

A horizontal decorative bar consisting of a solid red segment on the left, followed by four segments with abstract, colorful patterns in shades of blue, orange, green, and purple.

Implementing Redundancy in Red Hat Network Satellite

By Vladimir Zlatkin

Abstract

This document will help you create two identical Red Hat® Network (RHN) Satellites connected to one external Oracle® database. These RHN Satellites will function in redundant mode.

November 2005

Table of Contents

Requirements.....	3
Red Hat Network.....	3
Hardware.....	3
Software.....	3
Installation.....	3
Monitoring Configuration.....	5
Redundancy.....	5
Configuring the clients	6
Conclusion	6
Appendix – Useful documentation.....	7
Appendix - Troubleshooting	7



Requirements

Before beginning, review the necessary requirements listed below:

Red Hat Network

Your RHN Hosted account must have two Red Hat® Enterprise Linux® channel entitlements and two RHN Satellite entitlements.

Hardware

This configuration requires two systems that will become RHN Satellites and a single system that will serve as an external Oracle® database. Hardware requirements for each of these systems can be found in the *Satellite Installation Guide*.

Software

As with hardware, software requirements for RHN Satellite, and the external Oracle® database can be found in the *Satellite Installation Guide*. A round-robin mechanism that switches between the two RHN Satellites in case of malfunction is also required. Possible solutions include Domain Name System (DNS) or Red Hat Cluster Suite. Both RHN Satellites must be given the same hostname.

Installation

You can use `rsync` to synchronize data or you can choose to mount directories via Network File System (NFS) from an external server. If you use the latter solution, you should set up an NFS server and mount `/var/satellite/` on both machines.

First, install the master RHN Satellite. If necessary, change the hostname of the first machine (the master Satellite) to be that of your high-availability solution. To accomplish this, it is sufficient to modify `/etc/hosts` and set that hostname to point to the local IP. In addition, you can either modify `/etc/sysconfig/network` and reboot the system or run `/bin/hostname` with the hostname as an argument. Next, follow the steps described in the *Satellite Installation Guide* to install your first RHN Satellite. Throughout the web interface part of the installation process, use the high-availability hostname.

Second, install the slave RHN Satellite. Change the hostname of the second machine; it should have the same hostname as the master Satellite. Follow the steps above to accomplish this. Next, follow the steps outlined in the *Satellite*

Installation Guide to install the second RHN Satellite. In the web interface portion of the installation, use the same entries as with the first Satellite; most importantly, use the same database references. Do not clear and repopulate the database. To prevent this, uncheck the check box labeled 'Clear and Repopulate.' Optionally, you can skip the 'Satellite Certificate Generation' phase of the installation and uncheck the check box labeled 'Perform Satellite Sync.'

Note: *The last installation step will not allow you to create an Organization Administrator account. The Organization Administrator account generated in the first installation functions for both Satellites.*

Third, set up your failover mechanism. The end result should be this: From an external machine you should be able to reach each of the Satellites with the same hostname. It is best to use the hostname of the master RHN Satellite server as the hostname for both machines. It is also a good idea to ensure that if you access the hostname from one of the Satellites, you return to localhost and not the other machine. You can do this by editing `/etc/hosts` or, preferably, with a firewall rule.

During the installation of the Satellites, you should have generated two SSL certificates. If this is not the case, use the `rhns-ssl-tool` to generate certificates. For these steps, you must have access to the SSL build tree, which by default is in `/root/ssl-build/`. (The *Client Configuration Guide* suggests a more secure location; you can move the SSL tree after completing the steps below.)

To synchronize the certificates, copy these files from the master Satellite to the slave Satellite:

```
/var/www/html/pub/*  
  
/root/ssl-build/<sat-hostname>/rhns-org-httpd-ssl-key-pair-  
<sat-host-name>-1.0-<version>.noarch.rpm  
  
/etc/jabberd/server.pem
```

Next, deploy the RPM on the slave server, optionally removing a previous version, by issuing the following command:

```
rpm -Uvh rhns-org-httpd-ssl-key-pair-<sat-host-name>-1.0-  
<version>.noarch.rpm
```

Monitoring Configuration

If you chose to enable monitoring scout functionality during the installation process, a list of the scouts is available in the web interface located by clicking the 'Monitoring' link in the top toolbar followed by clicking 'Scout Config Push.' Due to the identical installations, both scouts are listed with the name 'RHN Monitoring Satellite.' To differentiate between the two instances, edit the string stored in the RHN_SAT_CLUSTER.DESCRPTION column in the database. Consult your database administrator to change this value. It is helpful to provide your DBA with the IP addresses of each RHN Satellite system and an appropriate description of each.

Redundancy

When you create monitoring probes, you must choose one and only one scout from which to run them. If during a maintenance window that scout becomes unavailable, so do the probes, and you will lose a critical early warning system. There are three ways to solve this problem:

The first option is to create identical copies of all probes and probe suites on each monitoring scout. Probe suites are advantageous because they allow you to deploy multiple, identical probes to many machines in a batch. When associating systems with probe suites, be sure to choose different monitoring scouts.

The second option is to delegate the task of gathering monitoring data to an RHN Proxy. To enable this feature, consult the *RHN Proxy Server Installation Guide*. When creating monitoring probes, associate them with the scout that runs on an RHN Proxy — do not associate them with either of the RHN Satellites.

The third option is to edit the database table that contains the association of the monitoring probe from the failed scout instance to the functional one. With the help of your DBA, you can change this value in the RHN_CHECK_PROBE table. In the RHN_SAT_CLUSTER table, look up by IP address the *recids* of the failed and functional satellites. Next, look in the SAT_CLUSTER_ID column of the RHN_CHECK_PROBE table and change all rows from the failed *recid* to the functional one. The commands will look something like this:

```
update rhn_check_probe set sat_cluster_id=<functional  
recid> where sat_cluster_id=<failed recid>
```

Afterward, schedule a 'Scout Config Push' from the Monitoring tab in the web interface.

On the second RHN Satellite, run these commands:

```
/sbin/service rhn-satellite restart
/sbin/service taskomatic stop
```

The installation is now complete. Make sure that this is the case with some sanity tests in the web interface and on the client side. At this point, extend or customize the Satellites to your needs: set up PAM authentication, create `cron` jobs, or set up a PXE environment. Replicate these features precisely between the two Satellites.

Configuring the clients

Now that you have set up redundant RHN Satellites, you will want to configure your client systems to take advantage of both machines.

The latest version of `up2date` is capable of connecting to multiple RHN Servers. If the first server becomes unavailable, `up2date` automatically uses the other. Specify a semi-colon separated list of RHN Satellites in the `/etc/sysconfig/rhn/up2date` file in the `serverURL` configuration parameter.

If you chose to enable monitoring capabilities with the steps listed above, you should also enable your client machines to be monitored by multiple scouts. To do this you will first need the public key of each scout. Browse to the 'Monitoring' tab of the web interface, followed by the 'Scout Config Push' subtab. Click on the description of a scout and the web interface displays the public key of that scout. Copy the string and add it to `/opt/nocpulse/.ssh/authorized_keys` on the client. Repeat these steps on the client for each scout that you have configured.

Conclusion

At this point you have two RHN Satellites connected to a single external Oracle® database. They are both fully functional and, depending on your clustering solution, may be used simultaneously. In the case of a master RHN Satellite failure, remember to start `taskomatic` on the slave Satellite by issuing the following command:

```
/sbin/service taskomatic start
```

Appendix – Useful documentation

RHN Satellite Server Installation Guide

<https://rhn.redhat.com/rhn/help/satellite/index.jsp>

Red Hat Network Reference Guide

<https://rhn.redhat.com/rhn/help/reference/index.jsp>

Proxy Guide

<https://rhn.redhat.com/help/proxy/>

Client Configuration Guide

<https://rhn.redhat.com/help/client-config/>

Red Hat Knowledgebase

<http://kbase.redhat.com>

Appendix - Troubleshooting

Following is a list of files and database locations that contain the RHN Satellite hostname:

RHN Push

`/etc/jabberd/c2s.xml`

`/etc/jabberd/sm.xml`

Database Tables

`RHN_CONFIG_MACRO`

`RHN_SAT_CLUSTER`

If you experience a difference in the functionality of your two RHN Satellites, make sure that you have replicated the following files:

SSL

`/var/www/html/pub/*`

`/root/ssl-build/<sat-hostname>/rhn-org-httpd-ssl-key-pair-<sat-host-name>-1.0-<version>.noarch.rpm`

RHN Push

`/etc/jabberd/c2s.xml`

`/etc/jabberd/sm.xml`

`/etc/jabberd/server.pem`

PAM

/etc/pam.d/rhn-satellite (this is the default file)

PXE

/tftpboot/*

Cron

/etc/cron*