, 네트워크 스텍을 종료하는 것과 같이 네트워크 패킷을 제어하는 순서. 커널 2.4 이전 버전은 패킷 필터링에 대해 ipchains에 의존하며 필터링 과정의 각각의 단계에서 패킷에 적용되는 규칙 목록을 사용합니다. 커널 2.4버전에서는 iptables (넷필터(netfilter)라고도 부름)를 소개하고 있으며, 이는 ipchains와 비슷하지만 범위을 확장하여 네트워크 패킷 필터링을 할 수 있는 제어 기능을 갖고 있습니다.

이 장에서는 기본적인 패킷 필터링을 중점적으로 다루고, ipchains와 iptables의 다른점에 대해 알아보며, iptables 명령을 가지고 사용 가능한 다양한 옵션과 시스템 재부팅에서 필터링 규칙이 어떻게 보존되는지에 대해 설명합니다.

Refer to 46.9.7절. "추가 자료" for instructions on how to construct iptables rules and setting up a firewall based on these rules.

경고

커널 2.4 버전과 그 이후 버전에서 기본 방화벽 메카니즘은 iptables이지만, ipchains이 이미 실행되고 있을 경우, iptables는 사용될 수 없습니다. ipchains가 부팅할 때에 나타나면, 커널은 오류를 발생하며 iptables를 시작할 수 없게 됩니다.

ipchains의 기능은 이러한 오류에 의해 영향을 받지 않습니다.

46.9.1. 패킷 필터링 (Packet Filtering)

Linux 커널은 패킷을 필터하기 위해 Netfilter 기능을 사용하며, 다른 패킷이 정지해 있는 동안 패킷의 일부분이 시스템을 통과하거나 전달받는 것을 허용합니다. 이러한 기능은 Linux 커널에 내장되어 있으며, 다음과 같이 세가지의 내장된 테이블 또는 규칙 목록을 가지고 있습니다:

- filter — 네트워크 패킷을 다루기 위한 기본 테이블.

- nat — 새로운 연결을 생성하는 패킷을 변경하기 위해 사용되며 네트워크 주소 변환 (Network Address Translation) (NAT)을 위해 사용됨.

- mangle — 특정 유형의 패킷 변경을 위해 사용됨.

각각의 테이블에는 내장된 chains 그룹이 있으며, 이는 netfilter에 의한 패킷에서 실행되는 작업에 해당합니다.

filter 테이블에 대한 내장된 chains은 다음과 같습니다:

- INPUT — 호스트를 대상으로 하는 네트워크 패킷에 적용합니다.
- OUTPUT — 로컬에서 생성된 네트워크 패킷에 적용합니다.
- FORWARD — 호스트를 통해 라우트된 네트워크 패킷에 적용합니다.

nat 테이블에 해당하는 내장된 chains은 다음과 같습니다:

- PREROUTING — 네트워크 패킷이 도착하면 이를 변경합니다.
- OUTPUT — 패킷이 보내어지기 전에 로컬에서 생성된 네트워크 패킷을 변경합니다.
- POSTROUTING — 패킷이 보내어지기 전에 네트워크 패킷을 변경합니다.

mangle 테이블에 해당하는 내장된 chains은 다음과 같습니다:

- INPUT — 호스트를 대상으로 하는 네트워크 패킷을 변경합니다.
- OUTPUT — 패킷이 보내어지기 전에 로컬에서 생성된 네트워크 패킷을 변경합니다.
- FORWARD — 호스트를 통하여 라우트된 네트워크 패킷을 변경합니다.
- PREROUTING — 패킷이 라우트되기 전에 들어오는 네트워크 패킷을 변경합니다.
- POSTROUTING — 패킷이 보내어지기 전에 네트워크 패킷을 변경합니다.

Linux 시스템에서 받거나 보내진 모든 네트워크 패킷은 최소 하나의 테이블에 적용되지만, 패킷은 chain의 마지막에 나타나기 전에 각각의 테이블안에 있는 여러 규칙에 적용될 수 도 있습니다. 이러한 규칙의 구조와 목적은 다양하나, 이는 주로 특정 IP 주소에서 들어오거나 IP 주소로 나가는 패킷을 확인하거나 또는 특정 프로토콜과 네트워크 서비스를 사용할 때 주소를 설정합니다.

알림

기본값으로 방화벽 규칙이 /etc/sysconfig/iptables 또는 /etc/sysconfig/ip6tables 파일에 저장되었습니다.

iptables 서비스는 Linux 시스템이 부팅할 때 DNS 관련 서비스 이전에 시작합니다. 이는 방화벽 규칙이 숫자로된 IP 주소 (예:192.168.0.1)만을 참조할 수 있다는 것을 의미합니다. 도메인명 (예: host.example.com)은 이러한 규칙에서 오류를 발생합니다.

Regardless of their destination, when packets match a particular rule in one of the tables, a target or action is applied to them. If the rule specifies an ACCEPT target for a matching packet, the packet skips the rest of the rule checks and is allowed to continue to its destination. If a rule specifies a DROP target, that packet is refused access to the system and nothing is sent back to the host that sent the packet. If a rule specifies a QUEUE target, the packet is passed to user-space. If a rule specifies the optional REJECT target, the packet is dropped, but an error packet is sent to the packet's originator.

모든 chain은 ACCEPT, DROP, REJECT, 또는 QUEUE에 기본 정책을 가지고 있습니다. chain에 있는 규칙 중 어느 것도 패킷에 적용되지 않을 경우, 패킷은 기본 정책에 따라 다루어 집니다.

iptables 명령은 이러한 테이블을 설정하며, 필요한 경우 새로운 테이블을 설정하기도 합니다.

46.9.2. IPTables과 IPChains의 다른점

ipchains과 iptables는 특정 규칙이나 규칙 모음에 일치하는 가에 따라 패킷을 필터하기 위해 Linux 커널안에서 작동하는 규칙의 chains을 사용합니다. 하지만, iptables는 관리자가 시스템에 지나친 복잡성을 구축하지 않고 제어하게 하는 패킷을 필터링하는 확장가능한 방법을 제공합니다.

다음과 같이 ipchains와 iptables의 다른점을 확인해 보시기 바랍니다:

Using iptables, each filtered packet is processed using rules from only one chain rather than multiple chains.

예를 들어 ipchains를 사용한 시스템으로 들어오는 FORWARD 패킷은 수신지에 도달하기 위해 INPUT, FORWARD, OUTPUT chain을 통과하게 됩니다. 하지만, iptables는 패킷의 수신지가 로컬 시스템일 경우 패킷을 INPUT chain으로만 보내며, 로컬 시스템이 패킷을 생성했을 경우 패킷을 OUTPUT chain으로만 보냅니다. 따라서 실제적으로 패킷을 다루는 chain 안에서 특정 패킷을 감지하도록 규칙을 고안해야 합니다.

The DENY target has been changed to DROP.

ipchains에서, chain에 있는 규칙과 일치하는 패킷은 거부 (DENY) 대상이 됩니다. 이러한 대상은 iptables에서 취소 (DROP) 대상으로 변경되어야 합니다.

Order matters when placing options in a rule.

ipchains에서 규칙 옵션의 순서는 중요하지 않습니다.

iptables 명령은 엄격한 구문을 가집니다. iptables 명령은 발생지 포트 또는 수신지 포트 이전에 프로토콜 (ICMP, TCP, 또는 UDP)을 지정할 것을 요구합니다.

Network interfaces must be associated with the correct chains in firewall rules.

예를들어, 들어오는 인터페이스는 (-i 옵션) INPUT 또는 FORWARD chain에서만 사용될 수 있습니다. 이와 비슷하게, 나가는 인터페이스는 (-o 옵션) FORWARD 또는 OUTPUT chain에서만 사용될 수 있습니다.

다시 말해, INPUT chain과 들어오는 인터페이스는 함께 작동하며, OUTPUT chain과 나가는 인터페이스는 함께 작동합니다. FORWARD chain은 들어오고 나가는 인터페이스 모두와 함께 작동합니다.

OUTPUT chain은 더이상 들어오는 인터페이스에 의해 사용되지 않으며, INPUT chain은 나가는 인터페이스를 통해 이동하는 패킷에 의해 보여지지 않습니다.

This is not a comprehensive list of the changes. Refer to 46.9.7절. "추가 자료" for more specific information.

46.9.3. IPTables에 대한 명령 옵션

패킷을 필터링하기 위한 규칙은 iptables 명령을 사용하여 생성됩니다. 주로 다음과 같은 패킷 내용이 기준으로서 사용됩니다.

- 패킷 유형 — 명령을 필터하는 패킷의 유형을 지정합니다.
- 패킷 발생지/수신지 — 패킷의 발생지 또는 수신지에 기반하여 어떤 패킷이 명령을 필터할 지를 지정합니다.
- 대상 — 위의 기준과 일치하여 어떤 작업을 패킷에 실행할 지를 지정합니다.

Refer to 46.9.3.4절. "IPTables 일치 옵션" and 46.9.3.5절. "대상 옵션" for more information about specific options that address these aspects of a packet.

특정 iptables 규칙과 함께 사용되는 옵션은 사용 가능한 규칙에 대한 모든 규칙의 목적 및 조건에 기반하여 논리적으로 그룹지어져야만 합니다. 이 장의 나머지 부분에서는 iptables 명령에 해당하는 일반적으로 사용되는 옵션에 대해 설명합니다.

46.9.3.1. IPTables 명령 옵션의 구조

여러 iptables 명령은 다음과 같은 구조를 가지고 있습니다:

```
iptables [-t <table-name>] <command> <chain-name> \
    <parameter-1> <option-1> \
    <parameter-n> <option-n>
```

<table-name> — Specifies which table the rule applies to. If omitted, the filter table is used.

<command> — Specifies the action to perform, such as appending or deleting a rule.

<chain-name> — Specifies the chain to edit, create, or delete.

<parameter>-<option> pairs — Parameters and associated options that specify how to process a packet that matches the rule.

iptables 명령의 길이 및 복잡성은 목적에 따라 현저하게 변경될 수 있습니다.

예를 들어, chain에서 규칙을 삭제하기 위한 명령의 길이가 짧을 수 도 있습니다:

iptables -D <chain-name> <line-number>

반면, 여러 특정 매개변수와 옵션을 사용하는 특정 서브넷에서 패킷을 필터링하는 규칙을 추가하는 명령은 다소 길어질 수 있습니다. iptables 명령을 구축할 때, 몇몇 매개 변수와 옵션은 유효한 규칙을 구축하기 위해 추가 매개 변수 및 옵션을 요구한다는 점을 기억해두시기 바랍니다. 이는 더 많은 매개변수를 요구하는 추가 매개 변수와 함께 캐스캐이딩(cascading) 효과를 창출할 수 있습니다. 다른 옵션을 요구하는 모든 매개변수 및 옵션을 만족시킬때 까지, 규칙은 유효하지 않게 됩니다.

iptables -h를 입력하여 iptables 명령 구조의 종합적인 목록을 봅니다.

46.9.3.2. 명령 옵션

명령 옵션은 특정 작업을 실행하기 위해 iptables를 사용합니다. 하나의 명령 옵션에 하나의 iptables 명령만이 허용됩니다. 도움말 명령을 제외하고 모든 명령은 대문자로 쓰여져 있습니다.

iptables 명령은 다음과 같습니다:

- -A — 특정 chain의 마지막에 규칙을 추가합니다. 아래에 설명된 -I 옵션과는 달리, 이는 정수 인자를 갖지 않으며, 항상 특정 chain의 마지막에 규칙을 추가합니다.

- -C — 사용자 지정 chain에 특정 규칙을 추가하기 전에 이를 확인합니다. 이러한 명령은 추가 매개 변수와 옵션을 요구하여 복잡한 iptables 규칙을 구성하게 할 수 있습니다.

- -D <integer> | <rule> — Deletes a rule in a particular chain by number (such as 5 for the fifth rule in a chain), or by rule specification. The rule specification must exactly match an existing rule.

- -E — 사용자 정의된 chain의 이름을 변경합니다. 사용자 정의된 chain은 기본값이 아닌 기존의 모든 chain입니다. (사용자 정의된 chain의 생성에 관한 자세한 정보는 아래의 -N 옵션을 참조하시기 바랍니다.) 이는 표면적인 변경 사항으로 테이블의 구조에 아무런 영향을 미치지 않습니다.

  알림

  기본 chain 중 하나의 이름을 변경하시려 할 경우, 시스템은 Match not found 오류를 보고하여, 기본 chain의 이름을 변경하실 수 없게 됩니다.

- -F — 선택한 chain을 삭제합니다, 이는 실제적으로 chain에 있는 모든 규칙을 삭제하게 됩니다. 어떤 chain도 지정되어 있지 않을 경우,이 명령은 모든 chain에 있는 모든 규칙을 삭제합니다.

- -h — 명령 구조 목록과 명령 매개 변수 및 옵션의 요약 설명을 제공합니다.

- -I [<integer>] — Inserts the rule in the specified chain at a point specified by a user-defined integer argument. If no argument is specified, the rule is inserted at the top of the chain.

  주의

  위에서 언급된것 처럼, chain에 있는 규칙의 순서는 어떤 규칙이 어떤 패킷에 적용되는냐에 따라 결정됩니다. 이는 -A 또는 -I 옵션을 사용하여 규칙을 추가할 때 중요합니다.

  이는 특히 정수 인자와 함께 -I 옵션을 사용하여 규칙을 추가할 때 중요합니다. chain에 규칙을 추가할 때 기존의 숫자를 지정하신 경우, iptables 명령은 기존의 규칙 이전에 (또는 위에) 새로운 규칙을 추가합니다.

- -L — 명령 이후에 지정된 chain에 있는 모든 규칙의 목록을 만듭니다. 기본 filter 테이블의 모든 chain에 있는 모든 규칙의 목록을 만들기 위해, chain이나 테이블을 지정하지 마십시오. 그렇지 않으면, 특정 테이블의 특정 chain에 있는 규칙의 목록을 만드는데 다음의 구문을 사용해야 합니다.

  ```
  iptables -L <chain-name> -t <table-name>
  ```

  Additional options for the -L command option, which provide rule numbers and allow more verbose rule descriptions, are described in 46.9.3.6절. “옵션 목록”.

- -N — 사용자 정의된 이름과 함께 새로운 chain을 생성합니다. 고유한 chain 명이 아닐 경우 오류 메세지가 나타납니다.

- -P — 특정 chain에 대해 기본 정책을 설정하여, 패킷이 규칙에 일치하지 않고 chain을 관통하게 되면, 허용 (ACCEPT) 또는 취소 (DROP)와 같은 특정 목표 규칙을 보내게 됩니다.

- -R — Replaces a rule in the specified chain. The rule's number must be specified after the chain's name. The first rule in a chain corresponds to rule number one.

- -X — 사용자 정의된 chain을 삭제합니다. 내장된 chain을 삭제하실 수 는 없습니다.

- -Z — 바이트를 설정하고 테이블에 대한 모든 chain에 있는 패킷 카운터를 0까지 설정합니다.

46.9.3.3. IPTables 매개 변수 옵션

특정 chain 안에서 규칙의 추가, 삭제, 삽입 또는 대체에 사용되는 특정 iptables 명령은 패킷 필터링 규칙을 구성하기 위해 다양한 매개 변수를 필요로 합니다.

- -c — 특정 규칙에 대해 카운터를 재설정합니다. 이러한 매개 변수는 어떤 카운터를 재설정할지를 지정하기 위해 PKTS 및 BYTES 옵션을 허용합니다.
- -d — 수신지의 호스트명, IP 주소, 또는 규칙에 일치하는 패킷의 네트워크를 설정합니다. 네트워크에 일치할 때, 다음의 IP 주소/넷마스크 형식이 지원됩니다:
  - N.N.N.N/M.M.M.M — N.N.N.N는 IP 주소 영역이며 M.M.M.M는 넷마스크에 해당합니다.
  - N.N.N.N/M — N.N.N.N는 IP 주소 영역이며 M는 비트마스크(bitmask)에 해당합니다.
- -f — 이러한 규칙은 분리된 패킷에만 적용됩니다.

  분리되지 않은 패킷에 일치하는 것을 지정하기 위해 이러한 매개 변수 뒤에 느낌표 기호 (!) 옵션을 사용합니다.

  > 알림
  >
  > 분리된 패킷이 IP 프로토콜의 표준이 될지라도 분리된 패킷과 분리되지 않은 패킷을 구별하는 것이 바람직합니다.
  >
  > Originally designed to allow IP packets to travel over networks with differing frame sizes, these days fragmentation is more commonly used to generate DoS attacks using mal-formed packets. It's also worth noting that IPv6 disallows fragmentation entirely.

- -i — eth0 또는 ppp0와 같이 들어오는 네트워크 인터페이스를 설정합니다. iptables 명령을 가지고, 이러한 매개 변수 옵션은 nat 및 mangle 테이블과 함께 filter 테이블과 PREROUTING chain을 사용할 때 INPUT과 FORWARD chain을 사용할 수 도 있습니다.

  이러한 매개 변수는 다음과 같은 특정 옵션도 지원합니다:

  - 느낌표 기호 (!) — 지시 사항을 바꾸어, 의미있는 특정 인터페이스를 이러한 규칙에서 제외합니다.
  - 플러스 기호 (+) — 특정 문자열에 일치하는 인터페이스를 일치시키기 위해 사용되는 와일드카드 문자. 예를 들어, 매개 변수 -i eth+는 이러한 규칙을 이더넷 인터페이스에 적용시킬 수 있지만, ppp0와 같은 다른 인터페이스는 여기서 제외됩니다.

  -i 매개 변수가 사용되었지만 어떤 인터페이스도 지정되지 않았을 경우, 모든 인터페이스는 이러한 규칙의 영향을 받게 됩니다.

- -j — 패킷이 특정 규칙에 일치할 때 지정된 대상으로 이동합니다.

  기준이 되는 대상은 ACCEPT, DROP, QUEUE, RETURN입니다.

확장된 옵션은 Red Hat Enterprise Linux iptables RPM 패키지를 사용하여 기본값으로 읽어진 모듈을 통해 사용하실 수 있습니다. 이러한 모듈에 있는 유효한 대상에는 LOG, MARK, REJECT, 등이 포함됩니다. 이러한 대상에 대한 보다 자세한 정보는 iptables 메뉴얼 페이지에서 참조하시기 바랍니다.

이러한 옵션은 특정 규칙을 현재 chain의 외부에 있는 사용자 정의된 chain에 일치하는 패킷에 지정하는 데 사용되며 그 결과 다른 규칙을 패킷에 적용시킬 수 있습니다.

어떤 대상도 지정되지 않았을 경우, 패킷은 아무런 작업을 실행하지 않고 이전 규칙으로 이동합니다. 하지만, 이러한 규칙에 대한 카운터는 한개로 증가합니다.

- -o — 규칙에 대하여 나가는 네트워크 인터페이스를 설정합니다. 이러한 옵션은 filter 테이블에 있는 OUTPUT 및 FORWARD chain과 nat 및 mangle 테이블에 있는 POSTROUTING chain에 대해서만 유효합니다. 이러한 매개 변수는 들어오는 네트워크 인터페이스 매개 변수 (-i)와 같은 옵션을 허용합니다.

- -p <protocol> — Sets the IP protocol affected by the rule. This can be either icmp, tcp, udp, or all, or it can be a numeric value, representing one of these or a different protocol. You can also use any protocols listed in the /etc/protocols file.

  The "all" protocol means the rule applies to every supported protocol. If no protocol is listed with this rule, it defaults to "all".

- -s — 수신지 (-d) 매개 변수와 같은 구문을 사용하여 특정 패킷에 대한 소스를 설정합니다.

46.9.3.4. IPTables 일치 옵션

Different network protocols provide specialized matching options which can be configured to match a particular packet using that protocol. However, the protocol must first be specified in the iptables command. For example, -p <protocol-name> enables options for the specified protocol. Note that you can also use the protocol ID, instead of the protocol name. Refer to the following examples, each of which have the same effect:

```
iptables -A INPUT -p icmp --icmp-type any -j ACCEPT
iptables -A INPUT -p 5813 --icmp-type any -j ACCEPT
```

서비스 정의는 /etc/services 파일에 있습니다. 가독성을 위해, 포트 번호를 사용하는 데신 서비스명을 사용하실 것을 권장합니다.

중요

허가없이 편집하는 것을 방지하기 위해 /etc/services 파일을 보호합니다. 이 파일을 편집하는 것이 가능하게 될 경우, 침입자는 사용자가 연결을 끊어야 하는 사용자의 컴퓨터에 있는 포트를 활성화할 수 있습니다. 이 파일의 보안을 위해 root 계정으로 다음의 명령을 입력하시기 바랍니다:

```
chown root.root /etc/services
chmod 0644 /etc/services
chattr +i /etc/services
```

파일의 이름을 변경하거나, 삭제 또는 링크로 연결하여 이를 방지할 수 있습니다.

46.9.3.4.1. TCP 프로토콜

이러한 일치 옵션은 TCP 프로토콜 (-p tcp)을 위해 사용가능합니다:

- --dport — 패킷을 위해 수신지 포트를 설정합니다.

  이러한 옵션을 설정하기 위해, www 또는 smtp와 같은 네트워크 서비스명; 포트 번호; 또는 포트 번호의 범위를 사용합니다.

  포트 번호의 범위를 지정하기 위해, 콜론을 사용하여 번호를 구별합니다(:). 예: -p tcp --dport 3000:3200. 사용할 수 있는 유효한 최대 범위는 0:65535입니다.

  네트워크 서비스 혹은 포트를 사용하지 않는 모든 패킷을 일치시키기 위해 --dport 옵션 뒤에 느낌표 기호 (!)를 사용합니다.

  네트워크 서비스명과 별칭 및 사용하는 포트 번호를 검색하기 위해, /etc/services 파일을 확인합니다.

  --destination-port 일치 옵션은 --dport와 의미가 같습니다.

- --sport — --dport와 같은 옵션을 사용하여 패킷의 소스 포트를 설정합니다. --source-port 일치 옵션은 --sport와 의미가 같습니다.

- --syn — 일반적으로 SYN 패킷이라고 불리우는, 통신을 초기화하기 위해 고안된 모든 TCP 패킷에 적용합니다. 데이터 페이로드(payload)를 전송하는 패킷에는 적용되지 않습니다.

  SYN이 아닌 모든 패킷에 일치시키기 위해 --syn 옵션 뒤에 느낌표 기호 (!)를 사용합니다.

- --tcp-flags <tested flag list> <set flag list> — Allows TCP packets that have specific bits (flags) set, to match a rule.

  --tcp-flags 일치 옵션은 두개의 매개 변수를 허용합니다. 첫 번째 매개 변수는 마스크로 패킷에서 검사하기 위해 콤마로 구분되는 플래그 목록입니다. 두번째 매개 변수는 규칙을 일치시키시 위해 설정되야 하는 콤마로 구분되는 플래그 목록입니다.

  가능한 플래그는 다음과 같습니다:

  - ACK
  - FIN

  - PSH
  - RST
  - SYN
  - URG
  - ALL
  - NONE

  예를 들어, 다음과 같은 명세 사항을 포함하고 있는 iptables 규칙은 SYN 플래그 모음을 갖고 있는 TCP 패킷에 일치하며 ACK 및 FIN 플래그는 설정되지 않았습니다:

  --tcp-flags ACK,FIN,SYN SYN

  일치 옵션의 효과를 바꾸기 위해 --tcp-flags 뒤에 느낌표 부호 (!)를 사용합니다.

- --tcp-option — 특정 패킷 안에서 설정할 수 있는 TCP-특정 옵션과 일치시키려 합니다. 이러한 일치 옵션은 느낌표 부호 (!)를 사용하여 바꿀수 도 있습니다.

46.9.3.4.2. UDP 프로토콜

이러한 일치 옵션은 UDP 프로토콜 (-p udp)에 대해 사용가능 합니다:

- --dport — 서비스명, 포트 번호, 또는 포트 번호의 영역을 사용하여, UDP 패킷의 수신지 포트를 지정합니다. --destination-port 일치 옵션은 --dport와 의미가 같습니다.
- --sport — 서비스명, 포트 번호, 또는 포트 번호의 영역을 사용하여, UDP 패킷의 소스 포트를 지정합니다. --source-port 일치 옵션은 --sport와 의미가 같습니다.

--dport 및 --sport 옵션으로 포트 번호의 영역을 지정하고 콜론 (:)을 사용하여 두 번호를 구별합니다. 예: -p tcp --dport 3000:3200. 허용된 유효한 최대 영역은 0:65535입니다.

46.9.3.4.3. ICMP 프로토콜

다음의 일치 옵션은 ICMP (Internet Control Message Protocol) (-p icmp)에 대해 사용 가능합니다:

- --icmp-type — 규칙과 일치하는 ICMP 유형의 이름이나 번호를 설정합니다. 사용가능한 ICMP 이름 목록은 iptables -p icmp -h 명령을 입력하여 검색하실 수 있습니다.

46.9.3.4.4. 추가 일치 옵션 모듈

추가 일치 옵션은 iptables 명령으로 불러온 모듈을 통해 사용하실 수 있습니다.

To use a match option module, load the module by name using the -m <module-name>, where <module-name> is the name of the module.

여러 모듈은 기본값으로 사용 가능합니다. 추가 기능을 제공하기 위해 모듈을 생성하실 수 있습니다.

다음은 가장 일반적으로 사용되는 모듈의 부분적인 목록입니다:

- limit 모듈 — 얼마나 많은 패킷을 특정 규칙에 일치하게 할지에 제한을 둡니다.

LOG 대상과 함께 사용할 때, limit 모듈은 여러 일치하는 패킷을 반복되는 메세지가 있는 시스템 로그로 채우거나 시스템 리소스를 사용하지 못하게 합니다.

Refer to 46.9.3.5절. “대상 옵션” for more information about the LOG target.

limit 모듈은 다음과 같은 옵션을 활성화합니다:

- --limit — Sets the maximum number of matches for a particular time period, specified as a <value>/<period> pair. For example, using --limit 5/hour allows five rule matches per hour.

  기간은 초, 분, 시, 일로 명시될 수 있습니다.

  번호 및 시간 편집기가 사용되지 않을 경우, 3/hour의 기본값이 사용될 것입니다.

- --limit-burst — 한번에 규칙에 일치하는 패킷의 수에 제한을 둡니다.

  이러한 옵션은 정수로 지정되었으며 --limit 옵션과 함께 사용하셔야 합니다.

  아무런 값이 지정되지 않은 경우, 기본값은 (5)로 간주됩니다.

- state 모듈 — 상태 일치를 활성화합니다.

  state 모듈은 다음과 같은 옵션을 활성화합니다:

  - --state — 패킷을 다음의 연결 상태와 일치시킵니다:

    - ESTABLISHED — 일치 패킷은 기존 연결에 있는 다른 패킷과 관련되어 있습니다. 클라이언트와 서버 사이의 연결을 유지하시기를 원하실 경우, 이를 허용하셔야 합니다.

    - INVALID — 일치 패킷은 알려진 연결에 묶어질 수 없습니다.

    - NEW — 일치 패킷은 새로운 연결을 생성하거나 또는 이전에 볼 수 없었던 양방향 연결의 일부분이 될 수 있습니다. 서비스에 새로운 연결을 허용하시려면 이를 허용하셔야 합니다.

    - RELATED — 일치 패킷은 기존 연결과 관련하여 새로운 연결을 시작합니다. 이에 대한 예로 FTP가 있으며, 이는 소통량 제어 (포트 21)를 위해 하나의 연결을 사용하며, 데이터 전송 (포트 20)과 분리된 연결을 사용합니다.

    이러한 연결 상태는 -m state --state INVALID,NEW와 같이 콤마로 구분하여 서로 함께 사용될 수 있습니다.

- mac 모듈 — 하드웨어 MAC 주소 일치를 활성화합니다.

  mac 모듈은 다음의 옵션을 활성화합니다:

  - --mac-source — 패킷을 보낸 네트워크 인터페이스 카드의 MAC 주소에 일치시킵니다. 규칙에서 MAC 주소를 제외시키기 위해, --mac-source 일치 옵션 뒤에 느낌표 기호 (!)를 입력합니다.

모듈에서 사용가능한 일치 옵션에 대한 보다 많은 정보는 iptables 메뉴얼 페이지에서 참조하시기 바랍니다.

46.9.3.5. 대상 옵션

패킷이 특정 규칙과 일치할 때, 규칙은 패킷을 알맞은 작업을 결정하게 하는 여러 다른 대상에 지정합니다. 각각의 chain은 기본 대상을 가지고 있으며, 이는 chain에서 어떤 규칙도 패킷에 일치하지 않거나 또는 패킷에 일치하는 어떤 규칙도 대상을 지정할 수 없는 경우입니다.

다음은 기준이 되는 대상입니다:

- <user-defined-chain> — A user-defined chain within the table. User-defined chain names must be unique. This target passes the packet to the specified chain.
- ACCEPT — 패킷의 수신지 또는 다른 chain을 통과하는 패킷을 허용합니다.
- DROP — 요청에 응답하지 않고 패킷을 취소합니다. 패킷을 보낸 시스템은 이러한 실패를 통보하지 않습니다.
- QUEUE — 패킷은 사용자 공간 프로그램으로 처리되기 위해 대기합니다.
- RETURN — 현재 chain에 있는 규칙에 대해 패킷을 확인하는 것을 중지합니다. RETURN 대상과 함께 패킷이 다른 chain에서 불려진 chain에 있는 규칙에 일치할 경우, 패킷은 확인이 중지된 곳에서 규칙을 조사하기 위해 첫 번째 chain으로 복귀합니다. RETURN 규칙이 내장된 chain에서 사용되고 패킷을 이전 chain으로 이동시킬 수 없을 경우, 현재 chain의 기본 대상이 사용됩니다.

In addition, extensions are available which allow other targets to be specified. These extensions are called target modules or match option modules and most only apply to specific tables and situations. Refer to 46.9.3.4.4절. "추가 일치 옵션 모듈" for more information about match option modules.

여러 확장된 대상 모듈이 있으며, 대부분 특정 테이블이나 상황에만 적용됩니다. Red Hat Enterprise Linux에서 기본값으로 포함된 가장 많이 사용되는 대상 모듈에는 다음과 같은 것이 있습니다:

- LOG — 이러한 규칙에 일치하는 모든 패킷을 기록합니다. 패킷이 커널에 의해 기록되므로, /etc/syslog.conf 파일은 이러한 로그 항목이 기록되는 장소를 결정합니다. 이는 기본값으로 /var/log/messages 파일에 위치하게 됩니다.

  기록하는 방법을 지정하기 위해 LOG 대상 뒤에 추가 옵션이 사용될 수 있습니다.

  - --log-level — 기록하는 작업에 대한 우선 순위를 설정합니다. 우선 순위 목록에 대한 정보는 syslog.conf 메뉴얼 페이지에서 참조하시기 바랍니다.
  - --log-ip-options — IP 패킷의 헤더에 설정된 모든 옵션을 기록합니다.
  - --log-prefix — 로그 행이 기록되기 전에 최대 29개로된 문자열을 만듭니다. 이는 syslog 필터의 기록을 위해 패킷 기록과 함께 사용하는 것이 유용합니다.

    > **알림**
    >
    > 이러한 옵션과 관련하여, log-prefix 값에 공백을 추가하셔야 합니다.

  - --log-tcp-options — TCP 패킷의 헤더에 설정된 모든 옵션을 기록합니다.
  - --log-tcp-sequence — 로그에 있는 패킷에 해당하는 TCP 순서 번호를 기록합니다.
- REJECT — 오류 패킷을 원격 시스템에 되돌려 보내고 패킷을 취소합니다.

  The REJECT target accepts --reject-with <type> (where <type> is the rejection type) allowing more detailed information to be returned with the error packet. The message port-unreachable is the default error type given if no other option is used. Refer to the iptables man page for a full list of <type> options.

nat 테이블을 사용하는 IP 매스커레이딩이나 또는 mangle 테이블을 사용하는 패킷 변경에 유용한 여러 기능을 포함하는 대상 확장기능은 iptables 메뉴얼 페이지에서 확인하실 수 있습니다.

46.9.3.6. 옵션 목록

The default list command, iptables -L [<chain-name>], provides a very basic overview of the default filter table's current chains. Additional options provide more information:

- -v — 각각의 chain에서 실행되는 패킷과 바이트의 수, 각각의 규칙과 일치하는 패킷과 바이트의 수, 그리고 어떤 인터페이스가 특정 규칙에 적용되는가와 같은 상세한 출력을 보여줍니다.
- -x — 패킷과 바이트 숫자를 정확한 값으로 확장합니다. 복잡한 시스템에서는 특정 chain이나 규칙에 의해 실행되는 패킷과 바이트의 수는 Kilobytes,Megabytes (Megabytes) 또는 Gigabytes로 단축됩니다. 이러한 옵션은 전체 숫자가 나타나도록 강제합니다.
- -n — 기본 호스트명 및 네트워크 서비스 포맷이 아닌 숫자로된 포맷에서 IP 주소와 포트 번호를 보여줍니다.
- --line-numbers — chain에 있는 숫자로된 순서 옆의 각각의 chain에 있는 규칙 목록을 만듭니다. 이러한 옵션은 chain에 있는 특정 규칙을 삭제하거나 또는 chain안에서 규칙을 삽입하기 위한 위치를 지정할 때에 유용합니다.
- -t <table-name> — Specifies a table name. If omitted, defaults to the filter table.

다음의 예는 여러 옵션의 사용을 보여줍니다. -x 옵션을 포함하여 나타나는 바이트 표시에 있어서 다른점에 주의하시기 바랍니다.

```
~]# iptables -L OUTPUT -v -n -x
Chain OUTPUT (policy ACCEPT 64005 packets, 6445791 bytes)
    pkts      bytes target     prot opt in     out     source               destination
    1593     133812 ACCEPT     icmp --  *      *       0.0.0.0/0            0.0.0.0/0

~]# iptables -L OUTPUT -v -n
Chain OUTPUT (policy ACCEPT 64783 packets, 6492K bytes)
    pkts bytes target     prot opt in     out     source               destination
    1819  153K ACCEPT     icmp --  *      *       0.0.0.0/0            0.0.0.0/0
~]#
```

46.9.4. IPTables 규칙 저장하기

iptables 명령과 함께 생성된 규칙은 메모리에 저장되어 있습니다. iptables 규칙 모음을 저장하기 전에 시스템을 재시작하실 경우, 모든 규칙이 삭제됩니다. 시스템 재부팅을 통한 넷필터 규칙은 저장되어야 합니다. 넷필터 규칙을 저장하기 위해 root로 다음의 명령을 입력하시기 바랍니다:

```
service iptables save
```

이는 iptables init 스크립트를 실행하고, /sbin/iptables-save 프로그램을 실행하며 현재 iptables 설정을 /etc/sysconfig/iptables에 기록합니다. 기존의 /etc/sysconfig/iptables 파일은 /etc/sysconfig/iptables.save로 저장됩니다.

다음번에 시스템이 부팅할 때, iptables init 스크립트는 /sbin/iptables-restore 명령을 사용하여 /etc/sysconfig/iptables에 저장된 규칙을 재적용합니다.

While it is always a good idea to test a new iptables rule before committing it to the /etc/sysconfig/iptables file, it is possible to copy iptables rules into this file from another system's version of this file. This provides a quick way to distribute sets of iptables rules to multiple machines.

배포, 백업, 또는 기타 다른 목적을 위해 iptables 규칙을 분리된 파일에 저장하실 수 있습니다. iptables 규칙을 저장하기 위해, root로 다음의 명령을 입력하시기 바랍니다:

```
iptables-save > <filename>
```

where <filename> is a user-defined name for your ruleset.

중요

/etc/sysconfig/iptables 파일을 다른 컴퓨터로 배포할 경우, 새로운 규칙을 실행하기 위해 /sbin/service iptables restart을 입력하시기 바랍니다.

알림

iptables 기능을 구성하는 테이블과 chain을 조작하기 위해 사용되는 iptables 명령 (/sbin/iptables)와 iptables 서비스 자체를 활성화 및 비활성화하는데 사용되는 iptables 서비스 (/sbin/iptables service) 사이의 다른점에 주의하시기 바랍니다.

46.9.5. IPTables 제어 스크립트

Red Hat Enterprise Linux에서 iptables 제어하는 두가지 기본적인 방법이 있습니다:

- Security Level Configuration Tool (system-config-securitylevel) — A graphical interface for creating, activating, and saving basic firewall rules. Refer to 46.8.2절. "Basic Firewall Configuration" for more information.
- /sbin/service iptables <option> — Used to manipulate various functions of iptables using its initscript. The following options are available:
  - start — 방화벽이 설정되어 있을 경우 (즉, /etc/sysconfig/iptables가 존재할 경우), 실행되는 모든 iptables는 완전히 정지되며 /sbin/iptables-restore 명령을 사용하기 시작합니다. 이러한 옵션은 ipchains 커널 모듈을 읽어오지 않았을 때에만 작동합니다. 이러한 모듈을 읽어왔는지를 확인하시려면, root로 다음의 명령을 입력하시기 바랍니다:

    ```
    lsmod | grep ipchains
    ```

    이러한 명령으로 아무런 출력 결과가 나타나지 않으면, 이는 모듈이 읽어지지 않았음을 의미합니다. 필요하신 경우, 모듈을 삭제하기 위해 /sbin/rmmod 명령을 사용하시기 바랍니다.
  - stop — 방화벽이 실행되고 있을 경우, 메모리에 있는 방화벽 규칙은 삭제되며, 모든 iptables 모듈 및 도움 프로그램이 제거됩니다.

/etc/sysconfig/iptables-config 설정 파일에서 IPTABLES_SAVE_ON_STOP 지시문이 기본값에서 yes로 변경된 경우, 현재 규칙은 /etc/sysconfig/iptables에 저장되며 기존의 규칙은 /etc/sysconfig/iptables.save 파일로 이동합니다.

Refer to 46.9.5.1절. "IPTables 제어 스크립트 설정 파일" for more information about the iptables-config file.

- restart — 방화벽이 실행되고 있는 경우, 메모리에 있는 방화벽 규칙은 삭제되며, 방화벽이 /etc/sysconfig/iptables에 설정되어 있을 경우 이를 다시 시작하게 됩니다. 이러한 옵션은 ipchains 커널 모듈을 읽어오지 않았을 경우에만 작동합니다.

  /etc/sysconfig/iptables-config 설정 파일에서 IPTABLES_SAVE_ON_RESTART 지시문이 기본값에서 yes로 변경된 경우, 현재 규칙은 /etc/sysconfig/iptables에 저장되며 기존의 규칙은 /etc/sysconfig/iptables.save 파일로 이동합니다.

  Refer to 46.9.5.1절. "IPTables 제어 스크립트 설정 파일" for more information about the iptables-config file.

- status — 방화벽의 상태를 보여주며 모든 활성 규칙 목록을 만듭니다.

  The default configuration for this option displays IP addresses in each rule. To display domain and hostname information, edit the /etc/sysconfig/iptables-config file and change the value of IPTABLES_STATUS_NUMERIC to no. Refer to 46.9.5.1절. "IPTables 제어 스크립트 설정 파일" for more information about the iptables-config file.

- panic — 모든 방화벽 규칙을 삭제합니다. 모든 설정 테이블에 대한 정책은 DROP에 설정됩니다.

  이러한 옵션은 서버가 절충될 수 있을 경우 유용합니다. 물리적으로 네트워크로 연결을 해제하거나 또는 시스템을 중지시키지 않고, 모든 네트워크 소통을 중지시키기 위해 이러한 옵션을 사용하실 수 있지만 컴퓨터 진단과 포렌식스를 위해 이를 준비상태로 두시는 것이 좋습니다.

- save — Saves firewall rules to /etc/sysconfig/iptables using iptables-save. Refer to 46.9.4절. "IPTables 규칙 저장하기" for more information.

Tip

To use the same initscript commands to control netfilter for IPv6, substitute ip6tables for iptables in the /sbin/service commands listed in this section. For more information about IPv6 and netfilter, refer to 46.9.6절. "IPTables 및 IPv6".

46.9.5.1. IPTables 제어 스크립트 설정 파일

iptables initscripts의 동작은 /etc/sysconfig/iptables-config 설정 파일에 의해 제어됩니다. 다음은 이러한 파일에 들어 있는 지시문 목록입니다:

- IPTABLES_MODULES — 방화벽이 활성화 상태일 때, 불러오려는 추가 iptables 모듈의 목록을 공백으로 구분하여 지정합니다. 이에는 연결 추적 및 NAT 도움말이 포함될 수 있습니다.

- IPTABLES_MODULES_UNLOAD — 재시작 및 정지 상태에서 모듈을 읽어오지 않습니다. 이러한 지시문은 다음과 같은 값을 허용합니다:

  - yes — 기본 값. 방화벽 재시작 또는 정지에 대한 올바른 상태를 실행하기 위해 이러한 옵션이 설정되어야 합니다.
  - no — 넷필터 모듈을 제거하는 데 문제가 있을 경우에만 이러한 옵션을 설정해야 합니다.

- IPTABLES_SAVE_ON_STOP — 방화벽이 중지되었을 때 현재 방화벽 규칙을 /etc/sysconfig/iptables에 저장합니다. 이러한 지시문은 다음과 같은 규칙을 허용합니다:
  - yes — 방화벽이 중지되었을 때 기존의 규칙을 /etc/sysconfig/iptables에 저장하고, 이전 버전을 /etc/sysconfig/iptables.save 파일로 이동시킵니다.
  - no — 기본값. 방화벽이 중지되었을 때 기존 규칙을 저장하지 않습니다.
- IPTABLES_SAVE_ON_RESTART — 방화벽을 재시작할 때 현재 방화벽 규칙을 저장합니다. 이러한 지시문은 다음과 같은 값을 허용합니다:
  - yes — 방화벽을 재시작할 때, 기존의 규칙을 /etc/sysconfig/iptables에 저장하고, 이전 버전을 /etc/sysconfig/iptables.save 파일로 이동시킵니다.
  - no — 기본값. 방화벽을 재시작할 때 기존 규칙을 저장하지 않습니다.
- IPTABLES_SAVE_COUNTER — 모든 chain 및 규칙에 있는 모든 패킷과 바이트를 저장하거나 복구합니다. 이러한 지시문은 다음과 같은 값을 허용합니다:
  - yes — 카운터 값을 저장합니다.
  - no — 기본값. 카운터 값을 저장하지 않습니다.
- IPTABLES_STATUS_NUMERIC — 도메인 또는 호스트명을 대신하여 숫자로된 포맷에서 IP 주소를 출력합니다. 이러한 지시문은 다음과 같은 값을 허용합니다:
  - yes — 기본값. 출력 상태안에서 IP 주소만을 복귀하게 합니다.
  - no — 출력 상태 안에서 도메인 또는 호스트명을 복귀하게 합니다.

46.9.6. IPTables 및 IPv6

iptables-ipv6 패키지가 설치되었을 경우, Red Hat Enterprise Linux에 있는 넷필터는 차세대 IPv6 인터넷 프로토콜을 거를수 있습니다. IPv6 넷필터를 조작하기 위해 사용되는 명령은 ip6tables입니다.

이 명령에 대한 대부분의 지시문은 iptables에서 사용된 것과 동일하지만, nat 테이블은 아직 지원되지 않습니다. 즉, 이는 매스커레이딩(masquerading) 및 포트 포워딩과 같은 IPv6 네트워크 주소 번역 작업을 아직 실행할 수 없음을 의미합니다.

ip6tables에 대한 규칙은 /etc/sysconfig/ip6tables 파일에 저장되어 있습니다. ip6tables initscripts에 의해 저장된 기존 규칙은 /etc/sysconfig/ip6tables.save 파일에 저장되어 있습니다.

ip6tables init 스크립트에 대한 설정 옵션은 /etc/sysconfig/ip6tables-config에 저장되어 있으며, 각각의 지시문에 해당하는 이름은 iptables 부분에 따라 다릅니다.

예, iptables-config 지시문 IPTABLES_MODULES: ip6tables-config 파일과 같은 것은 IP6TABLES_MODULES 입니다.

46.9.7. 추가 자료

iptables과 함께 패킷 필터링에 관한 추가 정보는 다음 소스에서 참조하시기 바랍니다.

- 46.8절. "Firewalls" — Contains a chapter about the role of firewalls within an overall security strategy as well as strategies for constructing firewall rules.

46.9.7.1. 설치된 문서

- man iptables — iptables에 대한 설명 및 대상, 옵션, 일치 확장에 대한 종합적인 목록이 포함되어 있습니다.

46.9.7.2. 유용한 웹사이트

- http://www.netfilter.org/ — 넷필터(netfilter)/iptables 프로젝트의 홈. iptables에 관한 정보는 물론, 특정 문제를 다루고 있는 FAQ 및 Linux IP 방화벽 관리자인 Rusty Russell에 의해 쓰여진 도움말 등이 들어있습니다. 사이트에 있는 HOWTO 문서에서는 기본적인 네트워킹 개념, 커널 패킷 필터링, NAT 설정과 같은 주제를 다루고 있습니다.

- http://www.linuxnewbie.org/nhf/Security/IPtables_Basics.html — Linux 커널을 통해 패킷을 이동하는 방법에 대한 소개와 함께 기본 iptables 명령을 구축하는 방법에 대한 소개.

보안 및 SELinux

47.1. Access Control Mechanisms (ACMs)

This section provides a basic introduction to Access Control Mechanisms (ACMs). ACMs provide a means for system administrators to control which users and processes can access different files, devices, interfaces, etc., in a computer system. This is a primary consideration when securing a computer system or network of any size.

47.1.1. Discretionary Access Control (DAC)

Discretionary Access Control (DAC) defines the basic access controls for objects in a filesystem. This is the typical access control provided by file permissions, sharing, etc. Such access is generally at the discretion of the owner of the object (file, directory, device, etc.).

DAC provides a means of restricting access to objects based on the identity of the users or groups (subjects) that try to access those objects. Depending on a subject's access permissions, they may also be able to pass permissions to other subjects.

47.1.2. Access Control Lists (ACLs)

Access Control Lists (ACLs) provide further control over which objects a subject can access. For more information, refer to 9장. Access Control Lists.

47.1.3. Mandatory Access Control (MAC)

Mandatory Access Control (MAC) is a security mechanism that restricts the level of control that users (subjects) have over the objects that they create. Unlike in a DAC implementation, where users have full control over their own files, directories, etc., MAC adds additional labels, or categories, to all file system objects. Users and processes must have the appropriate access to these categories before they can interact with these objects.

In Red Hat Enterprise Linux, MAC is enforced by SELinux. For more information, refer to 47.2절. “Introduction to SELinux” .

47.1.4. Role-based Access Control (RBAC)

Role-based Access Control (RBAC) is an alternative method of controlling user access to file system objects. Instead of access being controlled by user permissions, the system administrator establishes Roles based on business functional requirements or similar criteria. These Roles have different types and levels of access to objects.

In contrast to DAC or MAC systems, where users have access to objects based on their own and the object's permissions, users in an RBAC system must be members of the appropriate group, or Role, before they can interact with files, directories, devices, etc.

From an administrative point of view, this makes it easier to control who has access to various parts of the file system, just by controlling their group memberships.

47.1.5. Multi-Level Security (MLS)

Multi-Level Security (MLS) is a specific Mandatory Access Control (MAC) security scheme. Under this scheme, processes are called Subjects. Files, sockets and other passive operating system entities are called Objects. For more information, refer to 47.6절. “Multi-Level Security (MLS)” .

47.1.6. Multi-Category Security (MCS)

Multi-Category Security (MCS) is an enhancement to SELinux, and allows users to label files with categories. MCS is an adaptation of MLS and re-uses much of the MLS framework in SELinux. For more information, refer to 47.4.1절. "Introduction"

47.2. Introduction to SELinux

Security-Enhanced Linux (SELinux) is a security architecture integrated into the 2.6.x kernel using the Linux Security Modules (LSM). It is a project of the United States National Security Agency (NSA) and the SELinux community. SELinux integration into Red Hat Enterprise Linux was a joint effort between the NSA and Red Hat.

47.2.1. SELinux Overview

SELinux provides a flexible Mandatory Access Control (MAC) system built into the Linux kernel. Under standard Linux Discretionary Access Control (DAC), an application or process running as a user (UID or SUID) has the user's permissions to objects such as files, sockets, and other processes. Running a MAC kernel protects the system from malicious or flawed applications that can damage or destroy the system.

SELinux defines the access and transition rights of every user, application, process, and file on the system. SELinux then governs the interactions of these entities using a security policy that specifies how strict or lenient a given Red Hat Enterprise Linux installation should be.

On a day-to-day basis, system users will be largely unaware of SELinux. Only system administrators need to consider how strict a policy to implement for their server environment. The policy can be as strict or as lenient as needed, and is very finely detailed. This detail gives the SELinux kernel complete, granular control over the entire system.

The SELinux Decision Making Process

When a subject, (for example, an application), attempts to access an object (for example, a file), the policy enforcement server in the kernel checks an access vector cache (AVC), where subject and object permissions are cached. If a decision cannot be made based on data in the AVC, the request continues to the security server, which looks up the security context of the application and the file in a matrix. Permission is then granted or denied, with an avc: denied message detailed in /var/log/messages if permission is denied. The security context of subjects and objects is applied from the installed policy, which also provides the information to populate the security server's matrix.

Refer to the following diagram:

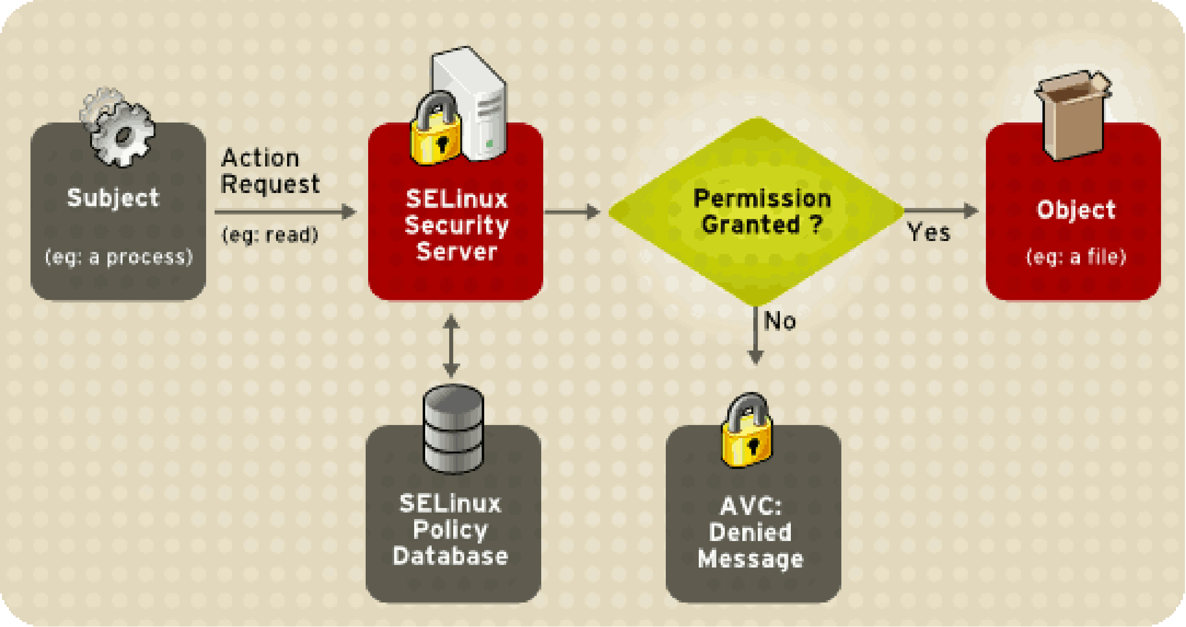

그림 47.1. SELinux Decision Process

SELinux Operating Modes

Instead of running in enforcing mode, SELinux can run in permissive mode, where the AVC is checked and denials are logged, but SELinux does not enforce the policy. This can be useful for troubleshooting and for developing or fine-tuning SELinux policy.

For more information about how SELinux works, refer to 47.2.3절. “Additional Resources” .

47.2.2. Files Related to SELinux

The following sections describe SELinux configuration files and related file systems.

47.2.2.1. The SELinux Pseudo-File System

The /selinux/ pseudo-file system contains commands that are most commonly used by the kernel subsystem. This type of file system is similar to the /proc/ pseudo-file system.

Administrators and users do not normally need to manipulate this component.

The following example shows sample contents of the /selinux/ directory:

```
-rw-rw-rw-  1 root root 0 Sep 22 13:14 access
dr-xr-xr-x  1 root root 0 Sep 22 13:14 booleans
--w-------  1 root root 0 Sep 22 13:14 commit_pending_bools
-rw-rw-rw-  1 root root 0 Sep 22 13:14 context
-rw-rw-rw-  1 root root 0 Sep 22 13:14 create
--w-------  1 root root 0 Sep 22 13:14 disable
-rw-r--r--  1 root root 0 Sep 22 13:14 enforce
-rw-------  1 root root 0 Sep 22 13:14 load
-r--r--r--  1 root root 0 Sep 22 13:14 mls
-r--r--r--  1 root root 0 Sep 22 13:14 policyvers
-rw-rw-rw-  1 root root 0 Sep 22 13:14 relabel
-rw-rw-rw-  1 root root 0 Sep 22 13:14 user
```

For example, running the cat command on the enforce file reveals either a 1 for enforcing mode or 0 for permissive mode.

47.2.2.2. SELinux Configuration Files

The following sections describe SELinux configuration and policy files, and related file systems located in the /etc/ directory.

47.2.2.2.1. The /etc/sysconfig/selinux Configuration File

There are two ways to configure SELinux under Red Hat Enterprise Linux: using the SELinux Administration Tool (system-config-selinux), or manually editing the configuration file (/etc/sysconfig/selinux).

The /etc/sysconfig/selinux file is the primary configuration file for enabling or disabling SELinux, as well as for setting which policy to enforce on the system and how to enforce it.

> **Note**
>
> The /etc/sysconfig/selinux contains a symbolic link to the actual configuration file, /etc/selinux/config.

The following explains the full subset of options available for configuration:

- SELINUX=enforcing|permissive|disabled — Defines the top-level state of SELinux on a system.
  - enforcing — The SELinux security policy is enforced.
  - permissive — The SELinux system prints warnings but does not enforce policy.

    This is useful for debugging and troubleshooting purposes. In permissive mode, more denials are logged because subjects can continue with actions that would otherwise be denied in enforcing mode. For example, traversing a directory tree in permissive mode produces avc: denied messages for every directory level read. In enforcing mode, SELinux would have stopped the initial traversal and kept further denial messages from occurring.
  - disabled — SELinux is fully disabled. SELinux hooks are disengaged from the kernel and the pseudo-file system is unregistered.

Tip

Actions made while SELinux is disabled may result in the file system no longer having the correct security context. That is, the security context defined by the policy. The best way to relabel the file system is to create the flag file /.autorelabel and reboot the machine. This causes the relabel to occur very early in the boot process, before any processes are running on the system. Using this procedure means that processes can not accidentally create files in the wrong context or start up in the wrong context.

It is possible to use the fixfiles relabel command prior to enabling SELinux to relabel the file system. This method is not recommended, however, because after it is complete, it is still possible to have processes potentially running on the system in the wrong context. These processes could create files that would also be in the wrong context.

Note

Additional white space at the end of a configuration line or as extra lines at the end of the file may cause unexpected behavior. To be safe, remove unnecessary white space.

- SELINUXTYPE=targeted|strict — Specifies which policy SELinux should enforce.
  - targeted — Only targeted network daemons are protected.

Important

The following daemons are protected in the default targeted policy: dhcpd, httpd (apache.te), named, nscd, ntpd, portmap, snmpd, squid, and syslogd. The rest of the system runs in the unconfined_t domain. This domain allows subjects and objects with that security context to operate using standard Linux security.

The policy files for these daemons are located in /etc/selinux/targeted/src/policy/domains/program. These files are subject to change as newer versions of Red Hat Enterprise Linux are released.

Policy enforcement for these daemons can be turned on or off, using Boolean values controlled by the SELinux Administration Tool (system-config-selinux).

Setting a Boolean value for a targeted daemon to 1 disables SELinux protection for the daemon. For example, you can set dhcpd_disable_trans to 1 to prevent init, which executes apps labeled dhcpd_exec_t, from transitioning to the dhcpd_t domain.

  Use the getsebool -a command to list all SELinux booleans. The following is an example of using the setsebool command to set an SELinux boolean. The -P option makes the change permanent. Without this option, the boolean would be reset to 1 at reboot.

  ```
  setsebool -P dhcpd_disable_trans=0
  ```

  - strict — Full SELinux protection, for all daemons. Security contexts are defined for all subjects and objects, and every action is processed by the policy enforcement server.

- SETLOCALDEFS=0|1 — Controls how local definitions (users and booleans) are set. Set this value to 1 to have these definitions controlled by load_policy from files in /etc/selinux/<policyname>. or set it to 0 to have them controlled by semanage.

> **Caution**
>
> You should not change this value from the default (0) unless you are fully aware of the impact of such a change.

47.2.2.2.2. The /etc/selinux/ Directory

The /etc/selinux/ directory is the primary location for all policy files as well as the main configuration file.

The following example shows sample contents of the /etc/selinux/ directory:

```
-rw-r--r--  1 root root  448 Sep 22 17:34 config
drwxr-xr-x  5 root root 4096 Sep 22 17:27 strict
drwxr-xr-x  5 root root 4096 Sep 22 17:28 targeted
```

The two subdirectories, strict/ and targeted/, are the specific directories where the policy files of the same name (that is, strict and targeted) are contained.

47.2.2.3. SELinux Utilities

The following are some of the commonly used SELinux utilities:

- /usr/sbin/setenforce — Modifies in real-time the mode in which SELinux runs.

  For example:

  setenforce 1 — SELinux runs in enforcing mode.

  setenforce 0 — SELinux runs in permissive mode.

  To actually disable SELinux, you need to either specify the appropriate setenforce parameter in /etc/sysconfig/selinux or pass the parameter selinux=0 to the kernel, either in /etc/grub.conf or at boot time.

- /usr/sbin/sestatus -v — Displays the detailed status of a system running SELinux. The following example shows an excerpt of sestatus -v output:

```
SELinux status:                 enabled
SELinuxfs mount:                /selinux
Current mode:                   enforcing
Mode from config file:          enforcing
Policy version:                 21
Policy from config file:        targeted

Process contexts:
Current context:                user_u:system_r:unconfined_t:s0
Init context:                   system_u:system_r:init_t:s0
/sbin/mingetty                  system_u:system_r:getty_t:s0
```

- /usr/bin/newrole — Runs a new shell in a new context, or role. Policy must allow the transition to the new role.

> **Note**
>
> This command is only available if you have the policycoreutils-newrole package installed, which is required for the strict and MLS policies.

- /sbin/restorecon — Sets the security context of one or more files by marking the extended attributes with the appropriate file or security context.

- /sbin/fixfiles — Checks or corrects the security context database on the file system.

Refer to the man page associated with these utilities for more information.

Refer to the setools or policycoreutils package contents for more information on all available binary utilities. To view the contents of a package, use the following command:

```
rpm -ql <package-name>
```

47.2.3. Additional Resources

Refer to the following resources for more detailed information on SELinux.

47.2.3.1. Installed Documentation

- /usr/share/doc/setools-<version-number>/ All documentation for utilities contained in the setools package. This includes all helper scripts, sample configuration files, and documentation.

47.2.3.2. Useful Websites

- http://www.nsa.gov/research/selinux/index.shtml Homepage for the NSA SELinux development team. Many resources are available in HTML and PDF formats. Although many of these links are not SELinux specific, some concepts may apply.

- http://docs.fedoraproject.org/ Homepage for the Fedora documentation project, which contains Fedora Core specific materials that may be more timely, since the release cycle is much shorter.
- http://selinux.sourceforge.net Homepage for the SELinux community.

47.3. SELinux의 전반적인 배경 및 역사

SELinux was originally a development project from the National Security Agency (NSA)[1] and others. It is an implementation of the Flask operating system security architecture.[2]The NSA integrated SELinux into the Linux kernel using the Linux Security Modules (LSM) framework. SELinux motivated the creation of LSM, at the suggestion of Linus Torvalds, who wanted a modular approach to security instead of just accepting SELinux into the kernel.

원래 persistent security ID (PSID)를 사용한 SELinux 실행은 ext2 inode의 사용되지 않는 영역에 저장되었습니다. 이렇게 숫자로된 표현 (즉, 읽을 수 없는 표현)은 보안 문맥 레이블로 SELinux에 의해 매핑되었습니다. 하지만, 이는 PSID를 지원하기 위해 각각의 파일 시스템을 수정해야 하므로, 스케일러블 솔루션이 아니지만 Linux 커널에서 업스트림을 지원할 수 있는 솔루션입니다.

The next evolution of SELinux was as a loadable kernel module for the 2.4.<x> series of Linux kernels. This module stored PSIDs in a normal file, and SELinux was able to support more file systems. This solution was not optimal for performance, and was inconsistent across platforms. Finally, the SELinux code was integrated upstream to the 2.6.x kernel, which has full support for LSM and has extended attributes (xattrs) in the ext3 file system. SELinux was moved to using xattrs to store security context information. The xattr namespace provides useful separation for multiple security modules existing on the same system.

업스트림에 대한 커널을 준비하기 위한 여러 작업과 차후의 SELinux 개발은NSA, Red Hat 및 SELinux 개발자 커뮤니티의 공동 노력으로 이루어 집니다.

For more information about the history of SELinux, the definitive website is http://www.nsa.gov/research/selinux/index.shtml.

47.4. Multi-Category Security (MCS)

47.4.1. Introduction

Multi-Category Security (MCS) is an enhancement to SELinux, and allows users to label files with categories. These categories are used to further constrain Discretionary Access Control (DAC) and Type Enforcement (TE) logic. They may also be used when displaying or printing files. An example of a category is "Company_Confidential". Only users with access to this category can access files labeled with the category, assuming the existing DAC and TE rules also permit access.

The term categories refers to the same non-hierarchical categories used by Multi-Level Security (MLS). Under MLS, objects and subjects are labeled with Security Levels. These Security Levels consist of a hierarchical sensitivity value (such as "Top Secret") and zero or more non-hierarchical categories (such as "Crypto"). Categories provide compartments within sensitivity levels and enforce the need-to-know security principle. Refer to 47.6절. “Multi-Level Security (MLS)” for more information about Multi-Level Security.

[1] The NSA is the cryptologic agency of the United States of America's Federal government, charged with information assurance and signals intelligence. You can read more about the NSA at their website, http://www.nsa.gov/about/.

[2] Flask grew out of a project that integrated the Distributed Trusted Operating System (DTOS) into the Fluke research operating system. Flask was the name of the architecture and the implementation in the Fluke operating system.

47.4.1.1. What is Multi-Category Security?

MCS is an adaptation of MLS. From a technical point of view, MCS is a policy change, combined with a few userland modifications to hide some of the unneeded MLS technology. Some kernel changes were also made, but only relating to making it easy to upgrade to MCS (or MLS) without invoking a full file system relabel.

The hope is to improve the quality of the system as a whole, reduce costs, leverage the open source process, increase transparency, and make the technology base useful to more than a small group of extremely special-case users.

47.4.2. Applications for Multi-Category Security

Beyond access control, MCS could be used to display the MCS categories at the top and bottom of printed pages. This may also include a cover sheet to indicate document handling procedures. It should also be possible to integrate MCS with future developments in SELinux, such as Security Enhanced X. Integration with a directory server will also make MCS support for email easier. This could involve users manually labeling outgoing emails or by attaching suitably labeled files. The email client can then determine whether the recipients are known to be cleared to access the categories associated with the emails.

47.4.3. SELinux Security Contexts

SELinux stores security contexts as an extended attribute of a file. The "security." namespace is used for security modules, and the security.selinux name is used to persistently store SELinux security labels on files. The contents of this attribute will vary depending on the file or directory you inspect and the policy the machine is enforcing.

Note

This is expected to change in the 2.6.15 kernel (and already has in the latest -mm kernels), so that getxattr(2) always returns the kernel's canonicalized version of the label.

You can use the ls -Z command to view the category label of a file:

```
~]# ls -Z gravityControl.txt
-rw-r--r--   user      user      user_u:object_r:tmp_t:Moonbase_Plans gravityControl.txt
```

You can use the gefattr(1) command to view the internal category value (c10):

```
~]# getfattr -n security.selinux gravityControl.txt
# file: gravityControl.txt
security.selinux="user_u:object_r:tmp_t:s0:c10\000"
```

Refer to 47.5절. “Getting Started with Multi-Category Security (MCS)” for details on creating categories and assigning them to files.

47.5. Getting Started with Multi-Category Security (MCS)

This section provides an introduction to using MCS labels to extend the Mandatory Access Control (MAC) capabilities of SELinux. It discusses MCS categories, SELinux user identities, and how they apply to Linux user accounts and files. It builds on the conceptual information provided in 47.4절. "Multi-Category Security (MCS)" , and introduces some basic examples of usage.

47.5.1. Introduction

MCS labeling from a user and system administrator standpoint is straightforward. It consists of configuring a set of categories, which are simply text labels, such as "Company_Confidential" or "Medical_Records", and then assigning users to those categories. The system administrator first configures the categories, then assigns users to them as required. The users can then use the labels as they see fit.

The names of the categories and their meanings are set by the system administrator, and can be set to whatever is required for the specific deployment. A system in a home environment may have only one category of "Private", and be configured so that only trusted local users are assigned to this category.

In a corporate environment, categories could be used to identify documents confidential to specific departments. Categories could be established for "Finance", "Payroll", "Marketing", and "Personnel". Only users assigned to those categories can access resources labeled with the same category.

After users have been assigned to categories, they can label any of their own files with any of the categories to which they have been assigned. For example, a home user in the system described above could label all of their personal files as "Private", and no service such as Apache or vsftp would ever be able to access those files, because they don't have access to the "Private" category.

MCS works on a simple principle: to access a file, a user needs to be assigned to all of the categories with which the file is labeled. The MCS check is applied after normal Linux Discretionary Access Control (DAC) and Type Enforcement (TE) rules, so it can only further restrict security.

47.5.2. Comparing SELinux and Standard Linux User Identities

SELinux maintains its own user identity for processes, separately from Linux user identities. In the targeted policy (the default for Red Hat Enterprise Linux), only a minimal number of SELinux user identities exist:

- system_u — System processes
- root — System administrator
- user_u — All login users

Use the semanage user -l command to list SELinux users:

```
~]# semanage user -l

                Labeling   MLS/          MLS/
SELinux User    Prefix     MCS Level   MCS Range               SELinux Roles

root            user       s0          s0-s0:c0.c1023        system_r sysadm_r user_r
system_u        user       s0          s0-s0:c0.c1023        system_r
user_u          user       s0          s0-s0:c0.c1023        system_r sysadm_r user_r
```

Refer to 47.8.3절. "Understanding the Users and Roles in the Targeted Policy" for more information about SELinux users and roles.

SELinux Logins

One of the properties of targeted policy is that login users all run in the same security context. From a TE point of view, in targeted policy, they are security-equivalent. To effectively use MCS, however, we need to be able to assign different sets of categories to different Linux users, even though they are all the same SELinux user (user_u). This is solved by introducing the concept of an SELinux login. This is used during the login process to assign MCS categories to Linux users when their shell is launched.

Use the semanage login -a command to assign Linux users to SELinux user identities:

```
~]# semanage login -a james
~]# semanage login -a daniel
~]# semanage login -a olga
```

Now when you list the SELinux users, you can see the Linux users assigned to a specific SELinux user identity:

```
~]# semanage login -l

Login Name                SELinux User              MLS/MCS Range

__default__              user_u                     s0
james                    user_u                     s0
daniel                   user_u                     s0
root                     root                       s0-s0:c0.c1023
olga                     user_u                     s0
```

Notice that at this stage only the root account is assigned to any categories. By default, the root account is configured with access to all categories.

Red Hat Enterprise Linux and SELinux are preconfigured with several default categories, but to make effective use of MCS, the system administrator typically modifies these or creates further categories to suit local requirements.

47.5.3. Configuring Categories

SELinux maintains a mapping between internal sensitivity and category levels and their human-readable representations in the setrans.conf file. The system administrator edits this file to manage and maintain the required categories.

Use the chcat -L command to list the current categories:

```
~]# chcat -L
s0
s0-s0:c0.c1023                     SystemLow-SystemHigh
s0:c0.c1023                        SystemHigh
```

To modify the categories or to start creating your own, modify the /etc/selinux/<selinuxtype>/setrans.conf file. For the example introduced above, add the Marketing, Finance, Payroll, and Personnel categories as follows (this example uses the targeted policy, and irrelevant sections of the file have been omitted):

```
~]# vi /etc/selinux/targeted/setrans.conf
s0:c0=Marketing
s0:c1=Finance
s0:c2=Payroll
```

```
s0:c3=Personnel
```

Use the chcat -L command to check the newly-added categories:

```
~]# chcat -L
s0:c0                          Marketing
s0:c1                          Finance
s0:c2                          Payroll
s0:c3                          Personnel
s0
s0-s0:c0.c1023                 SystemLow-SystemHigh
s0:c0.c1023                    SystemHigh
```

Note

After you make any changes to the setrans.conf file, you need to restart the MCS translation service before those changes take effect. Use the following command to restart the service:

```
~]# service mcstrans restart
```

47.5.4. Assigning Categories to Users

Now that the required categories have been added to the system, you can start assigning them to SELinux users and files. To further develop the example above, assume that James is in the Marketing department, Daniel is in the Finance and Payroll departments, and Olga is in the Personnel department. Each of these users has already been assigned an SELinux login.

Use the chcat command to assign MCS categories to SELinux logins:

```
~]# chcat -l -- +Marketing james
~]# chcat -l -- +Finance,+Payroll daniel
~]# chcat -l -- +Personnel olga
```

You can also use the chcat command with additional command-line arguments to list the categories that are assigned to users:

```
~]# chcat -L -l daniel james olga
daniel: Finance,Payroll
james: Marketing
olga: Personnel
```

You can add further Linux users, assign them to SELinux user identities and then assign categories to them as required. For example, if the company director also requires a user account with access to all categories, follow the same procedure as above:

```
# Create a user account for the company director (Karl)
~]# useradd karl
~]# passwd karl
Changing password for user karl.
New UNIX password:
Retype new UNIX password:
passwd: all authentication tokens updated successfully.

# Assign the user account to an SELinux login
```

```
~]# semanage login -a karl

# Assign all the MCS categories to the new login
~]# chcat -l -- +Marketing,+Finance,+Payroll,+Personnel karl
```

Use the chcat command to verify the addition of the new user:

```
~]# chcat -L -l daniel james olga karl
daniel: Finance,Payroll
james: Marketing
olga: Personnel
karl: Marketing,Finance,Payroll,Personnel
```

Note

MCS category access is assigned during login. Consequently, a user does not have access to newly-assigned categories until they log in again. Similarly, if access to a category is revoked, this is only apparent to the user after the next login.

47.5.5. Assigning Categories to Files

At this point we have a system that has several user accounts, each of which is mapped to an SELinux user identity. We have also established a number of categories that are suitable for the particular deployment, and assigned those categories to different users.

All of the files on the system, however, still fall under the same category, and are therefore accessible by everyone (but still according to the standard Linux DAC and TE constraints). We now need to assign categories to the various files on the system so that only the appropriate users can access them.

For this example, we create a file in Daniel's home directory:

```
[daniel@dhcp-133 ~]$ echo "Financial Records 2006" > financeRecords.txt
```

Use the ls -Z command to check the initial security context of the file:

```
[daniel@dhcp-133 ~]$ ls -Z financeRecords.txt
-rw-r--r--  daniel daniel user_u:object_r:user_home_t        financeRecords.txt
```

Notice that at this stage the file has the default context for a file created in the user's home directory (user_home_t) and has no categories assigned to it. We can add the required category using the chcat command. Now when you check the security context of the file, you can see the category has been applied.

```
[daniel@dhcp-133 ~]$ chcat -- +Finance financeRecords.txt
[daniel@dhcp-133 ~]$ ls -Z financeRecords.txt
-rw-r--r--  daniel daniel root:object_r:user_home_t:Finance financeRecords.txt
```

In many cases, you need to assign more than one category to a file. For example, some files may need to be accessible to users from both the Finance and Payroll departments.

```
[daniel@dhcp-133 ~]$ chcat -- +Payroll financeRecords.txt
```

```
[daniel@dhcp-133 ~]$ ls -Z financeRecords.txt
-rw-r--r--  daniel daniel root:object_r:user_home_t:Finance,Payroll financeRecords.txt
```

Each of the categories that have been assigned to the file are displayed in the security context. You can add and delete categories to files as required. Only users assigned to those categories can access that file, assuming that Linux DAC and TE permissions would already allow the access.

If a user who is assigned to a different category tries to access the file, they receive an error message:

```
[olga@dhcp-133 ~]$ cat financeRecords.txt
cat: financeRecords.txt: Permission Denied
```

Note

Refer to the man pages for semanage and chcat for more information on the available options for these commands.

47.6. Multi-Level Security (MLS)

Protecting sensitive or confidential data is paramount in many businesses. In the event such information is made public, businesses may face legal or financial ramifications. At the very least, they will suffer a loss of customer trust. In most cases, however, they can recover from these financial and other losses with appropriate investment or compensation.

The same cannot be said of the defense and related communities, which includes military services, intelligence organizations and some areas of police service. These organizations cannot easily recover should sensitive information be leaked, and may not recover at all. These communities require higher levels of security than those employed by businesses and other organizations.

Having information of different security levels on the same computer systems poses a real threat. It is not a straight-forward matter to isolate different information security levels, even though different users log in using different accounts, with different permissions and different access controls.

Some organizations go as far as to purchase dedicated systems for each security level. This is often prohibitively expensive, however. A mechanism is required to enable users at different security levels to access systems simultaneously, without fear of information contamination.

47.6.1. Why Multi-Level?

The term multi-level arises from the defense community's security classifications: Confidential, Secret, and Top Secret.

Individuals must be granted appropriate clearances before they can see classified information. Those with Confidential clearance are only authorized to view Confidential documents; they are not trusted to look at Secret or Top Secret information. The rules that apply to data flow operate from lower levels to higher levels, and never the reverse. This is illustrated below.

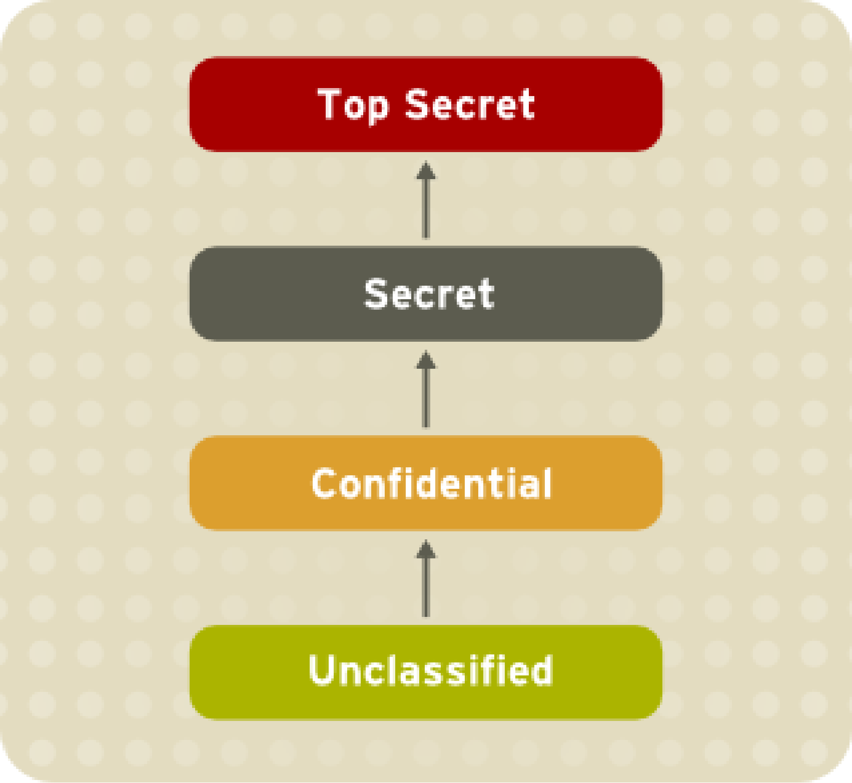

그림 47.2. Information Security Levels

47.6.1.1. The Bell-La Padula Model (BLP)

SELinux, like most other systems that protect multi-level data, uses the BLP model. This model specifies how information can flow within the system based on labels attached to each subject and object. Refer to the following diagram:

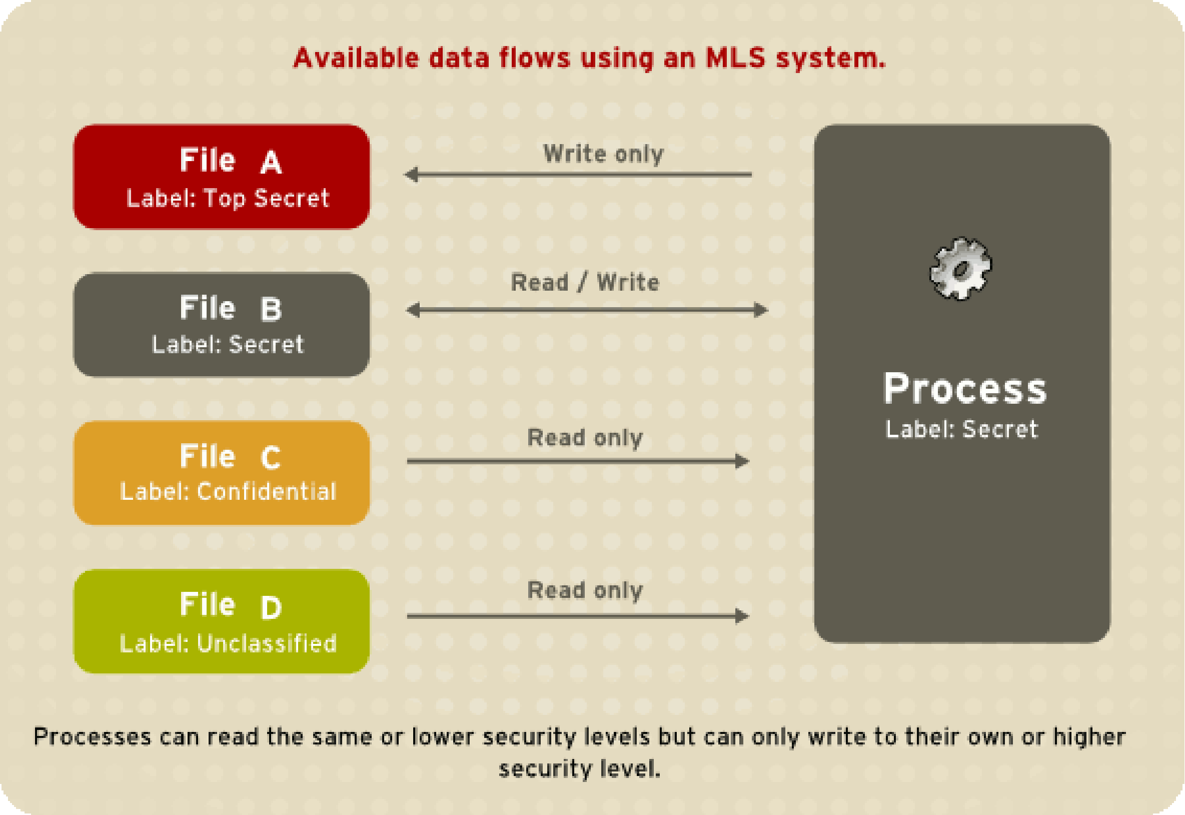

그림 47.3. Available data flows using an MLS system

Under such a system, users, computers, and networks use labels to indicate security levels. Data can flow between like levels, for example between "Secret" and "Secret", or from a lower level to a higher level. This means that users at level "Secret" can share data with one another, and can also retrieve information from Confidential-level (i.e., lower-level), users. However, data cannot flow from a higher level to a lower level. This prevents processes at the "Secret" level from viewing information classified as "Top Secret". It also prevents processes at a higher level from accidentally writing information to a lower level. This is referred to as the "no read up, no write down" model.

47.6.1.2. MLS and System Privileges

MLS access rules are always combined with conventional access permissions (file permissions). For example, if a user with a security level of "Secret" uses Discretionary Access Control (DAC) to block access to a file by other users, this also blocks access by users with a security level of "Top Secret". A higher security clearance does not automatically give permission to arbitrarily browse a file system.

Users with top-level clearances do not automatically acquire administrative rights on multi-level systems. While they may have access to all information on the computer, this is different from having administrative rights.

47.6.2. Security Levels, Objects and Subjects

As discussed above, subjects and objects are labeled with Security Levels (SLs), which are composed of two types of entities:

1. Sensitivity: — A hierarchical attribute such as "Secret" or "Top Secret".

2. Categories: — A set of non-hierarchical attributes such as "US Only" or "UFO".

An SL must have one sensitivity, and may have zero or more categories.

Examples of SLs are: { Secret / UFO, Crypto }, { Top Secret / UFO, Crypto, Stargate } and { Unclassified }

Note the hierarchical sensitivity followed by zero or more categories. The reason for having categories as well as sensitivities is so that sensitivities can be further compartmentalized on a need-to-know basis. For example, while a process may be cleared to the "Secret" sensitivity level, it may not need any type of access to the project "Warp Drive" (which could be the name of a category).

Note

1. Security Levels on objects are called Classifications.

2. Security Levels on subjects are called Clearances.

Thus, objects are labeled with a Classification, while subjects operate with a specific Clearance. Security Levels can have also Ranges, but these are beyond the scope of this introduction.

47.6.3. MLS Policy

SELinux uses the Bell-La Padula BLP model, with Type Enforcement (TE) for integrity. In simple terms, MLS policy ensures that a Subject has an appropriate clearance to access an Object of a particular classification.

For example, under MLS, the system needs to know how to process a request such as: Can a process running with a clearance of { Top Secret / UFO, Rail gun } write to a file classified as { Top Secret / UFO } ?

The MLS model and the policy implemented for it will determine the answer. (Consider, for example, the problem of information leaking out of the Rail gun category into the file).

MLS meets a very narrow (yet critical) set of security requirements based around the way information and personnel are managed in rigidly controlled environments such as the military. MLS is typically difficult to work with and does not map well to general-case scenarios.

Type Enforcement (TE) under SELinux is a more flexible and expressive security scheme, which is in many cases more suitable than MLS.

There are, however, several scenarios where traditional MLS is still required. For example, a file server where the stored data may be of mixed classification and where clients connect at different clearances. This results in a large number of Security Levels and a need for strong isolation all on a single system.

This type of scenario is the reason that SELinux includes MLS as a security model, as an adjunct to TE.

47.6.4. LSPP Certification

Efforts are being made to have Linux certified as an MLS operating system. The certification is equivalent to the old B1 rating, which has been reworked into the Labeled Security Protection Profile[3] under the Common Criteria[4] scheme.

[3] http://www.commoncriteriaportal.org/files/ppfiles/lspp.pdf

47.7. SELinux Policy Overview

This chapter is an overview of SELinux policy, some of its internals, and how it works. It discusses the policy in general terms, while 47.8절. "Targeted Policy Overview" focuses on the details of the targeted policy as it ships in Red Hat Enterprise Linux. This chapter starts with a brief overview of what policy is and where it resides.

Following on from this, the role of SELinux during the boot process is discussed. This is followed by discussions on file security contexts, object classes and permissions, attributes, types, access vectors, macros, users and roles, constraints, and a brief discussion summarizing special kernel interfaces.

47.7.1. What is the SELinux Policy?

The SELinux Policy is the set of rules that guide the SELinux security engine. It defines types for file objects and domains for processes. It uses roles to limit the domains that can be entered, and has user identities to specify the roles that can be attained. In essence, types and domains are equivalent, the difference being that types apply to objects while domains apply to processes.

47.7.1.1. SELinux Types

A type is a way of grouping items based on their similarity from a security perspective. This is not necessarily related to the unique purpose of an application or the content of a document. For example, a file can have any type of content and be for any purpose, but if it belongs to a user and exists in that user's home directory, it is considered to be of a specific security type, user_home_t.

These object types are considered alike because they are accessible in the same way by the same set of subjects. Similarly, processes tend to be of the same type if they have the same permissions as other subjects. In the targeted policy, programs that run in the unconfined_t domain have an executable file with a type such as sbin_t. From an SELinux perspective, this means they are all equivalent in terms of what they can and cannot do on the system.

For example, the binary executable file object at /usr/bin/postgres has the type postgresql_exec_t. All of the targeted daemons have their own *_exec_t type for their executable applications. In fact, the entire set of PostgreSQL executables such as createlang, pg_dump, and pg_restore have the same type, postgresql_exec_t, and they transition to the same domain, postgresql_t, upon execution.

47.7.1.1.1. Using Policy Rules to Define Type Access

The SELinux policy defines various rules which determine how each domain may access each type. Only what is specifically allowed by the rules is permitted. By default, every operation is denied and audited, meaning it is logged in the $AUDIT_LOG file. In Red Hat Enterprise Linux, this is set to /var/log/messages. The policy is compiled into binary format for loading into the kernel security server, and each time the security server makes a decision, it is cached in the AVC to optimize performance.

The policy can be defined either by modifying the existing files or by adding local Type Enforcement (TE) and File Context (FC) files to the policy tree. These new policies can be loaded into the kernel in real time. Otherwise, the policy is loaded during the boot process by init, as explained in 47.7.3절. "The Role of Policy in the Boot Process" . Ultimately, every system operation is determined by the policy and the type-labeling of the files.

[4] http://www.commoncriteriaportal.org/files/ppfiles/lspp.pdf

Important

After loading a new policy, it is recommended that you restart any services that may have new or changed labeling. Generally speaking, this is only the targeted daemons, as listed in 47.8.1절. “What is the Targeted Policy?” .

47.7.1.2. SELinux and Mandatory Access Control

SELinux is an implementation of Mandatory Access Control (MAC). Depending on the security policy type, SELinux implements either Type Enforcement (TE), Roles Based Access Control (RBAC) or Bell-La Padula Model Multi-Level Security (MLS).

The policy specifies the rules in the implemented environment. It is written in a language created specifically for writing security policy. Policy writers use m4 macros to capture common sets of low-level rules. A number of m4 macros are defined in the existing policy, which facilitate the writing of new policy. These rules are preprocessed into many additional rules as part of building the policy.conf file, which is compiled into the binary policy.

Access rights are divided differently among domains, and no domain is required to act as a master for all other domains. Moving between domains is controlled by the policy, through login programs, userspace programs such as newrole, or by requiring a new process execution in the new domain. This movement between domains is referred to as a transition.

47.7.2. Where is the Policy?

There are two components to the policy: the binary tree and the source tree. The binary tree is provided by the selinux-policy-<policyname> package and supplies the binary policy file.

Alternatively, the binary policy can be built from source when the selinux-policy-devel package is installed.

Note

Information on how to edit, write and compile policy is currently outside the scope of this document.

47.7.2.1. Binary Tree Files

- /etc/selinux/targeted/ — this is the root directory for the targeted policy, and contains the binary tree.
- /etc/selinux/targeted/policy/ — this is the location of the binary policy file policy.<xx>. In this guide, the variable SELINUX_POLICY is used for this directory.
- /etc/selinux/targeted/contexts/ — this is the location of the security context information and configuration files, which are used during runtime by various applications.
- /etc/selinux/targeted/contexts/files/ — contains the default contexts for the entire file system. This is referenced by restorecon when performing relabeling operations.

- /etc/selinux/targeted/contexts/users/ — in the targeted policy, only the root file is in this directory. These files are used for determining context when a user logs in. For example, for the root user, the context is user_u:system_r:unconfined_t.

- /etc/selinux/targeted/modules/active/booleans* — this is where the runtime Booleans are configured.

> Note
>
> These files should never be manually changed. You should use the getsebool, setsebool and semanage tools to manipulate runtime Booleans.

47.7.2.2. Source Tree Files

For developing policy modules, the selinux-policy-devel package includes all of the interface files used to build policy. It is recommended that people who build policy use these files to build the policy modules.

This package installs the policy interface files under /usr/share/selinux/devel/include and has make files installed in /usr/share/selinux/devel/Makefile.

To help applications that need the various SELinux paths, libselinux provides a number of functions that return the paths to the different configuration files and directories. This negates the need for applications to hard-code the paths, especially since the active policy location is dependent on the SELINUXTYPE setting in /etc/selinux/config.

For example, if SELINUXTYPE is set to strict, the active policy location is under /etc/selinux/strict.

To view the list of available functions, use the following command:

```
man 3 selinux_binary_policy_path
```

> Note
>
> This man page is available only if you have the libselinux-devel RPM installed.
>
> The use of libselinux and related functions is outside the scope of this document.

47.7.3. The Role of Policy in the Boot Process

SELinux plays an important role during the early stages of system start-up. Because all processes must be labeled with their correct domain, init performs some essential operations early in the boot process to maintain synchronization between labeling and policy enforcement.

1. After the kernel has been loaded during the boot process, the initial process is assigned the predefined initial SELinux ID (initial SID) kernel. Initial SIDs are used for bootstrapping before the policy is loaded.

2. /sbin/init mounts /proc/, and then searches for the selinuxfs file system type. If it is present, that means SELinux is enabled in the kernel.

3. If init does not find SELinux in the kernel, or if it is disabled via the selinux=0 boot parameter, or if /etc/selinux/config specifies that SELINUX=disabled, the boot process proceeds with a non-SELinux system.

   At the same time, init sets the enforcing status if it is different from the setting in /etc/selinux/config. This happens when a parameter is passed during the boot process, such as enforcing=0 or enforcing=1. The kernel does not enforce any policy until the initial policy is loaded.

4. If SELinux is present, /selinux/ is mounted.

5. init checks /selinux/policyvers for the supported policy version. The version number in /selinux/policyvers is the latest policy version your kernel supports. init inspects /etc/selinux/config to determine which policy is active, such as the targeted policy, and loads the associated file at $SELINUX_POLICY/policy.<version>.

   If the binary policy is not the version supported by the kernel, init attempts to load the policy file if it is a previous version. This provides backward compatibility with older policy versions.

   If the local settings in /etc/selinux/targeted/booleans are different from those compiled in the policy, init modifies the policy in memory based on the local settings prior to loading the policy into the kernel.

6. By this stage of the process, the policy is fully loaded into the kernel. The initial SIDs are then mapped to security contexts in the policy. In the case of the targeted policy, the new domain is user_u:system_r:unconfined_t. The kernel can now begin to retrieve security contexts dynamically from the in-kernel security server.

7. init then re-executes itself so that it can transition to a different domain, if the policy defines it. For the targeted policy, there is no transition defined and init remains in the unconfined_t domain.

8. At this point, init continues with its normal boot process.

The reason that init re-executes itself is to accommodate stricter SELinux policy controls. The objective of re-execution is to transition to a new domain with its own granular rules. The only way that a process can enter a domain is during execution, which means that such processes are the only entry points into the domains.

For example, if the policy has a specific domain for init, such as init_t, a method is required to change from the initial SID, such as kernel, to the correct runtime domain for init. Because this transition may need to occur, init is coded to re-execute itself after loading the policy.

This init transition occurs if the domain_auto_trans(kernel_t, init_exec_t, <target_domain_t>) rule is present in the policy. This rule states that an automatic transition occurs on anything executing in the kernel_t domain that executes a file of type init_exec_t. When this execution occurs, the new process is assigned the domain <target_domain_t>, using an actual target domain such as init_t.

47.7.4. Object Classes and Permissions

SELinux defines a number of classes for objects, making it easier to group certain permissions by specific classes. For example:

- File-related classes include filesystem for file systems, file for files, and dir for directories. Each class has its own associated set of permissions.

  The filesystem class can mount, unmount, get attributes, set quotas, relabel, and so forth. The file class has common file permissions such as read, write, get and set attributes, lock, relabel, link, rename, append, etc.

- Network related classes include tcp_socket for TCP sockets, netif for network interfaces, and node for network nodes.

  The netif class, for example, can send and receive on TCP, UDP and raw sockets (tcp_recv, tcp_send, udp_send, udp_recv, rawip_recv, and rawip_send.)

The object classes have matching declarations in the kernel, meaning that it is not trivial to add or change object class details. The same is true for permissions. Development work is ongoing to make it possible to dynamically register and unregister classes and permissions.

Permissions are the actions that a subject can perform on an object, if the policy allows it. These permissions are the access requests that SELinux actively allows or denies.

47.8. Targeted Policy Overview

This chapter is an overview and examination of the SELinux targeted policy, the current supported policy for Red Hat Enterprise Linux.

Much of the content in this chapter is applicable to all types of SELinux policy, in terms of file locations and the type of content in those files. The difference lies in which files exist in the key locations and their contents.

47.8.1. What is the Targeted Policy?

The SELinux policy is highly configurable. For Red Hat Enterprise Linux 5, Red Hat supports a single policy, the targeted policy. Under the targeted policy, every subject and object runs in the unconfined_t domain except for the specific targeted daemons. Objects that are in the unconfined_t domain have no restrictions and fall back to using standard Linux security, that is, DAC. The daemons that are part of the targeted policy run in their own domains and are restricted in every operation they perform on the system. This way daemons that are exploited or compromised in any way are contained and can only cause limited damage.

For example, the http and ntp daemons are both protected in the default targeted policy, and run in the httpd_t and ntpd_t domains, respectively. The ssh daemon, however, is not protected in this policy, and consequently runs in the unconfined_t domain.

Refer to the following sample output, which illustrates the various domains for the daemons mentioned above:

```
user_u:system_r:httpd_t          25129 ?        00:00:00 httpd
user_u:system_r:ntpd_t           25176 ?        00:00:00 ntpd
system_u:system_r:unconfined_t          25245 ? 00:00:00 sshd
```

The Strict Policy

The opposite of the targeted policy is the strict policy. In the strict policy, every subject and object exists in a specific security domain, and all interactions and transitions are individually considered within the policy rules.

The strict policy is a much more complex environment, and does not ship with Red Hat Enterprise Linux. This guide focuses on the targeted policy that ships with Red Hat Enterprise Linux, and the components of SELinux used by the targeted daemons.

The targeted daemons are as follows: dhcpd; httpd; mysqld; named; nscd; ntpd; portmap; postgres; snmpd; squid; syslogd; and winbind.

Note

Depending on your installation, only some of these daemons may be present.

47.8.2. Files and Directories of the Targeted Policy

Refer to 47.7.2절. "Where is the Policy?" for a list of the common files and directories used by SELinux.

47.8.3. Understanding the Users and Roles in the Targeted Policy

This section covers the specific roles enabled for the targeted policy. The unconfined_t type exists in every role, which significantly reduces the usefulness of roles in the targeted policy. More extensive use of roles requires a change to the strict policy paradigm, where every process runs in an individually considered domain.

Effectively, there are only two roles in the targeted policy: system_r and object_r. The initial role is system_r, and everything else inherits that role. The remaining roles are defined for compatibility purposes between the targeted policy and the strict policy.[5]

Three of the four roles are defined by the policy. The fourth role, object_r, is an implied role and is not found in policy source. Because roles are created and populated by types using one or more declarations in the policy, there is no single file that declares all roles. (Remember that the policy itself is generated from a number of separate files.)

system_r
: This role is for all system processes except user processes:

```
system_r (28 types)
    dhcpd_t
    httpd_helper_t
    httpd_php_t
    httpd_suexec_t
    httpd_sys_script_t
    httpd_t
```

[5] Any role could have been chosen for the targeted policy, but system_r already had existing authorization for the daemon domains, simplifying the process. This was done because no mechanism currently exists to alias roles.

```
        httpd_unconfined_script_t
        initrc_t
        ldconfig_t
        mailman_cgi_t
        mailman_mail_t
        mailman_queue_t
        mysqld_t
        named_t
        ndc_t
        nscd_t
        ntpd_t
        pegasus_t
        portmap_t
        postgresql_t
        snmpd_t
        squid_t
        syslogd_t
        system_mail_t
        unconfined_t
        winbind_helper_t
        winbind_t
        ypbind_t
```

user_r

This is the default user role for regular Linux users. In a strict policy, individual users might be used, allowing for the users to have special roles to perform privileged operations. In the targeted policy, all users run in the unconfined_t domain.

object_r

In SELinux, roles are not utilized for objects when RBAC is being used. Roles are strictly for subjects. This is because roles are task-oriented and they group together entities which perform actions (for example, processes). All such entities are collectively referred to as subjects. For this reason, all objects have the role object_r, and the role is only used as a placeholder in the label.

sysadm_r

This is the system administrator role in a strict policy. If you log in directly as the root user, the default role may actually be staff_r. If this is true, use the newrole -r sysadm_r command to change to the SELinux system administrator role to perform system administration tasks. In the targeted policy, the following retain sysadm_r for compatibility:

```
sysadm_r (6 types)
        httpd_helper_t
        httpd_sys_script_t
        initrc_t
        ldconfig_t
        ndc_t
        unconfined_t
```

There is effectively only one user identity in the targeted policy. The user_u identity was chosen because libselinux falls back to user_u as the default SELinux user identity. This occurs when there is no matching SELinux user for the Linux user who is logging in. Using user_u as the single user in the targeted policy makes it easier to change to the strict policy. The remaining users exist for compatibility with the strict policy.[6]

[6] A user aliasing mechanism would also work here, to alias all identities from the strict policy to a single user identity in the targeted policy.

The one exception is the SELinux user root. You may notice root as the user identity in a process's context. This occurs when the SELinux user root starts daemons from the command line, or restarts a daemon originally started by init.

SELinux를 사용하여 작업하기

SELinux는 시스템 관리자 및 최종 사용자에게 새로운 보안 패러다임 및 실행 도구를 제공합니다. 이 장에서 다루게 될 도구 및 기술 사항은 최종 사용자, 관리자, 분석가에 의해 실행되는 표준 운용에 중점을 둡니다.

48.1. End User Control of SELinux

In general, end users have little interaction with SELinux when Red Hat Enterprise Linux is running the targeted policy. This is because users are running in the domain of unconfined_t along with the rest of the system except the targeted daemons.

In most situations, standard DAC controls prevent you from performing tasks for which you do not have the required access or permissions before SELinux is consulted. Consequently, it is likely that you will never generate an avc: denied message.

The following sections cover the general tasks and practices that an end user might need to perform on a Red Hat Enterprise Linux system. These tasks apply to users of all privilege levels, not only to end users.

48.1.1. Moving and Copying Files

In file system operations, security context must now be considered in terms of the label of the file, the process accessing it, and the directories where the operation is happening. Because of this, moving and copying files with mv and cp may have unexpected results.

Copying Files: SELinux Options for cp

Unless you specify otherwise, cp follows the default behavior of creating a new file based on the domain of the creating process and the type of the target directory. Unless there is a specific rule to set the label, the file inherits the type from the target directory.

Use the -Z user:role:type option to specify the required label for the new file.

The -p (or --preserve=mode,ownership,timestamps) option preserves the specified attributes and, if possible, additional attributes such as links.

```
touch bar foo
ls -Z bar foo
-rw-rw-r--  auser   auser   user_u:object_r:user_home_t   bar
-rw-rw-r--  auser   auser   user_u:object_r:user_home_t   foo
```

If you use the cp command without any additional command-line arguments, a copy of the file is created in the new location using the default type of the creating process and the target directory. In this case, because there is no specific rule that applies to cp and /tmp, the new file has the type of the parent directory:

```
cp bar /tmp
ls -Z /tmp/bar
-rw-rw-r--  auser   auser   user_u:object_r:tmp_t   /tmp/bar
```

The type tmp_t is the default type for temporary files.

Use the -Z option to specify the label for the new file:

```
cp -Z user_u:object_r:user_home_t foo /tmp
ls -Z /tmp/foo
-rw-rw-r--  auser   auser   user_u:object_r:user_home_t   /tmp/foo
```

Moving Files: SELinux Options for mv

Moving files with mv retains the original type associated with the file. Care should be taken using this command as it can cause problems. For example, if you move files with the type user_home_t into ~/public_html, then the httpd daemon is not able to serve those files until you relabel them. Refer to 48.1.3절. "Relabeling a File or Directory" for more information about file labeling.

표 48.1. Behavior of mv and cp Commands

Command	Behavior
mv	The file retains its original label. This may cause problems, confusion, or minor insecurity. For example, the tmpwatch program running in the sbin_t domain might not be allowed to delete an aged file in the /tmp directory because of the file's type.
cp	Makes a copy of the file using the default behavior based on the domain of the creating process (cp) and the type of the target directory.
cp -p	Makes a copy of the file, preserving the specified attributes and security contexts, if possible. The default attributes are mode, ownership, and timestamps. Additional attributes are links and all.
cp -Z <user:role:type>	Makes a copy of the file with the specified labels. The -Z option is synonymous with --context.

48.1.2. Checking the Security Context of a Process, User, or File Object

Checking a Process ID

In Red Hat Enterprise Linux, the -Z option is equivalent to --context, and can be used with the ps, id, ls, and cp commands. The behavior of the cp command with respect to SELinux is explained in 표 48.1. "Behavior of mv and cp Commands".

The following example shows a small sample of the output of the ps command. Most of the processes are running in the unconfined_t domain, with a few exceptions.

```
[user@localhost ~]$ ps auxZ
LABEL                           USER       PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME
 COMMAND
system_u:system_r:init_t       root          1  0.0  0.1   2032   620 ?        Ss   15:09   0:00 init [5]
system_u:system_r:kernel_t      root          2  0.0  0.0      0     0 ?        S    15:09   0:00 [migration/0]
system_u:system_r:kernel_t      root          3  0.0  0.0      0     0 ?        SN   15:09   0:00 [ksoftirqd/0]

user_u:system_r:unconfined_t    user      3122  0.0  0.6   6908  3232 ?        S    16:47   0:01 /usr/libexec/gconfd-2 5
user_u:system_r:unconfined_t    user      3125  0.0  0.1   2540   588 ?        S    16:47   0:00 /usr/bin/gnome-keyring-
daemon
```

```
user_u:system_r:unconfined_t     user      3127  0.0   1.4   33612   6988 ?         Sl    16:47    0:00 /usr/libexec/gnome-
settings-daemon
user_u:system_r:unconfined_t     user      3144  0.1   1.4   16528   7360 ?         Ss    16:47    0:01 metacity --sm-client-
id=default1
user_u:system_r:unconfined_t     user      3148  0.2   2.9   79544  14808 ?         Ss    16:47    0:03 gnome-panel --sm-client-
id default2
```

Checking a User ID

You can use the -Z option with the id command to determine a user's security context. Note that with this command you cannot combine -Z with other options.

```
[root@localhost ~]# id -Z
user_u:system_r:unconfined_t
```

Note that you cannot use the -Z option with the id command to inspect the security context of a different user. That is, you can only display the security context of the currently logged-in user:

```
[user@localhost ~]$ id
uid=501(user) gid=501(user) groups=501(user) context=user_u:system_r:unconfined_t
[user@localhost ~]$ id root
uid=0(root) gid=0(root) groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
[user@localhost ~]$ id -Z root
id: cannot display context when selinux not enabled or when displaying the id
of a different user
```

Check a File ID

You can use the -Z option with the ls command to group common long-format information. You can display mode, user, group, security context, and filename information.

```
cd /etc
ls -Z h* -d
drwxr-xr-x  root root  system_u:object_r:etc_t           hal
-rw-r--r--  root root  system_u:object_r:etc_t           host.conf
-rw-r--r--  root root  user_u:object_r:etc_t             hosts
-rw-r--r--  root root  system_u:object_r:etc_t           hosts.allow
-rw-r--r--  root root  system_u:object_r:etc_t           hosts.canna
-rw-r--r--  root root  system_u:object_r:etc_t           hosts.deny
drwxr-xr-x  root root  system_u:object_r:hotplug_etc_t   hotplug
drwxr-xr-x  root root  system_u:object_r:etc_t           hotplug.d
drwxr-xr-x  root root  system_u:object_r:httpd_sys_content_t htdig
drwxr-xr-x  root root  system_u:object_r:httpd_config_t  httpd
```

48.1.3. Relabeling a File or Directory

You may need to relabel a file when moving or copying into special directories related to the targeted daemons, such as ~/public_html directories, or when writing scripts that work in directories outside of /home.

There are two general types of relabeling operations:

- Deliberately changing the type of a file

- Restoring files to the default state according to policy

There are also relabeling operations that an administrator performs. These are covered in 48.2.2절. “Relabeling a File System” .

Tip

The majority of SELinux permission control in the targeted policy is Type Enforcement (TE). Consequently, you can generally ignore the user and role information in a security label and focus on just changing the type. You do not normally need to consider the role and user settings on files.

Note

If relabeling affects the label on a daemon's executable, you should restart the daemon to be sure it is running in the correct domain. For example, if /usr/sbin/mysqld has the wrong security label, and you address this by using a relabeling operation such as restorecon, you must restart mysqld after the relabeling operation. Setting the executable file to have the correct type (mysqld_exec_t) ensures that it transitions to the proper domain when started.

Use the chcon command to change a file to the correct type. You need to know the correct type that you want to apply to use this command. The directories and files in the following example are labeled with the default type defined for file system objects created in /home:

```
cd ~
ls -Zd public_html/
drwxrwxr-x  auser  auser  user_u:object_r:user_home_t public_html/

ls -Z web_files/
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    1.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    2.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    3.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    4.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    5.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    index.html
```

If you move these files into the public_html directory, they retain the original type:

```
mv web_files/* public_html/
ls -Z public_html/
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    1.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    2.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    3.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    4.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    5.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t    index.html
```

To make these files viewable from a special user public HTML folder, they need to have a type that httpd has permissions to read, presuming the Apache HTTP Server is configured for UserDir and the Boolean value httpd_enable_homedirs is enabled.

```
chcon -R -t httpd_user_content_t public_html/
ls -Z public_html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t   1.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t   2.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t   3.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t   4.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t   5.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t   index.html

ls -Z public_html/ -d
drwxrwxr-x  auser  auser  user_u:object_r:httpd_user_content_t  public_html/
```

Tip

If the file has no label, such as a file created while SELinux was disabled in the kernel, you need to give it a full label with chcon system_u:object_r:shlib_t foo.so. Otherwise, you will receive an error about applying a partial context to an unlabeled file.

Use the restorecon command to restore files to the default values according to the policy. There are two other methods for performing this operation that work on the entire file system: fixfiles or a policy relabeling operation. Each of these methods requires superuser privileges. Cautions against both of these methods appear in 48.2.2절. “Relabeling a File System” .

The following example demonstrates restoring the default user home directory context to a set of files that have different types. The first two sets of files have different types, and are being moved into a directory for archiving. Their contexts are different from each other, and are incorrect for a standard user's home directory:

```
ls -Z /tmp/
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t              /tmp/file1
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t              /tmp/file2
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t              /tmp/file3

mv /tmp/{1,2,3} archives/
mv public_html/* archives/
ls -Z archives/
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t              file1
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t    file1.html
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t              file2
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t    file2.html
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t              file3
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t    file3.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t    file4.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t    file5.html
-rw-rw-r--  auser  auser  user_u:object_r:httpd_user_content_t  index.html
```

The archives/ directory already has the default type because it was created in the user's home directory:

```
ls -Zd archives/
drwxrwxr-x  auser  auser  user_u:object_r:user_home_t  archives/
```

Using the restorecon command to relabel the files uses the default file contexts set by the policy, so these files are labeled with the default label for their current directory.

```
/sbin/restorecon -R archives/
ls -Z archives/
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file1
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file1.html
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file2
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file2.html
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file3
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file3.html
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file4.html
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    file5.html
-rw-rw-r--  auser  auser  system_u:object_r:user_home_t    index.html
```

48.1.4. Creating Archives That Retain Security Contexts

You can use either the tar or star utilities to create archives that retain SELinux security contexts. The following example uses star to demonstrate how to create such an archive. You need to use the appropriate -xattr and -H=exustar options to ensure that the extra attributes are captured and that the header for the *.star file is of a type that fully supports xattrs. Refer to the man page for more information about these and other options.

The following example illustrates the creation and extraction of a set of html files and directories. Note that the two directories have different labels. Unimportant parts of the file context have been omitted for printing purposes (indicated by ellipses '...'):

```
ls -Z public_html/ web_files/

public_html/:
-rw-rw-r--  auser  auser  ...httpd_user_content_t 1.html
-rw-rw-r--  auser  auser  ...httpd_user_content_t 2.html
-rw-rw-r--  auser  auser  ...httpd_user_content_t 3.html
-rw-rw-r--  auser  auser  ...httpd_user_content_t 4.html
-rw-rw-r--  auser  auser  ...httpd_user_content_t 5.html
-rw-rw-r--  auser  auser  ...httpd_user_content_t index.html
web_files/:
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  1.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  2.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  3.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  4.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  5.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  index.html
```

The following command creates the archive, retaining all of the SELinux security contexts:

```
star -xattr -H=exustar -c -f all_web.star public_html/ web_files/
star: 11 blocks + 0 bytes (total of 112640 bytes = 110.00k).
```

Use the ls command with the -Z option to validate the security context:

```
ls -Z all_web.star
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t \  all_web.star
```

You can now copy the archive to a different directory. In this example, the archive is copied to /tmp. If there is no specific policy to make a derivative temporary type, the default behavior is to acquire the tmp_t type.

```
cp all_web.star /tmp/ cd /tmp/

ls -Z all_web.star
-rw-rw-r--  auser  auser  user_u:object_r:tmp_t  all_web.star
```

Now you can expand the archives using star and it restores the extended attributes:

```
star -xattr -x -f all_web.star
star: 11 blocks + 0 bytes (total of 112640 bytes = 110.00k).

ls -Z /tmp/public_html/ /tmp/web_files/
/tmp/public_html/:
-rw-rw-r--  auser  auser  ...httpd_sys_content_t 1.html
-rw-rw-r--  auser  auser  ...httpd_sys_content_t 2.html
-rw-rw-r--  auser  auser  ...httpd_sys_content_t 3.html
-rw-rw-r--  auser  auser  ...httpd_sys_content_t 4.html
-rw-rw-r--  auser  auser  ...httpd_sys_content_t 5.html
-rw-rw-r--  auser  auser  ...httpd_sys_content_t index.html
/tmp/web_files/:
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  1.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  2.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  3.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  4.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  5.html
-rw-rw-r--  auser  auser  user_u:object_r:user_home_t  \ index.html
```

Caution

If you use an absolute path when you create an archive using star, the archive expands on that same path. For example, an archive made with this command restores the files to /var/log/httpd/:

```
star -xattr -H=exustar -c -f httpd_logs.star /var/log/httpd/
```

If you attempt to expand this archive, star issues a warning if the files in the path are newer than the ones in the archive.

48.2. Administrator Control of SELinux

In addition to the tasks often performed by users in 48.1절. "End User Control of SELinux", SELinux administrators could be expected to perform a number of additional tasks. These tasks typically require root access to the system. Such tasks are significantly easier under the targeted policy. For example, there is no need to consider adding, editing, or deleting Linux users from the SELinux users, nor do you need to consider roles.

This section covers the types of tasks required of an administrator who maintains Red Hat Enterprise Linux running SELinux.

48.2.1. Viewing the Status of SELinux

The sestatus command provides a configurable view into the status of SELinux. The simplest form of this command shows the following information:

```
~]# sestatus
SELinux status:                 enabled
SELinuxfs mount:                /selinux
Current mode:                   enforcing
Mode from config file:          enforcing
Policy version:                 21
Policy from config file:        targeted
```

The -v option includes information about the security contexts of a series of files that are specified in /etc/sestatus.conf:

```
~]# sestatus -v
SELinux status:                 enabled
SELinuxfs mount:                /selinux
Current mode:                   enforcing
Mode from config file:          enforcing
Policy version:                 21
Policy from config file:        targeted

Process contexts:
Current context:                user_u:system_r:unconfined_t
Init context:                   system_u:system_r:init_t
/sbin/mingetty                  system_u:system_r:getty_t
/usr/sbin/sshd                  system_u:system_r:unconfined_t:s0-s0:c0.c1023

File contexts:
Controlling term:               user_u:object_r:devpts_t
/etc/passwd                     system_u:object_r:etc_t
/etc/shadow                     system_u:object_r:shadow_t
/bin/bash                       system_u:object_r:shell_exec_t
/bin/login                      system_u:object_r:login_exec_t
/bin/sh                         system_u:object_r:bin_t -> system_u:object_r:shell_exec_t
/sbin/agetty                    system_u:object_r:getty_exec_t
/sbin/init                      system_u:object_r:init_exec_t
/sbin/mingetty                  system_u:object_r:getty_exec_t
/usr/sbin/sshd                  system_u:object_r:sshd_exec_t
/lib/libc.so.6                  system_u:object_r:lib_t -> system_u:object_r:lib_t
/lib/ld-linux.so.2              system_u:object_r:lib_t -> system_u:object_r:ld_so_t
```

The -b displays the current state of booleans. You can use this in combination with grep or other tools to determine the status of particular booleans:

```
~]# sestatus -b | grep httpd | grep on$
httpd_builtin_scripting         on
httpd_disable_trans             on
httpd_enable_cgi                on
httpd_enable_homedirs           on
httpd_unified                   on
```

48.2.2. Relabeling a File System

You may never need to relabel an entire file system. This usually occurs only when labeling a file system for SELinux for the first time, or when switching between different types of policy, such as changing from the targeted to the strict policy.

Relabeling a File System Using init

The recommended method for relabeling a file system is to reboot the machine. This allows the init process to perform the relabeling, ensuring that applications have the correct labels when they are started and that they are started in the right order. If you relabel a file system without rebooting, some processes may continue running with an incorrect context. Manually ensuring that all the daemons are restarted and running in the correct context can be difficult.

Use the following procedure to relabel a file system using this method.

```
touch /.autorelabel
reboot
```

At boot time, init.rc checks for the existence of /.autorelabel. If this file exists, SELinux performs a complete file system relabel (using the /sbin/fixfiles -f -F relabel command), and then deletes /.autorelabel.

Relabeling a File System Using fixfiles

It is possible to relabel a file system using the fixfiles command, or to relabel based on the RPM database:

Use the following command to relabel a file system only using the fixfiles command:

```
fixfiles relabel
```

Use the following command to relabel a file system based on the RPM database:

```
fixfiles -R <packagename> restore
```

Using fixfiles to restore contexts from packages is safer and quicker.

Caution

Running fixfiles on the entire file system without rebooting may make the system unstable.

If the relabeling operation applies a new policy that is different from the policy that was in place when the system booted, existing processes may be running in incorrect and insecure domains. For example, a process could be in a domain that is not an allowed transition for that process in the new policy, granting unexpected permissions to that process alone.

In addition, one of the options to fixfiles relabel prompts for approval to empty /tmp/ because it is not possible to reliably relabel /tmp/. Since fixfiles is run as root, temporary files that applications are relying upon are erased. This could make the system unstable or behave unexpectedly.

48.2.3. Managing NFS Home Directories

In Red Hat Enterprise Linux 5, most targeted daemons do not interact with user data and are not affected by NFS-mounted home directories. One exception is the Apache HTTP Server. For example, CGI scripts that are on the mounted file system have the nfs_t type, which is not a type that httpd_t is allowed to execute.

If you are having problems with the default type of nfs_t, try mounting the home directories with a different context:

```
mount -t nfs -o context=user_u:object_r:user_home_dir_t \
 fileserver.example.com:/shared/homes/ /home
```

Caution

48.2.9절. "Specifying the Security Context of Entire File Systems" explains how to mount a directory so that httpd can execute scripts. If you do this for user home directories, it gives the Apache HTTP Server increased access to those directories. Remember that a mountpoint label applies to the entire mounted file system.

Future versions of the SELinux policy address the functionality of NFS.

48.2.4. Granting Access to a Directory or a Tree

Similar to standard Linux DAC permissions, a targeted daemon must have SELinux permissions to be able to descend the directory tree. This does not mean that a directory and its contents need to have the same type. There are many types, such as root_t, tmp_t, and usr_t that grant read access for a directory. These types are suitable for directories that do not contain any confidential information, and that you want to be widely readable. They could also be used for a parent directory of more secured directories with different contexts.

If you are working with an avc: denied message, there are some common problems that arise with directory traversal. For example, many programs run a command equivalent to ls -l / that is not necessary to their operation but generates a denial message in the logs. For this you need to create a dontaudit rule in your local.te file.

When trying to interpret AVC denial messages, do not be misled by the path=/ component. This path is not related to the label for the root file system, /. It is actually relative to the root of the file system on the device node. For example, if your /var/ directory is located on an LVM (Logical Volume Management [1]) device, /dev/dm-0, the device node is identified in the message as dev=dm-0. When you see path=/ in this example, that is the top level of the LVM device dm-0, not necessarily the same as the root file system designation /.

48.2.5. Backing Up and Restoring the System

Refer to the explanation in 48.1.4절. "Creating Archives That Retain Security Contexts" .

48.2.6. Enabling or Disabling Enforcement

[1] LVM is the grouping of physical storage into virtual pools that are partitioned into logical volumes.

You can enable and disable SELinux enforcement at runtime or configure it to start in the correct mode at boot time, using the command line or GUI. SELinux can operate in one of three modes: disabled, meaning not enabled in the kernel; permissive, meaning SELinux is running and logging but not controlling permissions; or enforcing, meaning SELinux is running and enforcing policy.

Use the setenforce command to change between permissive and enforcing modes at runtime. Use setenforce 0 to enter permissive mode; use setenforce 1 to enter enforcing mode.

The sestatus command displays the current mode and the mode from the configuration file referenced during boot:

```
~]# sestatus | grep -i mode
Current mode:           permissive
Mode from config file:  permissive
```

Note that changing the runtime enforcement does not affect the boot time configuration:

```
~]# setenforce 1
~]# sestatus | grep -i mode
Current mode:           enforcing
Mode from config file:  permissive
```

You can also disable enforcing mode for a single daemon. For example, if you are trying to troubleshoot the named daemon and SELinux, you can turn off enforcing for just that daemon.

Use the getsebool command to get the current status of the boolean:

```
~]# getsebool named_disable_trans
named_disable_trans --> off
```

Use the following command to disable enforcing mode for this daemon:

```
~]# setsebool named_disable_trans 1
~]# getsebool named_disable_trans
named_disable_trans --> on
```

Note

This sets the runtime value only. Use the -P option to make the change persistent across reboots.

Any *_disable_trans booleans that are set to "on" invoke the conditional that prevents the process from transitioning to the domain on execution.

Use the following command to find which of these booleans are set:

```
~]# getsebool -a | grep disable.*on
httpd_disable_trans=1
mysqld_disable_trans=1
ntpd_disable_trans=1
```

You can set any number of boolean values using the setsebool command:

```
setsebool -P httpd_disable_trans=1 mysqld_disable_trans=1 ntpd_disable_trans=1
```

You can also use togglesebool <boolean_name> to change the value of a specific boolean:

```
~]# getsebool httpd_disable_trans
httpd_disable_trans --> off
~]# togglesebool httpd_disable_trans
httpd_disable_trans: active
```

You can configure all of these settings using system-config-selinux. The same configuration files are used, so changes appear bidirectionally.

Changing a Runtime Boolean

Use the following procedure to change a runtime boolean using the GUI.

> Note
>
> Administrator privileges are required to perform this procedure.

1. On the System menu, point to Administration and then click Security Level and Firewall to display the Security Level Configuration dialog box.

2. Click the SELinux tab, and then click Modify SELinux Policy.

3. In the selection list, click the arrow next to the Name Service entry, and select the Disable SELinux protection for named daemon check box.

4. Click OK to apply the change. Note that it may take a short time for the policy to be reloaded.

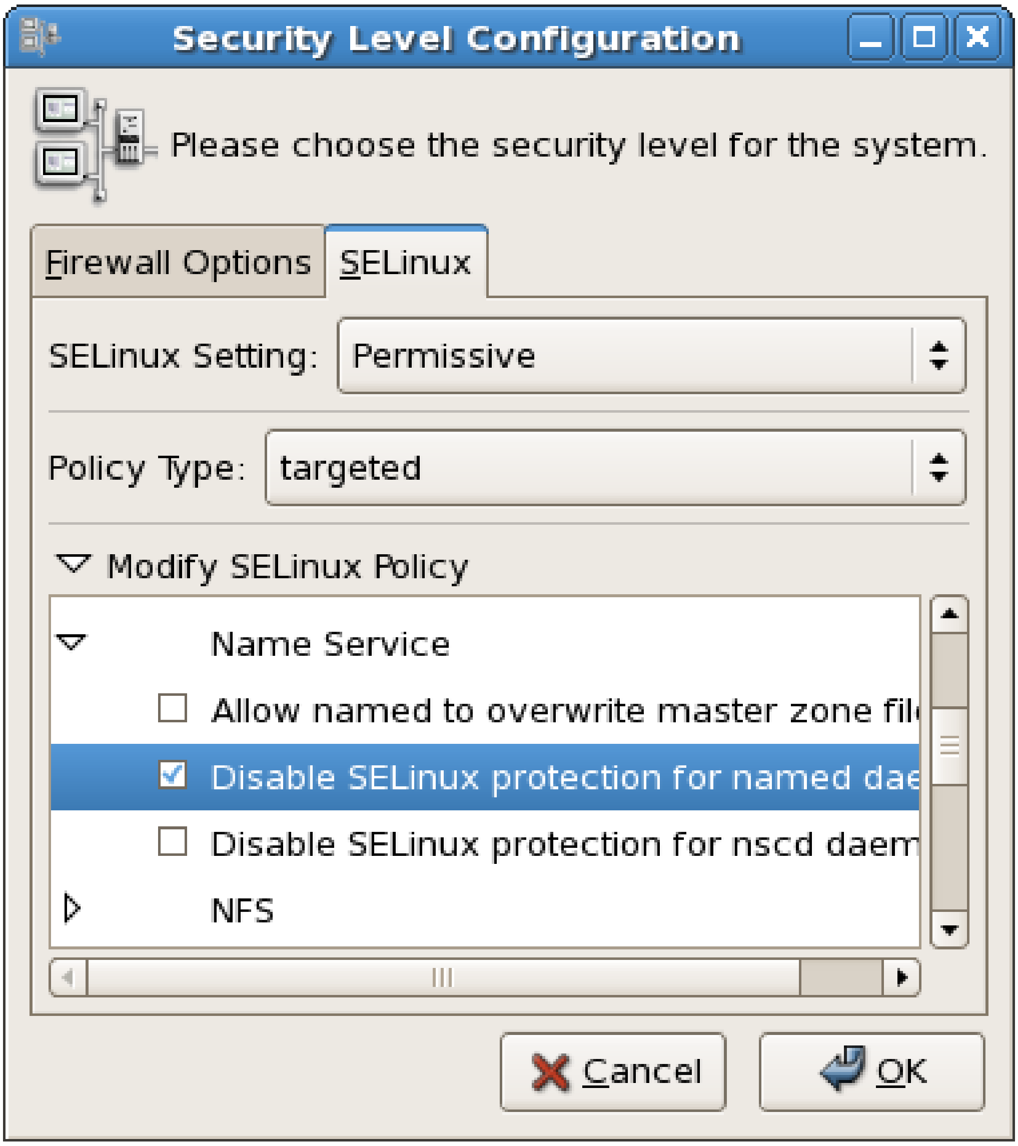

Please choose the security level for the system.

SELinux Setting: Permissive

Policy Type: targeted

□ Allow named to overwrite master zone fil

☑ Disable SELinux protection for named dae

□ Disable SELinux protection for nscd daem

그림 48.1. Using the Security Level Configuration dialog box to change a runtime boolean.

If you want to control these settings with scripts, you can use the setenforce(1), getenforce(1), and selinuxenabled(1) commands.

48.2.7. Enable or Disable SELinux

Important

Changes you make to files while SELinux is disabled may give them an unexpected security label, and new files will not have a label. You may need to relabel part or all of the file system after re-enabling SELinux.

From the command line, you can edit the /etc/sysconfig/selinux file. This file is a symlink to /etc/selinux/config. The configuration file is self-explanatory. Changing the value of SELINUX or SELINUXTYPE changes the state of SELinux and the name of the policy to be used the next time the system boots.

```
~]# cat /etc/sysconfig/selinux
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#        enforcing - SELinux security policy is enforced.
#        permissive - SELinux prints warnings instead of enforcing.
#        disabled - SELinux is fully disabled.
SELINUX=permissive
# SELINUXTYPE= type of policy in use. Possible values are:
#        targeted - Only targeted network daemons are protected.
#        strict - Full SELinux protection.
SELINUXTYPE=targeted

# SETLOCALDEFS= Check local definition changes
SETLOCALDEFS=0
```

Changing the Mode of SELinux Using the GUI

Use the following procedure to change the mode of SELinux using the GUI.

> **Note**
>
> You need administrator privileges to perform this procedure.

1. On the System menu, point to Administration and then click Security Level and Firewall to display the Security Level Configuration dialog box.

2. Click the SELinux tab.

3. In the SELinux Setting select either Disabled, Enforcing or Permissive, and then click OK.

4. If you changed from Enabled to Disabled or vice versa, you need to restart the machine for the change to take effect.

Changes made using this dialog box are immediately reflected in /etc/sysconfig/selinux.

48.2.8. Changing the Policy

This section provides a brief introduction to using customized policies on your system. A full discussion of this topic is beyond the scope of this document.

To load a different policy on your system, change the following line in /etc/sysconfig/selinux:

```
SELINUXTYPE=<policyname>
```

where <policyname> is the policy name directory under /etc/selinux/. This assumes that you have the custom policy installed. After changing the SELINUXTYPE parameter, run the following commands:

```
touch /.autorelabel
reboot
```

Use the following procedure to load a different policy using the system-config-selinux utility:

Note

You need administrator privileges to perform this procedure.

1. Ensure that the complete directory structure for the required policy exists under /etc/selinux.

2. On the System menu, point to Administration and then click Security Level and Firewall to display the Security Level Configuration dialog box.

3. Click the SELinux tab.

4. In the Policy Type list, select the policy that you want to load, and then click OK. This list is only visible if more than one policy is installed.

5. Restart the machine for the change to take effect.

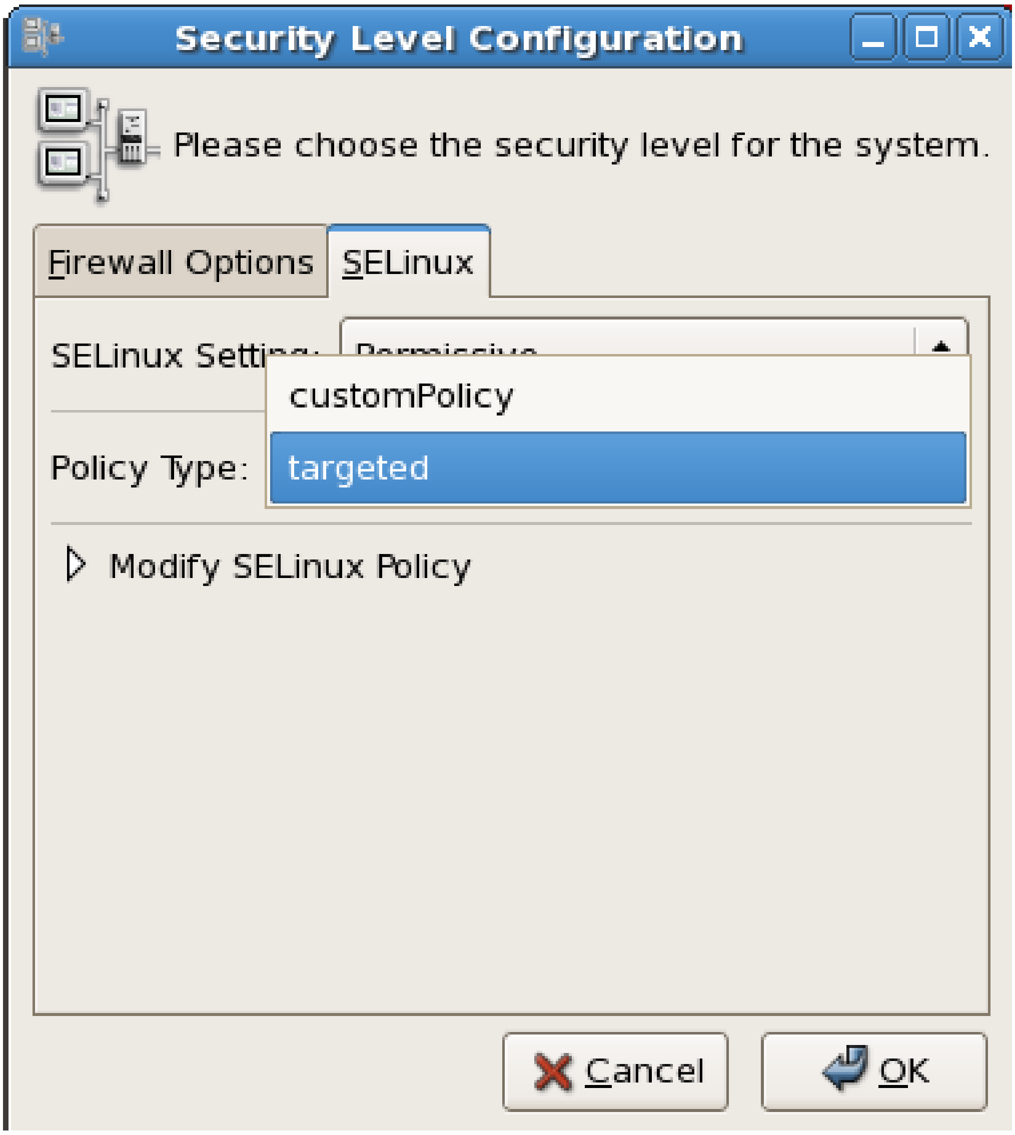

그림 48.2. Using the Security Level Configuration dialog box to load a custom policy.

48.2.9. Specifying the Security Context of Entire File Systems

You can use the mount -o context= command to set a single context for an entire file system. This might be a file system that is already mounted and that supports xattrs, or a network file system that obtains a genfs label such as cifs_t or nfs_t.

For example, if you need the Apache HTTP Server to read from a mounted directory or loopback file system, you need to set the type to httpd_sys_content_t:

```
mount -t nfs -o context=system_u:object_r:httpd_sys_content_t \
 server1.example.com:/shared/scripts /var/www/cgi
```

Tip

When troubleshooting httpd and SELinux problems, reduce the complexity of your situation. For example, if you have the file system mounted at /mnt and then symbolically linked to /var/www/html/foo, you have two security contexts to be concerned with. Because one security context is of the object class file and the other of type lnk_file, they are treated differently by the policy and unexpected behavior may occur.

48.2.10. Changing the Security Category of a File or User

Refer to 47.5.5절. "Assigning Categories to Files" and 47.5.4절. "Assigning Categories to Users" for information about adding and changing the security categories of files and users.

48.2.11. Running a Command in a Specific Security Context

You can use the runcon command to run a command in a specific context. This is useful for scripting or for testing policy, but care should be taken to ensure that it is implemented correctly.

For example, you could use the following command to run a script to test for mislabeled content. The arguments that appear after the command are considered to be part of the command. (In this example, ~/bin/contexttest is a user-defined script.)

```
runcon -t httpd_t ~/bin/contexttest -ARG1 -ARG2
```

You can also specify the entire context, as follows:

```
runcon user_u:system_r:httpd_t ~/bin/contexttest
```

48.2.12. Useful Commands for Scripts

The following is a list of useful commands introduced with SELinux, and which you may find useful when writing scripts to help administer your system:

getenforce
: This command returns the enforcing status of SELinux.

setenforce [Enforcing | Permissive | 1 | 0]

This command controls the enforcing mode of SELinux. The option 1 or Enforcing tells SELinux to enter enforcing mode. The option 0 or Permissive tells SELinux to enter passive mode. Access violations are still logged, but not prevented.

selinuxenabled

This command exits with a status of 0 if SELinux is enabled, and 1 if SELinux is disabled.

```
~]# selinuxenabled
~]# echo $?
0
```

getsebool [-a] [boolean_name]

This command shows the status of all booleans (-a) or a specific boolean (<boolean_name>).

setsebool [-P] <boolean_name> value | bool1=val1 bool2=val2 ...

This command sets one or more boolean values. The -P option makes the changes persistent across reboots.

togglesebool boolean ...

This command toggles the setting of one or more booleans. This effects boolean settings in memory only; changes are not persistent across reboots.

48.2.13. Changing to a Different Role

You use the newrole command to run a new shell with the specified type and/or role. Changing roles is typically only meaningful in the strict policy; the targeted policy is generally restricted to a single role. Changing types may be useful for testing, validation, and development purposes.

```
newrole -r <role_r> -t <type_t> [-- [ARGS]...]
```

The ARGS are passed directly to the shell specified in the user's entry in the /etc/passwd file.

Note

The newrole command is part of the policycoreutils-newrole package, which is required if you install the strict or MLS policy. It is not installed by default in Red Hat Enterprise Linux.

48.2.14. When to Reboot

The primary reason for rebooting the system from an SELinux perspective is to completely relabel the file system. On occasion you might need to reboot the system to enable or disable SELinux.

48.3. Analyst Control of SELinux

This section describes some common tasks that a security analyst might need to perform on an SELinux system.

48.3.1. Enabling Kernel Auditing

As part of an SELinux analysis or troubleshooting exercise, you might choose to enable complete kernel-level auditing. This can be quite verbose, because it generates one or more additional audit messages for each AVC audit message. To enable this level of auditing, append the audit=1 parameter to your kernel boot line, either in the /etc/grub.conf file or on the GRUB menu at boot time.

This is an example of a full audit log entry when httpd is denied access to ~/public_html because the directory is not labeled as Web content. Notice that the time and serial number stamps in the audit(...) field are identical in each case. This makes it easier to track a specific event in the audit logs:

```
Jan 15 08:03:56 hostname kernel: audit(1105805036.075:2392892): \
 avc:  denied  { getattr } for  pid=2239 exe=/usr/sbin/httpd \
 path=/home/auser/public_html dev=hdb2 ino=921135 \
 scontext=user_u:system_r:httpd_t \
 tcontext=system_u:object_r:user_home_t tclass=dir
```

The following audit message tells more about the source, including the kind of system call involved, showing that httpd tried to stat the directory:

```
Jan 15 08:03:56 hostname kernel: audit(1105805036.075:2392892): \
 syscall=195 exit=4294967283 a0=9ef88e0 a1=bfecc0d4 a2=a97ff4 \
 a3=bfecc0d4 items=1 pid=2239 loginuid=-1 uid=48 gid=48 euid=48 \
 suid=48 fsuid=48 egid=48 sgid=48 fsgid=48
```

The following message provides more information about the target:

```
Jan 15 08:03:56 hostname kernel: audit(1105805036.075:2392892): \
 item=0 name=/home/auser/public_html inode=921135 dev=00:00
```

The serial number stamp is always identical for a particular audited event. The time stamp may or may not be identical.

Note

If you are using an audit daemon for troubleshooting, the daemon may capture audit messages into a location other than /var/log/messages, such as /var/log/audit/audit.log.

48.3.2. Dumping and Viewing Logs

The Red Hat Enterprise Linux 5 implementation of SELinux routes AVC audit messages to /var/log/messages. You can use any of the standard search utilities (for example, grep), to search for lines containing avc or audit.

Customizing SELinux Policy

49.1. Introduction

In earlier releases of Red Hat Enterprise Linux it was necessary to install the selinux-policy-targeted-sources packages and then to create a local.te file in the /etc/selinux/targeted/src/policy/domains/misc directory. You could use the audit2allow utility to translate the AVC messages into allow rules, and then rebuild and reload the policy.

The problem with this was that every time a new policy package was released it would have to execute the Makefile in order to try to keep the local policy.

In Red Hat Enterprise Linux 5, this process has been completely revised. The "sources" rpm packages have been completely removed, and policy packages are treated more like the kernel. To look at the sources used to build the policy, you need to install the source rpm, selinux-policy-XYZ.src.rpm. A further package, selinux-policy-devel, has also been added, which provides further customization functionality.

49.1.1. Modular Policy

Red Hat Enterprise Linux introduces the concept of modular policy. This allows vendors to ship SELinux policy separately from the operating system policy. It also allows administrators to make local changes to policy without worrying about the next policy install. The most important command that was added was semodule.

semodule is the tool used to manage SELinux policy modules, including installing, upgrading, listing and removing modules. You can also use semodule to force a rebuild of policy from the module store and/or to force a reload of policy without performing any other transaction. semodule acts on module packages created by semodule_package. Conventionally, these files have a .pp suffix (policy package), although this is not mandated in any way.

49.1.1.1. Listing Policy Modules

To list the policy modules on a system, use the semodule -l command:

```
~]# semodule -l
amavis   1.1.0
ccs      1.0.0
clamav   1.1.0
dcc      1.1.0
evolution        1.1.0
iscsid   1.0.0
mozilla  1.1.0
mplayer  1.1.0
nagios   1.1.0
oddjob   1.0.1
pcscd    1.0.0
pyzor    1.1.0
razor    1.1.0
ricci    1.0.0
smartmon         1.1.0
```

Note

This command does not list the base policy module, which is also installed.

The /usr/share/selinux/targeted/ directory contains a number of policy package (*.pp) files. These files are included in the selinux-policy rpm and are used to build the policy file.

49.2. Building a Local Policy Module

The following section uses an actual example to demonstrate building a local policy module to address an issue with the current policy. This issue involves the ypbind init script, which executes the setsebool command, which in turn tries to use the terminal. This is generating the following denial:

```
type=AVC msg=audit(1164222416.269:22): avc:  denied  { use } for  pid=1940 comm="setsebool" name="0" dev=devpts
 ino=2 \
 scontext=system_u:system_r:semanage_t:s0 tcontext=system_u:system_r:init_t:s0 tclass=fd
```

Even though everything still works correctly (that is, it is not preventing any applications form running as intended), it does interrupt the normal work flow of the user. Creating a local policy module addresses this issue.

49.2.1. Using audit2allow to Build a Local Policy Module

The audit2allow utility now has the ability to build policy modules. Use the following command to build a policy module based on specific contents of the audit.log file:

ausearch -m AVC --comm setsebool | audit2allow -M mysemanage

The audit2allow utility has built a type enforcement file (mysemanage.te). It then executed the checkmodule command to compile a module file (mysemanage.mod). Lastly, it uses the semodule_package command to create a policy package (mysemanage.pp). The semodule_package command combines different policy files (usually just the module and potentially a file context file) into a policy package.

49.2.2. Analyzing the Type Enforcement (TE) File

Use the cat command to inspect the contents of the TE file:

```
~]# cat mysemanag.te
module mysemanage 1.0;

require {
 class fd use;
 type init_t;
 type semanage_t;
 role system_r;
};

allow semanage_t init_t:fd use;
```

The TE file is comprised of three sections. The first section is the module command, which identifies the module name and version. The module name must be unique. If you create an semanage module using the name of a pre-existing module, the system would try to replace the existing module package with the newly-created version. The last part of the module line is the version. semodule can update module packages and checks the update version against the currently installed version.

The next block of the TE file is the require block. This informs the policy loader which types, classes and roles are required in the system policy before this module can be installed. If any of these fields are undefined, the semodule command will fail.

Lastly are the allow rules. In this example, you could modify this line to dontaudit, because semodule does not need to access the file descriptor.

49.2.3. Loading the Policy Package

The last step in the process of creating a local policy module is to load the policy package into the kernel.

Use the semodule command to load the policy package:

```
~]# semodule -i mysemanage.pp
```

This command recompiles the policy file and regenerates the file context file. The changes are permanent and will survive a reboot. You can also copy the policy package file (mysemanage.pp) to other machines and install it using semodule.

The audit2allow command outputs the commands it executed to create the policy package so that you can edit the TE file. This means you can add new rules as required or change the allow rule to dontaudit. You could then recompile and repackage the policy package to be installed again.

There is no limit to the number of policy packages, so you could create one for each local modification you want to make. Alternatively, you could continue to edit a single package, but you need to ensure that the "require" statements match all of the allow rules.

References

The following references are pointers to additional information that is relevant to SELinux and Red Hat Enterprise Linux but beyond the scope of this guide. Note that due to the rapid development of SELinux, some of this material may only apply to specific releases of Red Hat Enterprise Linux.

Books

SELinux by Example

Mayer, MacMillan, and Caplan

Prentice Hall, 2007

Tutorials and Help

Understanding and Customizing the Apache HTTP SELinux Policy

http://docs.fedoraproject.org/selinux-apache-fc3/

Tutorials and talks from Russell Coker

http://www.coker.com.au/selinux/talks/ibmtu-2004/

Generic Writing SELinux policy HOWTO

https://sourceforge.net/docman/display_doc.php?docid=21959[amp]group_id=21266[1]

Red Hat Knowledgebase

http://kbase.redhat.com/

General Information

NSA SELinux main website

http://www.nsa.gov/research/selinux/index.shtml

NSA SELinux FAQ

http://www.nsa.gov/research/selinux/faqs.shtml

Fedora SELinux FAQ

http://docs.fedoraproject.org/selinux-faq/

SELinux NSA's Open Source Security Enhanced Linux

http://www.oreilly.com/catalog/selinux/

Technology

An Overview of Object Classes and Permissions

http://www.tresys.com/selinux/obj_perms_help.html

Integrating Flexible Support for Security Policies into the Linux Operating System (a history of Flask implementation in Linux)

http://www.nsa.gov/research/_files/selinux/papers/freenix01/freenix01.shtml

Implementing SELinux as a Linux Security Module

http://www.nsa.gov/research/selinux/index.shtmlpapers/module-abs.cfm

A Security Policy Configuration for the Security-Enhanced Linux

http://www.nsa.gov/research/_files/selinux/papers/policy/policy.shtml

[1] https://sourceforge.net/docman/display_doc.php?docid=21959[amp]group_id=21266

Community

SELinux community page

http://selinux.sourceforge.net

IRC

irc.freenode.net, #rhel-selinux

History

Quick history of Flask

http://www.cs.utah.edu/flux/fluke/html/flask.html

Full background on Fluke

http://www.cs.utah.edu/flux/fluke/html/index.html

부 VIII. Red Hat 교육 및 자격증

Red Hat 교육 과정 및 자격증은 Linux 및 모든 IT 분야에서 최고로 간주되어 지고 있습니다. 경험이 풍부한 Red Hat 전문가에 의해 교육이 제공되며, 실전 환경에서의 역량을 측정하는 자격증 프로그램은 고용주와 IT 전문가에 의해 각광을 받고 있습니다.

필요와 목적에 따라 적절한 자격증을 선택하시기 바랍니다. 실력이 우수하든지, 우수하지 않든지, UNIX 사용 경험이 있든지 또는 사용 경험이 없든지 간에 이에 상관 없이, Red Hat 교육 및 자격증 프로그램은 여러분에게 적합한 교육 내용을 제공해 드립니다.

Red Hat 교육 및 자격증

51.1. 세 가지 교육 방법

정규 과정

정규 과정은 세계적으로 125곳이 넘는 교육 센터에서 제공됩니다. Red Hat 교육 과정은 실습 위주 교육이기 때문에 최소 한 대의 전용 시스템 또는 일부 교육 과정에서 다섯 대 이상의 시스템을 사용할 수 있습니다. 모든 강사는 숙련된 Red Hat Certified Engineer(RHCE)로서 모든 교육 과정을 다룰 수 있습니다.

교육 과정 스케쥴은 http://www.redhat.com/explore/training을 참조하시기 바랍니다.

출장방문 교육

Onsite training is delivered by Red Hat at your facility for teams of 12 to 16 people per class. Red Hat's technical staff will assist your technical staff prior to arrival to ensure the training venue is prepared to run Red Hat Enterprise Linux, Red Hat or JBoss courses, and/or Red Hat certification exams. Onsites are a great way to train large groups at once. Open enrollment can be leveraged later for incremental training.

더 자세한 정보는 http://www.redhat.com/explore/onsite에서 참조하시기 바랍니다.

이러닝(eLearning)

Fully updated for Red Hat Enterprise Linux 4! No time for class? Red Hat's e—Learning titles are delivered online and cover RHCT and RHCE track skills. Our growing catalog also includes courses on the latest programming languages, scripting and ecommerce.

더 자세한 교육 과정 목록을 보시려면 http://www.redhat.com/explore/elearning을 참조하시기 바랍니다.

51.2. Microsoft Certified Professional 자원 센터

개인 포트폴리오에 Red Hat 자격을 추가하려는 Microsoft® Certified Professional을 위한 전용 정보와 기회가 마련되어 있습니다.

오늘 바로 확인하시기 바랍니다: http://www.redhat.com/explore/manager

자격증 과정

Red Hat Certified Technician® (RHCT®)

RHCT가 소개된 후 이제 3년째에 접어들고 있습니다. Red Hat Certified Technician은 리눅스 분야에서 가장 빠르게 발전하는 자격증이며 현재 15,000명이 넘는 RHCT가 있습니다. RHCT는 리눅스 자격증 과정의 첫 단계이며 UNIX®/ Linux 환경으로 전환하려는 사용자에게 적합한 기초 자격증 과정입니다.

Red Hat은 리눅스 및 모든 IT 분야에서 최고의 교육 및 자격증 과정을 제공합니다. 모든 과정은 숙련된 Red Hat 전문가에 의해 진행되며 자격증 프로그램은 실제 작동 중인 시스템에서의 활용 능력을 검증하기 때문에 고용업체 및 IT 전문가에 의한 수요가 매우 높습니다.

리눅스 지식 및 목표에 맞는 적당한 자격증 과정을 선정할 수 있습니다. Red Hat은 고급, 초급 또는 유닉스/리눅스에 대한 무경험자를 비롯한 모든 대상에 맞는 교육 및 자격증 과정을 제공합니다.

Red Hat Certified Engineer® (RHCE®)

Red Hat Certified Engineer began in 1999 and has been earned by more than 20,000 Linux experts. Called the "crown jewel of Linux certifications," independent surveys have ranked the RHCE program #1 in all of IT.

Red Hat Certified Security Specialist (RHCSS)

An RHCSS has RHCE security knowledge plus specialized skills in Red Hat Enterprise Linux, Red Hat Directory Server and SELinux to meet the security requirements of today's enterprise environments. RHCSS is Red Hat's newest certification, and the only one of its kind in Linux.

Red Hat Certified Architect (RHCA)

고급 교육을 원하는 RHCE는 Enterprise Architect 과정에 등록하여 새롭게 소개된 Red Hat Certified Architect(RHCA) 자격을 갖고 능력을 행사할 수 있습니다. RHCA는 Red Hat Certified Technician(RHCT) 및 Red Hat Certified Engineer(RHCE)의 마지막 단계이며 리눅스 분야에서 최고로 인정되는 자격증입니다.

52.1. 무료 모의 테스트

자동화된 무료 모의 테스트에서 리눅스 지식을 검사하여 Red Hat 교육 과정 수준을 선정할 수 있습니다.

이는 무료로 제공되며 어떠한 의무 조건도 없습니다. 10분 정도가 소요됩니다. http://www.redhat.com/explore/assess

RH033: Red Hat Linux Essentials

http://www.redhat.com/training/rhce/courses/rh033.html

53.1. 교육 과정 소개

RHCT 및 RHCE 자격증을 위한 교육 과정의 첫 단계이며 리눅스 또는 유닉스를 처음 사용하거나 다른 운영 체제에서 명령 행을 다루어 보지 않은 사용자에게 필요합니다. RH033에서는 Red Hat Enterprise Linux 환경에서 필요한 기본 요소에 대해 배우며 시스템 관리자의 역할을 수행할 수 있는 기본 능력을 갖추게 합니다.

53.1.1. 선수 조건

다른 컴퓨터 시스템, 메뉴, 그래픽형식 사용자 인터페이스를 사용할 수 있는 정도의 경험이 필요합니다.

53.1.2. 목표

Red Hat 시스템에 일반 명령 행 프로세스를 사용하고 데스크탑 생산 역할을 수행할 수 있는 활용 능력을 갖추며 시스템 관리(RH133) 과정을 배울 준비가 되어있는 Red Hat Enterprise Linux의 고급 사용자로 양성합니다.

53.1.3. 교육 대상

리눅스를 처음 사용하고 유닉스 또는 명령행에 대한 경험이 없으며 Red Hat Linux 시스템을 사용하고 제어하는데 필요한 기본 능력을 개발하고 발전시키려는 사용자를 대상으로 합니다.

53.1.4. 교육 목적

1. 리눅스 파일 시스템의 이해
2. 일반 파일 관리 작업 수행
3. GNOME 인터페이스 사용 및 구성
4. 명령 행에서 주요 리눅스 명령어 사용
5. GNOME GUI를 사용하여 일반 작업 수행
6. vi 편집기를 사용하여 텍스트 문서 열기, 편집하기, 저장하기
7. 파일 사용 권한
8. X Windows 시스템 구성
9. 정규 표현식 패턴 일치 및 I/O 방향전환
10. 시스템에 패키지 설치, 업그레이드, 삭제, 쿼리
11. 사용자에게 필요한 네트워크 유틸리티
12. 고급 사용자 유틸리티

53.1.5. 다음 교육 과정

RH133 Red Hat Linux Sys. Admin.

RH253 Red Hat Linux Net. and Sec. Admin

RH300 Red Hat Linux RHCE Rapid Track

"I would enthusiastically recommend this course to anyone interested in Linux."——Mike Kimmel, ITT Systems Division

RH035: Red Hat Linux Essentials for Windows Professionals

http://www.redhat.com/training/rhce/courses/rh035.html

54.1. 교육 과정 소개

RH035는 유닉스 또는 리눅스 경험이 없는 Windows® 전문가를 대상으로 하며 기본적인 Red Hat Enterprise Linux 시스템 관리 기술을 다룹니다. 첫 날에는 포트폴리오에 리눅스 관리 능력을 추가할 수 있는 Windows에서 리눅스로의 개념적이고 실용적인 전환 과정에 대해 소개합니다. 나머지 4일은 높은 인기의 RH033 과정을 도입하여 Red Hat Enterprise Linux 환경에 필요한 기본 기술을 소개하고 크로스 플랫폼 시스템 관리자로서의 차후 역할을 수행할 수 있는 능력을 발전시킵니다. 이 교육 과정에서는 RHCT 및 RHCE를 위한 기본적인 내용을 다룹니다.

54.1.1. 선수 조건

유닉스 또는 리눅스 경험이 없지만 기술자 또는 시스템 관리자 수준에서 Windows OS 제품을 사용한 작업 수행 경험을 갖추고 있는 IT 전문가이어야 합니다.

54.1.2. 목표

그래픽형식 도구를 사용한 일부 시스템 관리 작업 및 일반 명령 행 프로세스를 수행할 수 있는 능력을 갖추며 더 난해한 Red Hat Enterprise Linux System Administration(RH133) 과정을 배울 준비가 되어있는 Red Hat Enterprise Linux의 고급 사용자로 양성합니다.

54.1.3. 교육 대상

그래픽형식 사용자 인터페이스에서 서버를 관리하는 Windows 기술자이지만 Red Hat Enterprise Linux 시스템도 효과적으로 관리하고 개인 능력을 발전시키려는 사용자를 대상으로 합니다.

54.1.4. 교육 목적

1. 그래픽형식 도구를 사용한 소프트웨어 설치, 네트워크 구성, 사용자 인증 구성, 다양한 서비스 설치 및 구성

2. 리눅스 파일 시스템의 이해

3. 명령 행에서 주요 리눅스 명령어 사용

4. 파일 사용 권한 이해

5. X Windows 시스템 구성

6. 정규 표현식 패턴 및 I/O 방향전환

54.1.5. 다음 교육 과정

RH133 Red Hat Linux Sys. Admin. (p. 8)

RH253 Red Hat Linux Net. and Sec. Admin. (p. 9)

RH300 Red Hat Linux RHCE Rapid Track (p. 10)

"All in all I would rate this training experience as one of the best I have ever attended, and I've been in this industry for over 15 years." — Bill Legge, IT Consultant

RH133: Red Hat Linux System Administration and Red Hat Certified Technician (RHCT) Certification

http://www.redhat.com/training/rhce/courses/rh133.html

55.1. 교육 과정 소개

RH133은 Red Hat 리눅스에서 현재 네트워크에 워크스테이션을 추가하고 구성할 수 있는 수준의 시스템 관리 능력을 집중적으로 다룹니다. 4.5일의 교육 과정에서는 Red Hat Enterprise Linux 시스템에서 집중적인 실습 위주의 교육을 진행하며 마지막 날에는 RH202 RHCT Certification Lab Exam 과정을 진행합니다.

55.1.1. 선수 조건

RH033 Red Hat Linux Essentials 또는 이에 해당하는 Red Hat 리눅스 경험이 필요합니다.

55.1.2. 목표

이 교육 과정을 완전히 이수하면, RHCT 시험에 합격하여 인증된 리눅스 시스템 기본 관리 지식을 갖출 수 있습니다. 시험은 수행 능력 기반의 실습 시험으로서 현재 사용되는 네트워크에 새로운 Red Hat 리눅스 시스템을 설치, 구성, 추가할 수 있는 실제 능력을 평가합니다.

55.1.3. 교육 대상

Red Hat 리눅스의 기본 요소를 이해하고 시스템 관리자가 되고자 다음 단계의 기술 교육이 필요한 리눅스 또는 유닉스 사용자를 대상으로 합니다.

55.1.4. 교육 목적

1. 대화식 또는 킥스타트를 사용한 Red Hat 리눅스 설치

2. 일반 시스템 하드웨어 제어 및 리눅스 프린트 하위시스템 관리

3. 리눅스 파일 시스템 생성 및 관리

4. 사용자 및 그룹 관리 수행

5. 현재 사용되는 네트워크에 워크스테이션 구성

6. 워크스테이션을 NIS, DNS, DHCP 서비스 클라이언트로 구성

7. at, cron, anacron을 사용한 자동화 작업

8. 테입 및 tar 아카이브로 파일 시스템 백업

9. RPM을 사용한 소프트웨어 패키지 관리

10. X Windows 시스템 및 GNOME d.e. 구성

11. 수행 능력, 메모리, 프로세스 관리

12. 기본적인 호스트 보안 구성

55.1.5. 다음 교육 과정

RH253 Red Hat Linux Net. and Sec. Admin. (p. 9)

RH202 RHCT EXAM - 리눅스 분야에서 가장 빠르게 발전하는 자격증

http://www.redhat.com/training/rhce/courses/rh202.html

1. RHCT 시험은 RH133 과정에 포함되어 있지만 별도로 취득할 수 도 있습니다 (미화 $349).

2. RHCT 시험은 RH133 과정의 5번째 날에 진행됩니다.

56.1. 교육 과정 소개

RHCT(Red Hat Certified Technician)는 수행 능력 기반의 실습 시험으로서 Red Hat Enterprise Linux를 설치, 구성, 문제 해결할 수 있는 실제 능력을 평가합니다. Certification Lab Exam은 RH133 과정에 포함되어 있지만 RH033 및 RH133 내용을 완전히 이수하신 분들은 바로 시험을 보실 수 있습니다.

56.1.1. 선수 조건

시험을 준비 중인 수험자는 RH033 및 RH133 과정을 이수하여야 하지만 선수 조건은 아닙니다.

RH253 Red Hat Linux Networking and Security Administration

57.1. 교육 과정 소개

RH253 과정에서는 Red Hat Enterprise Linux에서 일반 네트워크 서비스를 구성하는 데 필요한 상세한 지식을 갖출 수 있으며 네트워크 및 내부 보안 작업에 관한 주제도 다루고 있습니다.

57.1.1. 선수 조건

RH133 Red Hat Linux System Administration 또는 이에 해당하는 Red Hat Enterprise Linux 경험과 LAN/WAN 구조 및 TCP/IP 인터네트워킹 경험이 필요합니다.

57.1.2. 목표

이 과정을 완전히 이수하면 Red Hat Enterprise Linux 서버를 설정하고 일반 네트워크 서비스 및 보안을 기본 수준으로 구성할 수 있습니다.

57.1.3. 교육 대상

Red Hat Enterprise Linux 시스템 관리에 대한 실제 경험이 있으며 네트워킹 서비스 및 보안에 대한 첫 교육 과정에 참여하여 Red Hat Enterprise Linux를 사용한 일반적인 네트워크 서비스 및 보안 관리 기술을 발전시키려는 리눅스 또는 유닉스 시스템 관리자를 대상으로 합니다.

57.1.4. 교육 목적

1. Red Hat 리눅스 서버 측면에서 네트워킹 서비스 설정 및 구성과 일반 네트워킹 서비스 기본 관리 능력 :DNS, NIS, Apache, SMB, DHCP, Sendmail, FTP를 비롯한 기타 서비스 (tftp, pppd, proxy)

2. 보안 소개

3. 보안 정책 개발

4. 내부 보안

5. 파일 및 파일 시스템 보안

6. 패스워드 보안

7. 커널 보안

8. 방화벽 기본 구성 요소

9. Red Hat 리눅스 기반 보안 도구

10. 공격 시도에 대한 대처

11. 보안 소스 및 방법

12. OSS 보안 도구 개요

57.1.5. 다음 교육 과정

RH302 RHCE Certification Exam

"This course was excellent. The teacher was fantastic—his depth of knowledge is amazing."——Greg Peters, Future Networks USA

RH300: RHCE Rapid Track Course (RHCE 시험 포함)

숙련된 유닉스/리눅스 사용자를 대상으로 하여 단기간에 RHCE 자격증을 취득하는 방법.

http://www.redhat.com/training/rhce/courses/rh300.html

58.1. 교육 과정 소개

5일 간의 교육 과정은 Red Hat 리눅스 시스템에서 실습 위주의 집중 교육을 제공하며 마지막 날에는 RHCE 자격증 시험을 진행합니다.

58.1.1. 선수 조건

RH033, RH133, RH253 또는 이에 해당하는 유닉스 경험이 필요합니다. 시스템 관리자 경험이 없거나 유닉스 또는 리눅스 환경의 고급 사용자가 아니면 RH300 과정에 등록할 수 없습니다.

58.1.2. 목표

이 과정을 완전히 이수하면 RHCE 시험에서 검증된 고급 Red Hat Linux 시스템 관리자가 될 수 있습니다.

58.1.3. 교육 대상

시스템 관리에 대한 실제 우수한 경험을 갖추고 있으며 단기 과정으로 RHCE 시험을 준비하려는 유닉스 또는 리눅스 시스템 관리자를 대상으로 합니다.

58.1.4. 교육 목적

1. 하드웨어 및 설치 (x86 구조)
2. 구성 및 관리
3. 대안적 설치 방법
4. 커널 서비스 및 구성
5. 표준 네트워킹 서비스
6. X Windows 시스템
7. 사용자 및 호스트 보안
8. 라우터, 방화벽, 클러스터, 문제 해결

58.1.5. 다음 교육 과정

Enterprise Architect 교육 과정 및 RHCA 자격증

RH302 RHCE EXAM

1. RHCE 시험은 RH300 과정에 포함되어 있지만, 별도로 취득할 수 도 있습니다.

2. RHCE 시험은 RH300 과정의 5번째 날에 진행됩니다.

http://www.redhat.com/training/rhce/courses/rhexam.html

59.1. 교육 과정 소개

RHCE는 IT 분야의 다른 자격증 프로그램과 다릅니다. RHCE 프로그램은 수행 능력 기반의 실습 시험으로서 Red Hat 리눅스의 설치, 구성, 디버깅, 주요 네트워킹 서비스 설정에서의 실제 능력을 평가합니다.

59.1.1. 선수 조건

RH300 교육 과정의 선수 조건과 같습니다. 더 자세한 정보는 RHCE 시험 준비 가이드에서 참조하시기 바랍니다: www.redhat.com/training/rhce/examprep.html

59.1.2. 내용

1. 섹션 I: 문제 해결 및 시스템 관리 (2.5시간)

2. 섹션 II: 설치 및 구성 (3시간)

"Seriously, this was an outstanding class. I feel very well prepared for the test tomorrow." — Logan Ingalls, Web developer, Texterity Inc., USA

RHS333: Red Hat Enterprise Security Network Services

가장 널리 사용되는 서비스에 대한 보안.

http://www.redhat.com/training/architect/courses/rhs333.html

60.1. 교육 과정 소개

Red Hat Enterprise Linux has gained considerable momentum as the operating system of choice for deploying network services such as web, ftp, email, and file sharing. Red Hat's RHCE curriculum provides training in deploying these services and on the essential elements of securing them.

60.1.1. 선수 조건

RH253, RH300을 이수하거나 RHCE 자격증 또는 이에 해당하는 실제 경험이 필요합니다. 본 과정은 보다 심화된 주제를 다루고 있으므로 서비스 설정 방법에 대한 핵심 사항들을 모두 숙지하고 계셔야 합니다.

60.1.2. 목표

RHCE 교과 과정에서 제공된 필수 보안 내용은 물론 가장 널리 활용되는 서비스에 관련된 보안 기능, 수용 능력, 위험성에 대해 심화 학습합니다.

60.1.3. 교육 대상

이 과정은 시스템 관리자, 컨설턴트 또는 네트워크 서버 계획, 구현, 관리에 관계된 다른 IT 전문가를 대상으로 합니다. 이는 Red Hat Enterprise Linux에서 작동하는 서비스를 중점으로 한 내용 및 실습으로 이루어져 있으며 기존 유닉스를 사용하는 시스템 관리자 및 다른 IT 전문가에게도 매우 유용할 것입니다.

60.1.4. 교육 목적

1. 기본적인 서비스 보안 이해

2. 암호화 이해

3. 시스템 동작 로깅

4. BIND 및 DNS 보안

5. 네트워크 사용자 인증 보안

6. NFS 보안 개선

7. 보안 쉘: OpenSSH

8. Sendmail 및 Postfix에서 이메일 보안

9. FTP 사용 관리

10. Apache 보안

11. 침입 대응의 기본

60.1.5. 다음 교육 과정

RH401 Red Hat Enterprise Deployment and System Mgmt. RH423 Red Hat Enterprise Directory Services and Authentication RH436 Red Hat Enterprise Storage Mgmt. RH442 Red Hat Enterprise System Monitoring and Performance Tuning

RH401: Red Hat Enterprise Deployment and Systems Management

Red Hat Enterprise Linux 활용 관리.

http://www.redhat.com/training/architect/courses/rh401.html

61.1. 교육 과정 소개

RH401은 4일 간의 실습 위주의 집중 교육 과정으로서 문제 해결 및 로드 밸런싱, 시스템 관리자를 위한 CVS, RPM 리빌드, 전용 프로그램 성능 개선 등 임무 중심적인 대규모 Red Hat Enterprise Linux 시스템을 활용 및 관리하는 데 필요한 핵심 기술 및 방법을 교육합니다.

61.1.1. 선수 조건

RH253 at a minimum, RHCE certification preferred, or comparable skills and knowledge. All prospective course participants without RHCE certification are encouraged to verify skills with Red Hat's free online pre—assessment tests. Note: Persons should not enroll in RH401 without meeting the above prerequisites.

RHCE 자격증이 없는 모든 수강자는 등록 시 반드시 Red Hat Global Learning Services에 연락하여 능력 평가를 받으셔야 합니다.

61.1.2. 목표

RH401 과정에서는 고급 시스템 관리자를 대상으로 다양한 용도의 대규모 Enterprise Linux 서버를 관리하며 문제 해결 및 로드 밸런싱을 필요로 하는 임무 중심적 프로그램을 관리할 수 있는 능력을 가르칩니다. 또한, RH401은 운영 체제를 기업체 업무용으로 관리하는 데 필요한 전문가 수준의 핵심 기술을 교육하는 것으로 정평이 나 있습니다. 따라서, Red Hat Enterprise Linux를 사용한 전체 기업체 시스템을 관리 팀에 의해 쉽게 유지될 수 있도록 효과적이고 효율적인 구현 및 관리 기술을 배울 수 있습니다.

61.1.3. 교육 대상

고급 Red Hat Enterprise Linux 시스템 관리자 및 기업체 환경과 임무 중심적 시스템에서 종사하는 IT 전문가를 대상으로 합니다.

61.1.4. 교육 목적

1. CVS를 사용한 구성 관리

2. 전용 RPM 패키지 개발

3. Red Hat Network 프록시 서버를 사용한 소프트웨어 관리

4. 호스트 프로비저닝 및 관리 시스템 조합

5. 성능 개선 및 분석

6. 고가용성 네트워크 로드 밸런싱 클러스터

7. 고가용성 프로그램 문제 해결 클러스터

61.1.5. 다음 교육 과정

RHS333 Enterprise Security: Securing Network Services

RH423 Red Hat Enterprise Directory Services and Authentication

RH436 Red Hat Enterprise Storage Mgmt.

RH442 Red Hat Enterprise System Monitoring and Performance Tuning

"After taking RH401 I am completely confident that I can implement enterprise—scale high—availability solutions end-to-end."——Barry Brimer, Bunge North America

RH423: Red Hat Enterprise Directory Services and Authentication

Red Hat Enterprise Linux 시스템에 필요한 디렉토리 서비스 관리 및 활용.

http://www.redhat.com/training/architect/courses/rh423.html

62.1. 교육 과정 소개

RH423은 기업 전반에 인증 및 정보 서비스를 제공하는 디렉토리 서비스의 이기종 통합에 대한 4일간의 집중 교육 및 실습을 제공합니다.

62.1.1. 선수 조건

RH253 at a minimum, RHCE certification preferred, or comparable skills and knowledge. All prospective course participants without RHCE certification are encouraged to verify skills with Red Hat's free online pre—assessment tests. Note: Persons should not enroll in RH423 without meeting the above prerequisites. All prospective course participants who do not possess RHCE certification are strongly advised to contact Red Hat Global Learning Services for a skills assessment when they enroll.

62.1.2. 목표

RH423에서는 고급 시스템 관리자를 대상으로 Red Hat Enterprise Linux 시스템에서 디렉토리 서비스를 관리하고 활용하는 방법을 배울 수 있으며 LDAP 기반 서비스의 기본 개념 이해, 구성, 관리에 대한 내용을 주로 다룹니다. 표준 네트워크 클라이언트 및 서비스를 디렉토리 서비스에 통합하여 디렉토리 서비스의 성능을 이용하도록 실습합니다. 또한, PAM(Pluggable Authentication Modules) 시스템이 인증 및 권한이 필요한 서비스와 어떻게 통합되는지를 교육합니다.

62.1.3. 교육 대상

고급 Red Hat Enterprise Linux 시스템 관리자 및 기업체 환경과 임무 중심적 시스템에서 종사하는 IT 전문가를 대상으로 합니다.

62.1.4. 교육 목적

1. LDAP의 기본 개념

2. OpenLDAP 서버 구성 및 관리

3. Using LDAP as a "white pages" directory service

4. LDAP를 사용한 사용자 인증 및 관리

5. 다중 LDAP 서버 통합

62.1.5. 다음 교육 과정

RHS333 Enterprise Security: Securing Network Services

RH401 Red Hat Enterprise Deployment and Systems Management

RH436 Red Hat Enterprise Storage Mgmt. (p. 16)

RH442 Red Hat Enterprise System Monitoring and Performance Tuning

SELinux Courses

63.1. RHS427: Introduction to SELinux and Red Hat Targeted Policy

http://www.redhat.com/training/security/courses/rhs427.html

SELinux을 소개하는 1일 단기 과정으로서 SELinux가 Red Hat 대상 정책과 어떻게 운영되는 지를 교육하며 SELinux의 고급 성능을 이용하는 프로그램 사용법에 대해 다룹니다. RH429의 첫날 과정이 RHS427로 구성되어 있습니다.

63.1.1. 교육 대상

컴퓨터 보안 전문가 및 리눅스 컴퓨터에서 보안 정책을 구현하는 IT 전문가를 대상으로 합니다. RHS429 과정은 RHCE 자격증 또는 이에 해당하는 지식이 필요합니다.

63.1.2. 교육 과정 요약

Red Hat Enterprise Linux의 주요 기능 중 하나는 바로 SELinux(Security Enhanced Linux)이며 이것은 사용자 및 프로세스가 시스템의 어느 부분을 사용할지에 대한 효율적인 제어 능력을 갖춘 강력한 커널 수준 보안 계층을 제공합니다. SELinux는 디폴트로 Red Hat Enterprise Linux에서 활성화되어 있으며 Red Hat에 대상 정책을 불러들이는 필수 엑세스 컨트롤 조합을 실행합니다. 이러한 엑세스 컨트롤은 대상 네트워크 서비스 보안을 강화하지만, 이전 버전의 Red Hat Enterprise Linux에서 작동하는 제3자 프로그램 및 스크립트 행동에 영향을 줄 수도 있습니다.

63.2. RHS429: Red Hat Enterprise SELinux Policy Administration

http://www.redhat.com/training/security/courses/rhs429.html

Red Hat Enterprise Linux의 주요 기능 중 하나는 바로 SELinux(Security Enhanced Linux)이며 이것은 사용자 및 프로세스가 시스템의 어느 부분을 사용할지에 대한 효율적인 제어 능력을 갖춘 강력한 커널 수준 보안 계층을 제공합니다. RHS429는 고급 시스템 관리자, 보안 관리자, 어플리케이션 프로그래머를 대상으로 SELinux 정책 작성에 대해 교육합니다. 이 과정에서 수강자는 SELinux가 어떻게 작동하는지 이해하며 SELinux 관리 및 SELinux 정책 작성 방법에 대해 실습할 수 있습니다.

RH436: Red Hat Enterprise Storage Management

Deploy and manage Red Hat's cluster file system technology.

사용 기술:

1. 다섯개의 서버

2. 스토리지 배열

http://www.redhat.com/training/architect/courses/rh436.html

64.1. 교육 과정 소개

RH436은 Red Hat 글로벌 파일 시스템(GFS)을 통해 최근 각광받는 기술인 공유 스토리지에 대한 4일간의 집중 실무 교육을 제공하며 Red Hat 클러스터 수트 및 GFS에서 실제 Red Hat Enterprise Linux 기술을 구현하는 방법을 집중적으로 다룹니다.

64.1.1. 선수 조건

RH253 at a minimum, RHCE certification preferred, or comparable skills and knowledge. All prospective course participants without RHCE certification are encouraged to verify skills with Red Hat's free online pre—assessment tests.

64.1.2. 목표

이 과정은 임무 중심적 컴퓨터 환경에 대한 고가용성 스토리지 데이터를 관리하고 활용할 수 있는 RHCE 수준의 능력을 갖춘 전문가를 대상으로 고안되었습니다. RH401에서 습득한 기술은 물론 클러스터 파일 시스템 및 GFS에 대한 집중적인 실무 교육이 진행됩니다.

64.1.3. 교육 대상

고급 Red Hat Enterprise Linux 시스템 관리자 및 기업체 환경과 임무 중심적 시스템에서 종사하는 IT 전문가를 대상으로 합니다.

64.1.4. 교육 목적

1. Red Hat Enterprise Linux 스토리지 관리 기술 검토

2. 데이터 스토리지 설계: 데이터 공유

3. 클러스터 수트 개요

4. 글로벌 파일 시스템(GFS) 개요

5. GFS 관리

6. 온라인 GFS 환경 수정: 데이터 용량 관리

7. GFS 모니터

8. GFS 수정 구현

9. 클러스터 수트 NFS를 DAS에서 GFS로 이전

10. GFS를 사용한 클러스터 수트

64.1.5. 다음 교육 과정

RHS333 Enterprise Security: Securing Network Services

RH401 Red Hat Enterprise Deployment and Systems Management

RH423 Red Hat Enterprise Directory Services and Authentication

RH442 Red Hat Enterprise System Monitoring and Performance Tuning

"The class gave me a chance to use some of the latest Linux tools, and was a reminder of the benefits of using Linux for high-availability systems."——Paul W. Frields, FBI — Operational Technology Division Quantico, VA, USA

RH442: Red Hat Enterprise System Monitoring and Performance Tuning

Red Hat Enterprise Linux 성능 개선 및 용량 관리 계획

http://www.redhat.com/training/architect/courses/rh442.html

65.1. 교육 과정 소개

RH442는 4일간의 고급 실무 교육으로서 시스템 구조, 성능 특성, 모니터링, 벤치마킹, 네트워크 성능 개선 등을 중점으로 다룹니다.

65.1.1. 선수 조건

RHCT at a minimum, RHCE certification recommended, or comparable skills and knowledge. All prospective course participants without RHCE certification are encouraged to verify skills with Red Hat's free online pre—assessment tests.

65.1.2. 목표

RH442는 Red Hat Enterprise Linux 성능 개선 및 용량 관리 계획에 대한 방법론을 제시하며 다음 사항을 다룹니다:

1. 시스템 성능에 따른 시스템 구조의 이해에 바탕을 둔 개념 소개

2. 성능 조작에 따른 결과 측정 방법 (벤치마킹)

3. 오픈 소스 벤치마킹 유틸리티

4. 시스템 성능 및 네트워크 성능 분석 방법

5. 특정 어플리케이션 로드 성능 개선

65.1.3. 교육 대상

RH442는 고급 Red Hat Enterprise Linux 시스템 관리자 및 기업체 환경과 임무 중심적 시스템에서 종사하는 IT 전문가를 대상으로 합니다.

65.1.4. 교육 목적

1. 시스템 성능과 관련된 시스템 구성요소 및 구조 개요

2. Translating manufacturers' hardware specifications into useful information

3. 표준 모니터링 도구를 사용하여 효과적인 상태 정보 수집 및 분석

4. SNMP를 사용한 성능 관련 데이터 수집

5. 오픈 소스 벤치마킹 유틸리티 사용

6. 네트워크 성능 개선

7. 어플리케이션 성능 개선

8. 특정 구성 개선

65.1.5. 다음 교육 과정

RHS333 Enterprise Security: Securing Network Services

RH401 Red Hat Enterprise Deployment and Systems Management

RH423 Red Hat Enterprise Directory Services and Authentication

RH436 Red Hat Enterprise Storage Mgmt.

Red Hat Enterprise Linux 개발자 교육 과정

66.1. RHD143: Red Hat Linux Programming Essentials

http://www.redhat.com/training/developer/courses/rhd143.html

RHD143은 Red Hat Enterprise Linux에서 어플리케이션 및 프로그램을 개발하는 데 필요한 주요 기술을 습득하는 실습 위주의 단기 과정입니다. 5일 동안 진행되는 과정은 실무 교육, 개념 설명, 예시, 실제 실습 및 프로그램 연습 등으로 구성됩니다. 이 과정을 이수하면 리눅스 시스템 기반 프로그램을 개발하는 데 필요한 주요 기술을 습득할 수 있습니다.

66.2. RHD221 Red Hat Linux Device Drivers

http://www.redhat.com/training/developer/courses/rhd221.html

이 과정은 숙련된 프로그래머를 대상으로 리눅스 시스템에 쓰이는 장치 드라이버를 어떻게 개발하는지 교육합니다. 이 과정을 이수하면 리눅스 구조, 하드웨어, 메모리 관리, 모듈 관리, 커널 소스 레이아웃을 이해할 수 있으며 문자, 블록, 네트워크 드라이버 개발에 필요한 주요 개념 및 기술을 습득할 수 있습니다.

66.3. RHD236 Red Hat Linux Kernel Internals

http://www.redhat.com/training/developer/courses/rhd236.html

이 과정은 프로세스 스케줄링, 메모리 관리, 파일 시스템, 주변 장치 관리 등을 비롯한 리눅스 커널 구조에 대한 자세한 교육을 제공합니다. 5일 동안 진행되는 이 교육은 실무 교육, 개념 설명, 예시, 실제 실습 및 프로그램 연습 등으로 구성됩니다.

66.4. RHD256 Red Hat Linux Application Development and Porting

http://www.redhat.com/training/developer/courses/rhd256.html

4일 간의 개발자 교육 과정으로서 이미 유닉스 시스템에서 개발 경험이 있는 숙련된 프로그래머를 대상으로 Red Hat Enterprise Linux에 새로운 어플리케이션 개발 및 기존 어플리케이션 포팅 기술을 교육합니다.

JBoss 교육 과정

67.1. RHD161 JBoss and EJB3 for Java

http://www.redhat.com/training/jboss/courses/rhd161.html

JBoss and EJB3 for Java Developers 교육 과정은 숙련된 Java 개발자를 대상으로 JBoss 어플리케이션 서버를 사용한 EJB3 및 J2EE 미들웨어 프로그래밍에 대한 지식을 교육합니다. 이 과정에서는 JBoss 어플리케이션 서버를 사용한 EJB3 및 J2EE에 대한 자세한 개념을 설명하며 EJB3 및 J2EE 어플리케이션 개발 및 활용 그리고 이 과정에 사용되는 도구들에 대한 실습을 제공합니다.

67.1.1. 선수 조건

기본 Java 프로그래밍 기술 및 OOAD 개념에 대한 지식이 필요합니다. 수강자는 반드시 다음에 관한 실무 지식 또는 경험을 갖추고 있어야 합니다:

1. 상속성, 폴리모피즘, 캡슐화에 대한 객체 지향 개념

2. 데이터 유형, 변수, 운영자, 문장, 흐름 제어에 관한 Java 문법

3. Java 클래스 작성, Java 인터페이스 및 추상 클래스 사용

67.2. RHD163 JBoss for Web Developers

http://www.redhat.com/training/jboss/courses/rhd163.html

JBoss for Web Developers 교육 과정은 JBoss Enterprise Middleware System(JEMS) 제품 스택에서 웹 계층화 기술을 중점으로 다룹니다. JBoss Portal에 대한 자세한 설명, 포틀릿 생성 및 활용 방법, 포틀릿과 JavaServer Faces(JSF)와 같은 다른 웹 계층화 프레임워크 통합, JBoss 어플리케이션 서버에 내장된 Tomcat 웹 컨테이너 등을 교육합니다. 교육 수강 시 JSP 및 Servlet 개발 또는 이에 관련된 경험이 필요하지만 포틀릿이나 JSF에 대한 경험은 없으셔도 무관합니다.

67.2.1. 선수 조건

이 과정은 Tomcat 컨테이너(Apache에 내장 또는 JBoss 어플리케이션 서버에 통합)가 사용되는 JBoss 어플리케이션 서버에서 기본 J2EE 웹 컨테이너(Servlet/JSP) 프로그래밍 기술 및 J2EE 웹 기반 다중 계층 어플리케이션 활용 경험이 필요합니다. 수강자는 다음 기술에 관한 개발 경험을 갖추고 있어야 합니다:

1. JNDI

2. Servlet 2.3/2.4 API

3. JSP 2.0 API

4. JBoss 어플리케이션 서버에서 J2EE 어플리케이션 개발 및 활용

5. 내장형(독립형) Tomcat 또는 통합형 Tomcat(JBossWeb)에서 웹 어플리케이션 활용

6. JDBC 및 EJB2.1 또는 EJB3.0에 대한 실무 지식

다른 관련 지식도 도움이 됩니다.

67.3. RHD167: JBoss - Hibernate Essentials

http://www.redhat.com/training/jboss/courses/rhd167.html

67.3.1. 선수 조건

1. 관계 종속 모델 이해

2. Java 언어 사용 능력

3. OOAD 개념 지식

4. UML 사용 능력

5. SQL 문법 이해

6. JDK 사용 및 명령 행에서 Java 실행 파일을 컴파일 및 실행할 수 있는 환경 생성

7. JDB 이해

J2EE 또는 Hibernate 지식은 필요하지 않으며 Hibernate 3.2 시리즈 기반 교육이 진행됩니다.

67.3.2. 교육 과정 요약

Hibernate Essentials 과정은 SQL 기반 데이터베이스 시스템을 사용하는 Java 개발자 또는 객체 지향 소프트웨어 개발에 입문하려는 데이터베이스 개발자를 대상으로 하여 Hibernate 또는 Java Persistence API 객체, 관련 종속 및 쿼리 서비스 구현 등을 교육합니다. 또한, ORM이 어떻게 시스템 성능에 영향을 미치고 SQL 데이터베이스 관리 시스템 및 종속 계층 성능을 향상시킬지에 대해 배우려는 데이터 관리자에게 매우 유용하며 Java Persistence JSR-220 sub-specification에 대한 JBoss Inc.구현 및 Hibernate 3라고 불리는 JBoss Inc. Hibernate 제품의 기본 API 버전 3.x에 대해 다룹니다.

67.4. RHD267: JBoss - Advanced Hibernate

http://www.redhat.com/training/jboss/courses/rhd267.html

JBoss Advanced Hibernate 교육 과정은 Hibernate O/R 매핑 프레임워크의 모든 기능 사용법을 교육합니다. 이 과정은 SQL 기반 데이터베이스 시스템을 사용하는 Java 개발자, 객체 지향 소프트웨어 개발에 입문하려는 데이터베이스 개발자, ORM이 어떻게 시스템 성능에 영향을 미치고 SQL 데이터베이스 관리 시스템 및 종속 계층 성능을 향상시킬지에 대해 배우려는 데이터 관리자를 주요 대상으로 하며 새로운 Hibernate 3 기능에 대해 가르칩니다.

67.4.1. 선수 조건

이 교육 과정을 등록하려면 다음과 같은 능력을 갖추고 있어야 합니다:

1. 기본 Hibernate 지식

2. Java 언어 사용 능력

3. OOAD 개념 지식

4. UML 사용 능력

5. SQL 문법 이해

6. JDK 사용 및 명령 행에서 Java 실행 파일을 컴파일 및 실행할 수 있는 환경 생성

7. JNDI 및 JDBC에 대한 지식 또는 실무 경험

8. 선수 조건은 아니지만 기본 EJB2.1 및 EJB3.0 지식 필요

9. Hibernate in Action(작가: Christian Bauer 및 Gavin King, 출판사: Manning)이라는 책을 미리 읽어보시면 좋습니다.

"The best part of the Advanced Hibernate course was networking with fellow engineers that had problems similar to my own, and working with a knowledgeable instructor to solve them."--Mike Pasternak, Consulting Engineer, United Switch & Signal

67.5. RHD261:JBoss for advanced J2EE developers

http://www.redhat.com/training/jboss/courses/rhd261.html

JBoss for Advanced J2EE Developers 교육 과정은 J2EE 전문가를 대상으로 JBoss 어플리케이션 서버 내부 구조를 이용하여 JBoss 어플리케이션 서버에서 J2EE 어플리케이션 성능 및 기능을 강화하는 방법을 교육합니다. 이 과정은 JMX를 비롯한 Microkernel 구조, 보안, 클러스터링, 정밀 개선과 같은 J2EE specification 관련 주제를 다룹니다.

67.5.1. 선수 조건

수강자는 JBoss for Advanced J2EE Developers 과정을 등록하기 전에 반드시 JBoss for Java Developers 과정을 이수하거나 Middleware Placement Exam에 합격해야 합니다. 개발자는 다음과 같은 실무 경험을 갖추고 있어야 합니다:

1. JNDI

2. JDBC

3. Servlets 및 JSP

4. Enterprise Java Bean

5. JMS

6. J2EE 보안 모델

7. JBoss 어플리케이션에서 J2EE 어플리케이션 개발 및 활용

8. ANT 및 XDoclet 경험 또는 이에 해당하는 능력

JMX 지식은 필수 조건은 아니지만 매우 유용할 것입니다. 이 과정은 JBoss 어플리케이션 서버 4.x 시리즈를 기반으로 합니다.

"I thought the training materials were well-organized, including both the handbook and the labs. The instructor frequently asked for feedback on material and pace. It was apparent that he cared about our understanding of the material."--Jeremy Prellwitz, SiRAS.com, USA

67.6. RH336: JBoss for Administrators

http://www.redhat.com/training/jboss/courses/rh336.html

67.6.1. 선수 조건

Windows 또는 리눅스(유닉스 기반) 운영체제에 대한 실무 경험이 필요합니다. 수강자는 반드시 다음과 같은 경험을 갖추고 있어야 합니다:

1. 디렉토리, 파일 생성 및 파일 저장 권한 수정

2. JDK 설치

3. 운영체제에 필요한 환경 변수 설정 (예, JAVA_HOME)

4. OS 기반 스크립트를 실행하여 Java 어플리케이션 시작

5. Java 아카이브 파일(jar 유틸리티) 생성 및 확장

J2EE 또는 JBoss 어플리케이션 서버 지식이 필요없습니다. 하지만, XML을 구성하여 Java 어플리케이션을 지원할 수 있는 능력을 반드시 갖추고 있어야 합니다.

67.6.2. 교육 과정 요약

JBoss for Administrators 교육 과정은 어플리케이션 지원을 담당하는 시스템 관리자, 구성 관리자, QA 관리자를 대상으로 JBoss 어플리케이션 서버(3.2 및 4.x 시리즈) 및 활용 어플리케이션을 구성하고 관리하는 능력을 교육합니다.

"The JBoss for Administrators course was a great balance of both lecture and labs. It is always nice to have hands on knowledge of the topics to make them seem more real and applicable."——Thomas Skowronek, Palm Harbor Homes, USA

67.7. RHD439: JBoss Clustering

http://www.redhat.com/training/jboss/courses/rhd439.html

Clustering 과정은 JBoss Enterprise Middleware System(JEMS)의 고가용성 서비스에 중점을 둔 4일간의 교육 과정입니다. JBoss 어플리케이션 서버가 JGroups 및 JBoss Cache에 적용되어 반응 및 문제 해결을 이루어내는 방법, JGroups 프로토콜 스택 구성, 개선, 구현 방법, 미들웨어 어플리케이션 구현에서 JBoss Cache 적용 방법, HTTP 로드 밸런싱에 필요한 mod_jk 사용 및 구성 방법 등에 대하여 배우게 됩니다. 또한, HA-JNDI, HA-JMS, HA-singleton와 같은 JBoss 어플리케이션 서버 고가용성 서비스에 대해서도 다루어 집니다.

67.7.1. 선수 조건

이 과정을 등록하기 전에 JBoss for Advanced J2EE Developers 과정을 이수하기를 권장합니다. 또한, 수강자는 최소 18개월의 J2EE 또는 다른 Java 미들웨어 기술을 사용한 실무 개발 경험 및 JBoss 어플리케이션 서버에 관한 일부 실무 경험을 갖추고 있어야 합니다. 최소 3년의 꾸준한 Java 프로그래밍 경험도 필요하며 기본적인 TCP/IP 개념도 이해하고 계셔야 합니다.

수강자는 다음 능력을 갖추고 있어야 합니다:

1. JTA, Transactions, Java Concurrency

2. EJB 2.1, JMS, Reliable Messaging Technologies

3. Apache httpd 및 일부 mod_jk 또는 mod-proxy 사용 경험

4. JBoss AS Microkernel 및 JMX 경험

5. TCP/IP, UDP, Multicasting 개념 이해

"The JBoss for Administrators course was very informative. Our instructor did a great job at answering our questions (some very specific to the student) while maintaining the course direction. I am very excited about applying what I have learned in the course."——Andy Beier, Arizona Statue University, USA

67.8. RHD449: JBoss jBPM

http://www.redhat.com/training/jboss/courses/rhd449.html

67.8.1. 설명

JBoss jBPM 과정은 비즈니스 분석가와 같이 일하여 비즈니스 프로세스를 jBPM 엔진을 사용한 J2EE 환경으로 통합하는 업무를 담당하는 시스템 아키텍처 및 개발자를 대상으로 합니다. 또한, JBoss jBPM 교육은 BPM 계획, 엔진 유형, 버즈워드 위치 설정에 관한 폭 넓은 지식을 전달합니다.

이 교육 과정을 모두 완료하면 다양한 실습 경험을 통하여 JBoss jBPM을 사용한 비즈니스 프로세스 개발 능력을 갖출 수 있습니다. 또한 이 과정에서 워크 플로우 엔진을 비교할 수 있는 충분한 지식을 습득할 수 있습니다.

67.8.2. 선수 조건

1. 이 과정을 등록하기 전에 반드시 Hibernate 어플리케이션 개발 경험 및 Hibernate에 쓰이는 간단한 Session Factory 구성, Hibernate Session 및 Transactional Demarcation 사용, Hibernate 객체에서 기본 쿼리 수행 능력을 갖추고 있어야 합니다.

2. Java 어플리케이션 개발 능력

3. 워크 플로우 및 비즈니스 프로세스 모델링(BPM)에 대한 사전 경험은 필요하지 않습니다.

4. JBoss 플러그인을 사용한 JBoss Eclipse 또는 Eclipse IDE 경험은 필수는 아니지만 유용합니다.

5. JUnit 테스트 프레임워크에 대한 기본 개념

67.9. RHD451 JBoss Rules

http://www.redhat.com/training/jboss/courses/rhd451.html

이 교육 과정은 Drools 3(JBoss Rules 3.0)에 사용되는 코어 엔진, 비즈니스 룰을 관리하는 데 필요한 다양한 기술 및 언어, J2SE 및 J2EE 어플리케이션에 룰 엔진 내장 방법 등을 교육합니다. 이 과정은 차후 JBoss Rules 배포판을 사용한 룰 관리 교육 과정에 필요한 필수 과정이 될 것입니다.

67.9.1. 선수 조건

1. 기본적인 Java 사용 능력

2. 스크립팅 엔진과 비교하여 추론 룰 엔진이 어떻게 구성되는지에 관한 개념 이해

3. Jboss Rules webinar 및 demo를 볼 수 있는 능력은 필수는 아니지만 유용합니다.

4. Java EE 통합 방법을 배우려는 수강자에게는 Java EE에 대한 경험이 필요합니다.

부록 A. Revision History

고침 11-0 Tue Feb 21 2012 **Jaromír Hradílek** jhradilek@redhat.com

Resolve BZ#749948: [Release Notes and Deployment Guide] Migration tooling from RHN Classic to Cert-based RHN for RHEL 5.
Resolve BZ#718608: MinorMod: FTP: Missing text fragment in vsftpd configuration documentation.
Resolve BZ#720387: MinorMod: The proc File System: Illogical parameter description.
Resolve BZ#720860: Update Deployment (Guide) in RHEL5 Build Tree.
Resolve BZ#760925: MinorMod: Network File System: Severely suboptimal timeo option in NFS mount examples (for TCP).
Resolve BZ#784754: MinorMod: Network Interfaces: typo - wrong tense in 15.3. Interface Control Scripts.
Resolve BZ#740916: MinorMod: The kdump Crash Recovery Service: Incorrect description of the crashkernel parameter.
Resolve BZ#767105: incorrect default action in kdump part.
Resolve BZ#714080: debug option for bonding can not be used in BONDING_OPTS in /etc/sysconfig/network-scripts/ifcfg-bondX.
Resolve BZ#769776: Documentation needs to be updated for "default shell" option in /etc/kdump.conf.
Resolve BZ#781441: /etc/securetty documentation is incorrect [rhel-5.7].

고침 10-0 Thu Jul 21 2011 **Jaromír Hradílek** jhradilek@redhat.com

Resolve BZ#720382: MinorMod: Network Interfaces: LINKDELAY parameter needs to be added to "Interface Configuration Files".
Resolve BZ#632028: MajorMod: Redundant Array of Independent Disks (RAID): Document mdadm Usage.
Resolve BZ#720009: MinorMod: LVM: Update screenshots in the "Manual LVM Partitioning" section.
Resolve BZ#711162: MinorMod: Network Interfaces: Incorrect static routes configuration.
Resolve BZ#707238: broadcast is calculated with ipcalc, not ifcalc.
Resolve BZ#678316: HOTPLUG network config file option is not documented.
Resolve BZ#562018: Ch.4 Redundant Array of Independent Disks (RAID) - screenshots need updating.
Resolve BZ#485033: iptables -p ALL --dport not allowed according to man 8 iptables.

고침 9-0 Thu Jan 13 2011 **Jaromír Hradílek** jhradilek@redhat.com

Resolve BZ#249485: 'fsid=num' is listed under NFS client options, but it is a server-only option.
Resolve BZ#253659: additional commands required when adding machines to domain.
Resolve BZ#453242: guide does not tell you which packages you need to run an NFS server.
Resolve BZ#504250: cell should have newline characters, it shouldn't be all on one line.
Resolve BZ#520650: /proc/loadavg documentation error.
Resolve BZ#584075: vsftp typo for text_userdb_names.
Resolve BZ#625384: bonding configuration SLAVE=bond0 is invalid.
Resolve BZ#644617: misspelled word.
Resolve BZ#645123: spelling Errors in Deployment Guide II.
Resolve BZ#595366: RFE: document Shared Subtrees.

고침 8-0 Thu July 30 2010 Douglas Silas dhensley@redhat.com

Resolve BZ#239313: document oom_adj and oom_score.

Resolve BZ#526502: correct quotaon instructions with proper, safe operating procedures.
Resolve BZ#551367: correct SELinux dhcpd_disable_trans description.
Resolve BZ#521215: clarify NFS interaction with portmapper, rpc.mountd, rpc.lockd and rpc.statd.
Resolve BZ#453875: various OpenSSH chapter corrections.
Resolve BZ#455162: correct zone example configuration file, description.
Resolve BZ#460767: make it a proper daemon.
Resolve BZ#600702: correct directories used for SSL key generation.

고침 7-0 Wed Sep 30 2009 Douglas Silas dhensley@redhat.com, **Jaromír Hradílek** jhradilek@redhat.com, Martin Prpic mprpic@redhat.com

Change heading titles to correspond with actual headings used in 'man rpm'.
Resolve BZ#499053: /usr/sbin/racoon is correct install path.
Remove any mention of 'pkgpolicy' in /etc/yum.conf as per BZ#237773.
Resolve BZ#455162: correct example zone file with regard to records, description.
Resolve BZ#510851: /proc/cmdline has confusing descriptions of sample output.
Resolve BZ#510847: page with multiple footnotes formatted incorrectly in online PDF.
Resolve BZ#214326: more detailed usage info concerning vsftpd banners and secueerity.
Resolve BZ#241314: formatting problems in screen elements.
Resolve BZ#466239: postfix connect-from-remote-host configuration fix.

고침 7-0 Mon Sep 14 2009 Douglas Silas dhensley@redhat.com

Resolve BZ#214326: Server Security FTP Banner instructions: questions re: vsftpd.conf.
Resolve BZ#466239: insert line into Postfix config file to allow connecting remotely.
Resolve BZ#499053: path for racoon daemon is /usr/sbin/racoon, not /sbin/racoon.
Resolve BZ#510847: missing footnotes in PDF output.
Resolve BZ#510851: rewrite /proc/cmdline minor section to make more sense.
Resolve BZ#515613: correct location of RHEL5 GPG keys and key details.
Resolve BZ#523070: various minor fixes; --redhatprovides to rpm -q --whatprovides.

고침 6-0 Wed Sep 02 2009 Douglas Silas dhensley@redhat.com

Resolve BZ#492539: "This directive is useful..." to "This directive must be used in machines containing more than one NIC to ensure...".
Resolve BZ#241314: re: kernel-pae and hugemem support on RHEL 4 and 5.
Resolve BZ#453071: incorrect tag use led to config files and other screen elements being displayed on single lines.
Resolve BZ#507987: clarify and correct statements about partitions being in use while resizing or removing.
Resolve BZ#462550: recommended amount of swap space, according to http://kbase.redhat.com/faq/docs/DOC-15252.
Resolve BZ#466239: line omitted from Postfix configuration meant connecting remotely failed
Resolving other MODIFIED BZs (fixed previously): 468483, 480324, 481246, 481247, 438823, 454841, 485187, 429989, 452065, 453466.

고침 5-0 Wed Jan 28 2009 Michael Hideo Smith mhideo@redhat.com

Resolves: #460981
Changing 64GB *tested* support to support for 16GB.

부록 B. Colophon

The manuals are written in DocBook XML v4.3 format.

Garrett LeSage created the admonition graphics (note, tip, important, caution, and warning). They may be freely redistributed with the Red Hat documentation.

Contributing Writers: John Ha (System Administration, Filesystems, Kernel), Joshua Wulf (Installation and Booting), Brian Cleary (Virtualization), David O'Brien (Security and SELinux), Michael Hideo (System Administration), Don Domingo (System Administration), Michael Behm (System Administration), Paul Kennedy (Storage), Melissa Goldin (Red Hat Network)

Honoring those who have gone before: Sandra Moore, Edward C. Bailey, Karsten Wade, Mark Johnson, Andrius Benokraitis, Lucy Ringland

Honoring engineering efforts: Jeffrey Fearn

Technical Editing: Michael Behm

Graphic Artist: Andrew Fitzsimon

The Red Hat Localization Team consists of the following people:

- East Asian Languages
  - Simplified Chinese
    - Tony Tongjie Fu
    - Simon Xi Huang
    - Leah Wei Liu
    - Sarah Saiying Wang
  - Traditional Chinese
    - Chester Cheng
    - Terry Chuang
    - Ben Hung-Pin Wu
  - Japanese
    - Kiyoto Hashida
    - Junko Ito
    - Noriko Mizumoto
    - Takuro Nagamoto
  - Korean
    - Eun-ju Kim
    - Michelle Kim
- Latin Languages

  - French
    - Jean-Paul Aubry
    - Fabien Decroux
    - Myriam Malga
    - Audrey Simons
    - Corina Roe
  - German
    - Jasna Dimanoski
    - Verena Furhuer
    - Bernd Groh
    - Daniela Kugelmann
    - Timo Trinks
  - Italian
    - Francesco Valente
  - Brazilian Portuguese
    - Glaucia de Freitas
    - Leticia de Lima
    - David Barzilay
  - Spanish
    - Angela Garcia
    - Gladys Guerrero
    - Yelitza Louze
    - Manuel Ospina
  - Russian
    - Yuliya Poyarkova
- Indic Languages
  - Bengali
    - Runa Bhattacharjee
  - Gujarati
    - Ankitkumar Rameshchandra Patel

    - Sweta Kothari
- Hindi
    - Rajesh Ranjan
- Malayalam
    - Ani Peter
- Marathi
    - Sandeep Shedmake
- Punjabi
    - Amanpreet Singh Alam
    - Jaswinder Singh
- Tamil
    - I Felix
    - N Jayaradha