The Amanda Backup System

NAME

amanda − Advanced Maryland Automatic Network Disk Archiver

SYNOPSIS

amdump *config*
amflush [**−f**] *config*
amcleanup *config*
amrecover [[**−C**] *config*] [*options*]
amrestore [*options*] *tapedevice* [*hostname* [*diskname*]]
amlabel *config label* [**slot** *slot*]
amcheck [*options*] *config*
amadmin *config command* [*options*]
amtape *config command* [*options*]
amverify *config*
amrmtape [*options*] *config label*
amstatus *config* [*options*]

DESCRIPTION

Amanda is the "Advanced Maryland Automatic Network Disk Archiver". This manual page gives an overview of the Amanda commands and configuration files for quick reference.

Here are all the Amanda commands. Each one has its own manual page. See them for all the gritty details.

amdump

Take care of automatic Amanda backups. It is normally executed by **cron** on a computer called the *tape server host* and requests backups of file systems located on *backup clients*. **Amdump** backs up all disks in the *disklist* file (discussed below) to tape or, if there is a problem, to a special *holding disk*. After all backups are done, **amdump** sends mail reporting failures and successes.

amflush

Flush backups from the holding disk to tape. **Amflush** is used after **amdump** has reported it could not write backups to tape for some reason. When this happens, backups stay in the holding disk. After the tape problem is corrected, run **amflush** to write backups from the holding disk to tape.

amcleanup

Clean up after an interrupted **amdump**. This command is only needed if **amdump** was unable to complete for some reason, usually because the tape server host crashed while **amdump** was running.

amrecover

Provides an interactive interface to browse the Amanda index files and select which tapes to recover files from. Can also run **amrestore** and the system restore program (e.g. **tar**) in some cases.

amrestore

Read an Amanda tape, searching for requested backups. **Amrestore** is suitable for everything from interactive restores of single files to a full restore of all partitions on a failed disk.

amlabel

Write an Amanda format label onto a tape. All Amanda tapes must be labeled with **amlabel**. **Amdump** and **amflush** will not write to an unlabeled tape (see TAPE MANAGEMENT below).

amcheck

Verify the correct tape is in the tape drive and all file systems on all backup client systems are ready to be backed up. Can optionally be run by **cron** before **amdump** so someone will get mail warning that backups will fail unless corrective action is taken.

amadmin

Take care of administrative tasks, like finding out which tapes are needed to restore a filesystem, forcing hosts to do full backups of selected disks, and looking at schedule balance information.

amtape
Take care of tape changer control operations, like loading particular tapes, ejecting tapes, and scanning the tape rack.

amverify
Check Amanda backup tapes for errors (GNU tar format backups only).

amrmtape
Delete a tape from the *tapelist* and Amanda database.

amstatus
Give the status of a running amdump.

CONFIGURATION

There are three user−editable files that control the behavior of Amanda. The first is *amanda.conf,* the main configuration file. It contains parameters to customize Amanda for your site. Second is the *disklist* file, which lists hosts and disk partitions to back up. Third is the *tapelist* file, which lists tapes that are currently active. These files are described in more detail in the following sections.

All files are stored in a *config* directory under /usr/amanda/etc/amanda. A site will often have more than one configuration. For example, a site might have a *normal* configuration for every−day backups and an *archive* configuration for infrequent full archival backups. The configuration files would be stored under /usr/amanda/etc/amanda/normal/ and /usr/amanda/etc/amanda/archive/, respectively.

All log and database files generated by Amanda go in corresponding directories somewhere. An example location is under /usr/adm/amanda. In our example, the files would go in /usr/adm/amanda/normal and /usr/adm/amanda/archive.

CONFIG FILE PARAMETERS

There are a number of configuration parameters that control the behavior of the Amanda programs. All have default values, so you need not specify the parameter in your *amanda.conf* if the default is suitable.

Lines starting with # are ignored, as are blank lines. Keywords are case insensitive, i.e. **mailto** and **MailTo** are treated the same.

Integer arguments may have one of the following (case insensitive) suffixes, some of which have a multiplier effect:

b byte bytes
Some number of bytes.

g gb gbytes gigabytes
Some number of gigabytes (bytes*1024*1024*1024).

k kb kbytes kilobytes
Some number of kilobytes (bytes*1024).

kps kbps
Some number of kilobytes per second (bytes*1024).

m mb meg mbytes megabytes
Some number of megabytes (bytes*1024*1024).

mps mbps
Some number of megabytes per second (bytes*1024*1024).

tape tapes
Some number of tapes.

day days
Some number of days.

week weeks
Some number of weeks (days*7).

You may also use the value **inf** where an integer is expected to mean an infinite amount.

Boolean arguments may have any of the values **y**, **yes**, **t**, **true** or **on** to indicate a true state, or **n**, **no**, **f**, **false** or **off** to indicate a false state. If no argument is given, **true** is assumed.

org *"string"*
Default: *DailySet1*. A descriptive name for the configuration. This string is used for mail reports. Each Amanda configuration should have a different value.

mailto *"string"*
Default: *operators*. A comma separated list of recipients for mail reports.

dumpcycle *int*
Default: *10 days*. The number of days in the backup cycle. Each disk will get a full backup no less often than this many days.

runspercycle *int*
Default: *same as dumpcycle*. The number of amdump runs in **dumpcycle** days. A value of 0 means the same value as **dumpcycle.** A value of -1 means guess the number of runs from the tapelist file, which is the number of tape used in the last **dumpcycle** days / **runtapes.**

tapecycle *int*
Default: *15 tapes*. The number of tapes in your active tape cycle. This must be at least one larger than the number of Amanda runs done during a dump cycle (see the **dumpcycle** parameter) times the number of tapes used per run (see the **runtapes** parameter).

For instance, if you have **dumpcycle** set to 14 days, do one Amanda run per day and have **runtapes** set to one, **tapecycle** must be at least 15 (14 days * one run/day * one tape/run + one tape).

In practice you should have five or more extra tapes to allow for schedule adjustments or disaster recovery.

dumpuser *"string"*
Default: *daemon*. The login name Amanda uses to run the backups. The backup client hosts must allow access from the tape server host as this user via *.rhosts* or *.amandahosts*, depending on how the Amanda software was built.

tapedev *"string"*
Default: */dev/null*. The path name of the tape device. If you have configured a tape changer (see the **tpchanger** option), this option is not used and should be commented out of your configuration file.

rawtapedev *"string"*
Default: */dev/null*. The path name of the raw tape device. It is only used if Amanda is compiled for Linux machines with floppy tapes. It is needed for QIC volume table operations.

tpchanger *"string"*
Default: *none*. The name of the tape changer. If you do not have a tape changer, this option is not used and should be commented out of your configuration file.

changerdev *"string"*
Default: */dev/null*. A tape change configuration parameter. Usage depends on the particular changer defined with the **tpchanger** option.

changerfile *"string"*
Default: */usr/adm/amanda/log/changer−status*. A tape change configuration parameter. Usage depends on the particular changer defined with the **tpchanger** option.

runtapes *int*
Default: *1*. The maximum number of tapes used in a single run. If you do not have a tape changer, this option is not used and should be commented out of your configuration file.

labelstr *"string"*
Default: *.**. The tape label constraint regular expression. All tape labels generated and used by this configuration must match the regular expression. If multiple configurations are run from the

same tape server host, it is helpful to constrain their labels to different strings (for example, "DAILY[0−9]*" vs. "ARCHIVE[0−9]*") to avoid overwriting each other's tapes.

tapetype *"string"*
Default: *EXABYTE*. The type of tape drive associated with **tapedev** or **tpchanger**. This refers to one of the defined **tapetype**s in the config file, which specify various tape parameters, like the **length**, **filemark** size, and **speed** of the tape media and device.

etimeout *int*
Default: *300 seconds*. Amount of time per disk on a given client that **planner** will wait to get the dump size estimates. For instance, with the default of 300 seconds and four disks on client A, **planner** will wait up to 20 minutes for that machine. A negative value will be interpreted as a total amount of time, instead of a per-disk value.

netusage *int*
Default: *300 Kbps*. The maximum network bandwidth allocated to Amanda, in Kbytes per second. See also the **interface** section.

inparallel *int*
Default: *10*. The maximum number of backups that Amanda will attempt to run in parallel. Amanda will stay within the constraints of network bandwidth and holding disk space available, so it doesn't hurt to set this number a bit high. Some contention can occur with larger numbers of backups, but this effect is relatively small on most systems.

maxdumps *int*
Default: *1*. The maximum number of backups from a single host that Amanda will attempt to run in parallel. See also the **dumptype** section and **inparallel** option.

bumpsize *int*
Default: *10 Mbytes*. The minimum savings required to trigger an automatic bump from one incremental level to the next. If Amanda determines that the next higher backup level will be this much smaller than the current level, it will do the next level. See also the **bumpmult** option.

bumpmult *float*
Default: *1.5*. The *bump multiplier*. Amanda multiplies **bumpsize** by this factor for each level. This prevents active filesystems from bumping too much by making it harder to bump to the next level. For example, with the default **bumpsize** and **bumpmult** set to 2.0, the bump threshold will be 10 Mbytes for level one, 20 Mbytes for level two, 40 Mbytes for level three, and so on.

bumpdays *int*
Default: *2*. To insure redundancy in the dumps, Amanda will keep filesystems at the same incremental level for at least **bumpdays** days, even if the bump threshold criteria are met.

diskfile *"string"*
Default: *disklist*. The file name for the *disklist* file.

infofile *"string"*
Default: */usr/adm/amanda/curinfo*. The file or directory name for the historical information database.

logdir *"string"*
Default: */usr/adm/amanda*. The directory for the **amdump** and **log** files.

indexdir *"string"*
Default */usr/adm/amanda/index*. The directory where index files are stored. Index files are only generated for filesystems whose **dumptype** has the **index** option enabled.

tapelist *"string"*
Default: *tapelist*. The file name for the active *tapelist* file.

reserve *number*
Default: *100*(percent) The amount of holding-disk space that should not be used for full backups if no tape is available. By default, when there is no tape to write to, degraded mode (incremental)

backups will be performed to the holding disk. If you wish to also allow full backups in this case, you may reduce the amount of holding disk space reserved for incrementals.

holdingdisk

Define parameters for a holding disk. The syntax is:

```
holdingdisk name {
  holdingdisk−option holdingdisk−value
  ...
}
```

The options and values are:

comment *"string"*
Default: *none*. A comment string describing this holding disk.

directory *"disk"*
Default: */dumps/amanda*. The path to this holding area.

use *int* Default: *10 Mbytes*. Amount of space that can be used in this holding disk area. If the value is negative, Amanda will use all available space minus that value.

chunksize *int*
Default: *-1*. Holding disk chunk size. Dumps larger than the specified size will be stored in multiple holding disk files. The size of each chunk will not exceed the specified value. However, even though dump images are split in the holding disk, they are concatenated as they are written to tape, so each dump image still corresponds to a single continuous tape section.

If you specify 0, Amanda will create holding disk chunks as large as ((INT_MAX/1024)-64) Kbytes.

A negative value causes filesystems larger than the absolute value to be dumped directly to tape, without using the holding disk. The default value, -1, is interpreted as -((INT_MAX/1024-64) Kbytes. which causes Amanda to dump filesystems larger than slightly less than 2 Gbytes (assuming 32 bit integers) directly to tape. The default is provided mainly for backward compatibility. You are encouraged to specify a positive value smaller than the maximum file size in the holding disk filesystem.

Each holding disk chunk includes a 32 Kbyte header, so the minimum chunk size is 64 Kbytes (but that would be really silly).

Operating systems that are limited to a maximum file size of 2 Gbytes actually cannot handle files that large. They must be at least one byte less than 2 Gbytes. Since Amanda works with 32 Kbyte blocks, and to handle the final read at the end of the chunk, the chunk size should be at least 64 Kbytes (2 * 32 Kbytes) smaller than the maximum file size, e.g. 2047 Mbytes.

DUMPTYPE SECTION

The config file may define multiple sets of backup options and refer to them by name from the **disklist** file. For instance, you might define one set of options for file systems that can benefit from high compression, another set that does not compress well, another set for file systems that should always get a full backup and so on.

A set of backup options are entered in a **dumptype** section, which looks like this:

```
define dumptype name {
  dumptype−option dumptype−value
  ...
}
```

Name is the name of this set of backup options. It is referenced from the *disklist* file.

Some of the options in a **dumptype** section are the same as global options. The global option value is used to set the default for all **dumptype** sections. For instance, setting **dumpcycle** to 50 in the main part of the config file causes all the **dumptype** sections to start with that value, but you may change it on a section by section basis.

The dumptype options and values are:

auth *"string"*
: Default: *bsd*. Type of authorization to perform between tape server and backup client hosts. May be **krb4** to use Kerberos−IV authorization.

comment *"string"*
: Default: *none*. A comment string describing this set of backup options.

comprate *float* [, *float*]
: Default: *0.50*, *0.50*. The expected full and incremental compression factor for dumps. Normally you should not set this because Amanda keeps the actual full and partial dump compression rates in the historical database.

compress [client|server] *"string"*
: Default: *client fast*. If Amanda does compression of the backup images, it can do so either on the backup client host before it goes to the holding disk or on the tape server host as it goes from the holding disk to tape. Which place to do compression (if at all) depends on how well the dump image usually compresses, the speed and load on the client or server, network capacity, holding disk capacity, tape hardware compression, etc.

For either type of compression, Amanda also allows you to select one of two styles of compression. **Best** is the best compression available, often at the expense of CPU overhead. **Fast** is often not as good a compression as **best**, but usually less CPU overhead.

So the **compress** options line may be one of:

```
compress none
compress [client] fast
compress [client] best
compress server fast
compress server best
```

Note that some tape devices do compression and this option has nothing to do with whether that is enabled, although you may want to disable Amanda compression if the drive is also going to do it. The **tapedev** path usually determines whether hardware compression is enabled.

dumpcycle *int*
: Default: *10 days*. The number of days in the backup cycle. Each disk using this set of options will get a full backup no less often than this many days. Setting this to zero forces a full backup each run.

exclude [list] *"string"*
: Default: *none*. The string is passed as a value to the GNU tar **−−exclude** or **−−exclude−list** option. In the case of **exclude list**, the file must be available on the backup client host.

holdingdisk *"boolean"*
: Default: *yes*. Whether a holding disk should be used for these backups or whether they should go directly to tape.

ignore *"boolean"*
: Default: *no*. Whether disks associated with this backup type should be backed up or not. This option is useful when you share a *disklist* file among several configurations, some of which should not back up all the listed file systems.

index *"boolean"*
Default: *no*. Whether an index (catalogue) of the backup should be generated and saved in **indexdir**.

kencrypt *"boolean"*
Default: *no*. Whether the backup image should be encrypted by Kerberos as it is sent across the network from the backup client host to the tape server host.

maxdumps *"int"*
Default: *1*. The maximum number of backups from a single host that Amanda will attempt to run in parallel. See also the **inparallel** option.

priority *"string"*
Default: *medium*. When there is no tape to write to, Amanda will do incremental backups in priority order to the holding disk. The priority may be **high**, **medium** or **low**.

program *"string"*
Default: *DUMP*. The type of backup to perform. Valid values are **DUMP** for the native operating system backup program, and **GNUTAR** to use GNU tar.

record *"boolean"*
Default: *yes*. Whether to ask the backup program to update its database (e.g. */etc/dumpdates*) of time stamps. You normally want this enabled for daily backups and turned off for periodic archival runs.

skip−full *"boolean"*
Default: *no*. If *true* and **planner** has scheduled a full backup, these disks will be skipped.

skip−incr *"boolean"*
Default: *no*. If *true* and **planner** has scheduled an incremental backup, these disks will be skipped.

starttime *"int"*
Default: *none*. Backups will not start until after this time of day. The value should be hh*100+mm, e.g. 6:30PM (18:30) would be entered as **1830**.

strategy *"string"*
Default: *standard*. Strategy to use when planning what level of backup to run next. Values are:

standard
The standard Amanda schedule.

nofull Never do full backups, only level 1 incrementals.

noinc Never do incremental backups, only full dumps.

skip Never do backups (useful when sharing the *disklist* file).

The following **dumptype** entries are predefined by Amanda:

```
define dumptype no-compress {
    compress none
}
define dumptype compress-fast {
    compress client fast
}
define dumptype compress-best {
    compress client best
}
define dumptype srvcompress {
    compress server fast
}
define dumptype bsd-auth {
```

```
    auth bsd
}
define dumptype krb4-auth {
    auth krb4
}
define dumptype no-record {
    record no
}
define dumptype no-hold {
    holdingdisk no
}
define dumptype no-full {
    skip-full yes
}
```

In addition to options in a **disktype** section, you may also put one or more other **disktype** names, which make this **disktype** inherit options from other **disktype**s. For instance, you might want two sections that are the same except for the **record** option:

```
define dumptype normal {
    comment "Normal backup, no compression, do indexing"
    no-compress
    index yes
    maxdumps 2
}
define dumptype testing {
    comment "Test backup, no compression, do indexing, no recording"
    normal
    no-record
}
```

TAPETYPE SECTION

The config file may define multiple types of tape media and devices. The information is entered in a **tapetype** section, which looks like this in the config file:

```
define tapetype name {
  tapetype-option tapetype-value
  ...
}
```

Name is the name of this type of tape medium/device. It is referenced from the **tapetype** option in the main part of the config file.

The tapetype options and values are:

comment *"string"*
: Default: *none*. A comment string describing this set of tape information.

filemark *"int"*
: Default: *1000 bytes*. How large a file mark (tape mark) is, measured in bytes. If you only know the size in some linear measurement (e.g. inches), convert it to bytes using the device density.

length *"int"*
: Default: *2000 kbytes*. How much data will fit on a tape.

speed *"int"*
: Default: *200*. How fast the drive will accept data, in bytes per second.

In addition to options, you may also put another **tapetype** name, which makes this **tapetype** inherit options from another **tapetype**. For instance, the only difference between a DLT4000 tape drive using Compact−III tapes and one using Compact−IV tapes is the length of the tape. So they could be entered as:

```
define tapetype DLT4000-III {
    comment "DLT4000 tape drives with Compact-III tapes"
    length 12500 mbytes         # 10 Gig tapes with some compression
    filemark 2000 kbytes
    speed 1536 kps
}
define tapetype DLT4000-IV {
    comment "DLT4000 tape drives with Compact-IV tapes"
    DLT4000-III
    length 25000 mbytes         # 20 Gig tapes with some compression
}
```

INTERFACE SECTION

The config file may define multiple types of network interfaces. The information is entered in an **interface** section, which looks like this:

```
define interface name {
  interface-option  interface-value
  ...
}
```

Name is the name of this type of network interface. It is referenced from the *disklist* file.

Note that these sections define network interface characteristics, not the actual interface that will be used. Nor do they impose limits on the bandwidth that will actually be taken up by Amanda. Amanda computes the bandwidth each file system backup will take based on the estimated size and time, then compares that plus any other running backups with the limit as another of the criteria when deciding whether to start the backup. Once a backup starts, it is free to use as much of the network as it can.

The interface options and values are:

comment *"string"*
: Default: *none*. A comment string describing this set of network information.

use *"int"*
: Default: *300 Kbps*. The speed of the interface in Kbytes per second.

In addition to options, you may also put another **interface** name, which makes this **interface** inherit options from another **interface**. At the moment, this is of little use.

DISKLIST FILE

The *disklist* file determines which disks will be backed up by Amanda. The file contains one line per disk:

hostname diskdevice dumptype [*spindle* [*interface*]]

Lines starting with # are ignored, as are blank lines. The fields have the following meanings:

hostname
: The name of the host to be backed up.

diskdevice
: The name of the disk device to be backed up. It may be a full device name, a device name without the */dev/* prefix, e.g. *sd0a*, or a mount point such as */usr*.

dumptype
: Refers to a *dumptype* defined in the *amanda.conf file*. *Dumptype*s specify backup related parameters, such as whether to compress the backups, whether to record backup results in */etc/dumpdates*, the disk's relative priority, etc.

spindle
: Default: *−1*. A number used to balance backup load on a host. Amanda will try not to run multiple backups on a given spindle.

interface
Default: *local*. The name of a network interface definition in the *config file*, used to balance network load.

TAPE MANAGEMENT

The *tapelist* file contains the list of tapes in active use. This file is maintained entirely by Amanda and should not be created or edited for normal operation. It contains lines of the form:

YYYYMMDD label

Where *YYYYMMDD* is the date the tape was written, and *label* is a label for the tape as written by **amlabel**.

Amdump and **amflush** will refuse to write to an unlabeled tape, or to a labeled tape that is considered active. There must be more tapes in active rotation (see the **tapecycle** option) than there are runs in the backup cycle (see the **dumpcycle** option) to prevent overwriting a needed backup file when a problem occurs.

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amadmin(8), amcheck(8), amcheckdb(8), amcleanup(8), amdump(8), amflush(8), amlabel(8), amoverview(8), amrecover(8), amrestore(8), amstatus(8), amtoc(8)

NAME

amadmin − administrative interface to control Amanda backups

SYNOPSIS

amadmin *config command* [*command options*]

DESCRIPTION

Amadmin performs various administrative tasks on the *config* Amanda configuration.

See the *amanda*(8) man page for more details about Amanda.

COMMANDS

All commands that take *hostname* [*disks*] parameters operate on all the disks in the *disklist* for that *hostname* if no disks are specified.

Disks are regular expressions, so "sd0" will match disks "sd0a" and "sd0g" and "/mnt" will match "/mnt", "/mnt/a", and "/mnt/b". To match only "/mnt", use "^/mnt$". This is the same mechanism used by *amrestore(8)*.

version
Show the current version and some compile time and runtime parameters.

force *hostname* [*disks*]
Force the *disks* on *hostname* to do a full level 0 backup during the next Amanda run.

unforce *hostname* [*disks*]
Undo a previous **force** command.

reuse [*tapelabel*]
The tapes *tapelabel* will be reuse in the cycle.

no-reuse [*tapelabel*]
The tapes *tapelabel* will never be reuse in the cycle.

find [**--sort hkdlb**] *hostname* [*disks*]
Display all backups currently on tape for each disk. The tape label, file number, and status are displayed.

The **--sort** option changes the sort order using the following flags:

h	host name
k	disk name
d	dump date
l	backup level
b	tape label

An uppercase letter reverses the sort order for that key. The default sort order is **hkdlb**.

delete *hostname disks*
Delete the specified *disks* from the Amanda database.

Note: if you do not also remove the disk from the *disklist* file, Amanda will treat it as a new disk on the next run.

tape Display the label of the tape Amanda is expecting to write to next. See also *amcheck*(8).

bumpsize
Display the current bump threshold parameters, calculated for all backup levels.

balance
Display the distribution of full backups throughout the dump schedule.

export [*hostname* [*disks*]]
Convert records from the Amanda database to a text format that may be transmitted to another Amanda machine and **import**ed.

import Convert **export**ed records read from standard input to a form Amanda uses and insert them into the database on this machine.

disklist [*hostname* [*disks*]]
Display the *disklist* information for each of the *disks* on *hostname* (or all hosts). Mostly used for debugging.

info *[* hostname *[* disks
Display the database record for each of the *disks* on *hostname* (or all hosts). Mostly used for debugging.

EXAMPLES

Request three specific file systems on *machine-a* get a full level 0 backup during the next Amanda run.

Note the use of "^/$" to get the root file system. Without the extra regular expression characters, just "/" would have matched all the file systems on *machine-a*.

```
$ amadmin DailySet1 force machine-a "^/$" /var /usr
amadmin: machine-a:/ is set to a forced level 0 tonight.
amadmin: machine-a:/var is set to a forced level 0 tonight.
amadmin: machine-a:/usr is set to a forced level 0 tonight.
```

Request all file systems on *machine-b* get a full level 0 backup during the next Amanda run.

```
$ amadmin DailySet1 force machine-b
amadmin: machine-b:/ is set to a forced level 0 tonight.
amadmin: machine-b:/var is set to a forced level 0 tonight.
amadmin: machine-b:/usr is set to a forced level 0 tonight.
amadmin: machine-b:/home is set to a forced level 0 tonight.
```

Undo the previous **force** request for */home* on *machine-b*. The other file systems will still get a full level 0 backup.

```
$ amadmin DailySet1 unforce machine-b /home
amadmin: force command for machine-b:/home cleared.
```

Locate tapes holding */var* from *machine-c*. The *file* column tells you which file on the tape has the backup (file number zero is a tape label). The *status* column tells you whether the backup was successful or had some type of error.

```
$ amadmin DailySet1 find machine-c /var
date        host      disk lv tape    file status
1997-11-09  machine-c /var  0 000110     9 OK
1997-11-08  machine-c /var  2 000109     2 OK
1997-11-07  machine-c /var  2 000108     2 OK
1997-11-06  machine-c /var  2 000107     2 OK
1997-11-05  machine-c /var  2 000106     3 OK
1997-11-04  machine-c /var  2 000105     2 OK
1997-11-03  machine-c /var  2 000104     2 OK
1997-11-02  machine-c /var  2 000103     2 OK
1997-11-01  machine-c /var  1 000102     5 OK
1997-10-31  machine-c /var  1 000101     3 OK
```

Forget about the */workspace* disk on *machine-d*. If you do not also remove the disk from the *disklist* file, Amanda will treat it as a new disk on the next run.

```
$ amadmin DailySet1 delete machine-d /workspace
amadmin: machine-d:/workspace deleted from database.
amadmin: NOTE: you'll have to remove these from the disklist yourself.
```

Find the next tape Amanda will use (in this case, *123456*).

```
$ amadmin DailySet1 tape
```

```
The next Amanda run should go onto tape 123456 or a new tape.
```

Show how well full backups are balanced across the dump cycle. The *due-date* column is the day the backups are due for a full backup. *#fs* shows the number of filesystems doing full backups that night, and *orig KB* and *out KB* show the estimated total size of the backups before and after any compression, respectively.

The *balance* column shows how far off that night's backups are from the average size (shown at the bottom of the balance column). Amanda tries to keep the backups within +/- 5%, but since the amount of data on each filesystem is always changing, and Amanda will never delay backups just to rebalance the schedule, it is common for the schedule to fluctuate by larger percentages. In particular, in the case of a tape or backup failure, a bump will occur the following night, which will not be smoothed out until the next pass through the schedule.

The last line also shows an estimate of how many Amanda runs will be made between full backups for a file system. In the example, a file system will probably have a full backup done every eight times Amanda is run (e.g. every eight days).

```
$ amadmin DailySet1 balance
 due-date  #fs   orig KB    out KB  balance
-------------------------------------------
11/10 Mon   21    930389    768753    +5.1%
11/11 Tue   29   1236272    733211    +0.2%
11/12 Wed   31   1552381    735796    +0.6%
11/13 Thu   23   1368447    684552    -6.4%
11/14 Fri   32   1065603    758155    +3.6%
11/15 Sat   14   1300535    738430    +0.9%
11/16 Sun   31   1362696    740365    +1.2%
11/17 Mon   30   1427936    773397    +5.7%
11/18 Tue   11   1059191    721786    -1.3%
11/19 Wed   19   1108737    661867    -9.5%
-------------------------------------------
TOTAL      241  12412187   7316312   731631  (estimated 8 runs per dumpcycle)
```

FILES

/usr/amanda/etc/amanda/*config*/amanda.conf

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8), amcheck(8), amdump(8), amrestore(8)

NAME

amcheck − run Amanda self−checks

SYNOPSIS

amcheck [**-M***address*] [**-mwsc**] *config*

DESCRIPTION

Amcheck runs a number of self−checks on both the Amanda tape server host and the Amanda backup client hosts.

On the tape server host, **amcheck** goes through the same tape checking used at the start of the nightly **amdump** run. This verifies the correct tape for the next run is in the correct drive. Holding disks are checked to insure they have enough free space.

Amcheck also causes a self−check to be run on all backup client hosts. This makes sure each host is up and running and that permissions on filesystems to be backed up are correct.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

-M*address*
: When **-m** is also set, mail the report to *address* instead of the **mailto** value from *amanda.conf*.

-s Only the tape server host checks are run.

-c Only the backup client host checks are run.

-w Enables a destructive check for write−protection on the tape (which would cause the subsequent **amdump** to fail). If the tape is writable, this check causes all data after the tape label to be erased (actually depends on the device driver: there is no portable non−destructive way to check for write−protection). The check is only made if the tape is otherwise correct.

-m Nothing is printed, but mail is sent if any errors are detected. The mail goes to the **mailto** address as specified in the *amanda.conf* file or the *address* value if **-M** is set.

EXAMPLES

In this example, both the tape server and backup client tests are run. The results are displayed on standard output.

```
% amcheck DailySet1
Amanda Tape Server Host Check
-----------------------------
/amanda2/amanda/work: 911475 KB disk space available, that’s plenty.
NOTE: skipping tape-writable test.
Tape VOL10 label ok.
Server check took 34.966 seconds.

Amanda Backup Client Hosts Check
--------------------------------
WARNING: northstar: selfcheck request timed out.  Host down?
WARNING: drinkme: selfcheck request timed out.  Host down?
WARNING: scruffy: selfcheck request timed out.  Host down?
Client check: 136 hosts checked in 51.945 seconds, 3 problems found.

(brought to you by Amanda 2.4.0)
```

In this example, if the line **mailto csd-amanda** is in *amanda.conf*, mail will be sent to **csd-amanda** if the tape check returns an error.

```
% amcheck −s −m DailySet1
```

MESSAGES

fatal slot *slot*: *error message*

(error) The tape changer detected some kind of fatal error while trying to load slot *slot*.

slot *slot*: *error message*

(warning) The tape changer detected some kind of non-fatal error (e.g. an empty slot was detected) while trying to load slot *slot*, or an error was detected trying to read the tape label.

slot *slot*: date *YYYYMMDD* label *label* (*result*)

(info) Tape *label* in slot *slot* was loaded and found to have been last written on *YYYYMMDD*. If the tape is new, the date field will be an *X*. The *result* may be one of:

exact label match

This is the expected tape.

no match

This label does not match the **labelstr** pattern in *amanda.conf*. Tape scanning will continue.

active tape

This tape is still active and cannot be overwritten. Tape scanning will continue.

first labelstr match

This tape is the first one that matches the **labelstr** pattern in *amanda.conf*. Tape scanning will continue if necessary.

labelstr match

This tape is the next one that matches the **labelstr** pattern in *amanda.conf*. Tape scanning will continue.

ERROR: statfs *disk*: *error message*

(error) An error was returned from the *statfs* system call on holding disk *disk* (maybe because it does not exist).

ERROR: *disk* is unwritable: *error message*

(error) An error was returned from the *access* system call on holding disk *disk*, probably because the caller does not have write permission.

WARNING: avail disk space unknown for *disk*. *N* KB requested.

(warning) **Amcheck** could not determine the amount of available space on holding disk *disk* to see if there were at least *N* KBytes available.

WARNING: *disk*: only *F* KB free (*R* KB requested).

(warning) *Amanda.conf* requested *R* KBytes of free space on holding disk *disk*, but only *F* KBytes were available. 10 MBytes is subtracted for each backup process (see the **inparallel** *amanda.conf* option) to allow for unexpected overruns.

Note: even though this message is listed as a warning, it causes **amcheck** to exit with a non-zero status.

disk: *N* KB disk space available, that's plenty.

(info) There was sufficient free space on holding disk *disk*.

WARNING: if a tape changer is not available, runtapes must be set to 1.

(warning) The **runtapes** *amanda.conf* option must be set to 1 if the **tpchanger** *amanda.conf* option is not set.

ERROR: *error message*.

(error) An error was detected while initializing the tape changer.

ERROR: *tape device*: *error message*.

(error) An error was detected while processing the tape label.

ERROR: cannot overwrite active tape *label*.
(error) Tape *label* is still active and cannot be used.

ERROR: label *label* doesn't match labelstr "*pattern*".
(error) The label on tape *label* does not match the **labelstr** *amanda.conf* option.

(expecting a new tape)
(info) The tape is not OK and a new tape was expected.

(expecting tape *label* or a new tape)
(info) The tape is not OK and either tape *label* or a new tape was expected.

ERROR: tape *label* label ok, but is not writable.
(error) Tape *label* is OK, but the write enable test failed.

Tape *label* is writable.
(info) Tape *label* is OK and the write enable test succeeded.

NOTE: skipping tape-writable test.
(info) The tape write test (see the **-w** option) was not enabled.

Tape *label* label ok.
(info) Tape *label* is OK for the next run.

Index dir "*directory*" doesn't exist or is not writable.
(warning) The directory used to collect backup index files (from the **indexdir** *amanda.conf* option) is not accessible.

Server check took *S* seconds.
(info) Reports how long the tape server host checks took.

ERROR: *host*: couldn't resolve hostname
(error) Could not look up backup client hostname *host*.

Client check: *H* hosts checked in *S* seconds, *N* problems found.
(info) Reports the number of backup client hosts checked, how long it took and the number of errors detected.

WARNING: *host*: selfcheck request timed out. Host down?
(warning) There was no response from *host* when trying to do the client checks.

ERROR: *host* NAK: *message*
(error) *Host* reported a negative acknowledgment error of *message* to the status check request.

ERROR: *host* NAK: [NAK parse failed]
(error) **Amcheck** could not parse the negative acknowledgment error from *host*. There might be an Amanda version mismatch between the host running **amcheck** and *host*.

ERROR: *host* [mutual-authentication failed]
(error) Kerberos authentication failed while contacting *host*.

ERROR: *host*: *message*
(error) Error *message* was reported by the status check on *host*.

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8), amdump(8)

NAME

amcheckdb − check Amanda database for tape consistency

SYNOPSIS

amcheckdb *config*

DESCRIPTION

Amcheckdb verifies every tape mentioned in the Amanda database is still valid in the *tapelist* file.

See the *amanda*(8) man page for more details about Amanda.

EXAMPLE

This shows a normal response:

```
# amcheckdb DailySet1
Ready.
```

This shows tape *DMP014* still listed in the database but no longer listed in the *tapelist* file:

```
# amcheckdb DailySet1
Tape DMP014 missing in /usr/amanda/etc/amanda/DailySet1/tapelist
Ready.
```

AUTHOR

Adrian T. Filipi-Martin <atf3r@cs.virginia.edu>

SEE ALSO

amadmin(8), amrmtape(8), amanda(8)

NAME

amcleanup − run the Amanda cleanup process after a failure

SYNOPSIS

amcleanup *config*

DESCRIPTION

Amcleanup generates the *Amanda Mail Report* and updates the Amanda databases after a system failure on a tape server host. This cleanup process is normally done automatically as part of the **amdump** program, but if **amdump** cannot complete for some reason (usually because of a host crash), **amcleanup** must be run some time later.

See the *amanda*(8) man page for more details about Amanda.

EXAMPLES

This example runs the Amanda cleanup process by hand after a failure.

```
% amcleanup DailySet1
```

Putting the following line in a system boot script (e.g. */etc/rc.local*) runs the Amanda cleanup process as part of the reboot, eliminating the need to run it by hand.

```
/usr/amanda/sbin/amcleanup DailySet1
```

If nothing needs to be done, **amcleanup** exits normally with the message:

```
amcleanup: no unprocessed logfile to clean up.
```

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8), amdump(8)

NAME

amdump − back up all disks in an Amanda configuration

SYNOPSIS

amdump *config*

DESCRIPTION

Amdump switches to the appropriate Amanda configuration directory, e.g. /usr/amanda/etc/amanda/*config*, then attempts to back up every disk specified by the *amanda.conf* file. **Amdump** is normally run by **cron**.

If file /usr/amanda/etc/amanda/*config*/hold exists, **amdump** will wait until it is removed before starting the backups. This allows scheduled backups to be delayed when circumstances warrant, for example, if the tape device is being used for some other purpose. While waiting, **amdump** checks for the hold file every five minutes.

See the *amanda*(8) man page for more details about Amanda.

EXAMPLE

Here is a typical crontab entry. It runs **amdump** every weeknight at 1 a.m. as user **bin**:

```
0 1 * * 1-5 bin /usr/amanda/sbin/amdump DailySet1
```

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8), amcheck(8), amcleanup(8), amrestore(8), amflush(8), cron(8)

NAME

amflush − flush Amanda backup files from holding disk to tape

SYNOPSIS

amflush [**-f**] *config*

DESCRIPTION

Amflush writes all Amanda backups from holding disk to tape, and updates the Amanda info database and tape list accordingly. The backups stay in the holding disk when something is wrong with the tape at the time **amdump** is run. When this happens, the problem must be corrected and **amflush** run by hand.

Amflush will look in the holding disk specified by the *amanda.conf* file in /usr/amanda/etc/amanda/*config* for any non-empty Amanda work directories. It then prompts you to select a directory. The work directories on the holding disk are named by the date at the time **amdump** was run, e.g. *19910215*.

Amflush normally detaches from the tty and runs as a background process. With the **-f** option, **amflush** stays in the foreground. This is useful if **amflush** is run as part of another script that, for example, advances the tape after the flush is completed.

See the *amanda*(8) man page for more details about Amanda.

EXAMPLE

Amflush will search for holding areas associated with the *PUCC* configuration. After you select which holding area to flush, **amflush** writes the data to tape, updates the databases and sends a mail report similar to *amdump*(8).

```
% amflush DailySet1
```

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8), amdump(8)

NAME

amlabel − label an Amanda tape

SYNOPSIS

amlabel [-a] [-f] *config label* [**slot** *slot*]

DESCRIPTION

Amlabel writes an Amanda label on the tape in the device specified by the *amanda.conf* file in /usr/amanda/etc/amanda/*config*. *Label* may be any string that does not contain whitespace and that matches the *amanda.conf* **labelstr** regular expression option. It is up to the system administrator to define a naming convention.

Amlabel will not write the label if the tape contains an active amanda tape or if the label specified is on an active tape. The **-f** (force) flag bypass these verification.

An optional slot may be specified after the tape label. If a tape changer is in use, *amlabel* will label the tape in the specified slot instead of the currently loaded tape.

The **-a** (add) option will add the tape to the configuration's tapelist file.

See the *amanda*(8) man page for more details about Amanda.

EXAMPLE

Write an Amanda label with the string "VOL1" on the tape loaded in the device named in the **tapedev** option in /usr/amanda/etc/amanda/DailySet1/amanda.conf:

```
% amlabel DailySet1 VOL1
```

Label the tape in slot 3 of the currently configured tape changer with the string "VOL2":

```
% amlabel DailySet1 VOL2 slot 3
```

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8)

NAME

amoverview − display file systems processed by Amanda over time

SYNOPSIS

amoverview [**-config** *config*] [**-hostwidth** *width*] [**-diskwidth** *width*] [**-verbose**]

DESCRIPTION

Amoverview displays a chart showing hosts and file systems processed by Amanda along with the backup level performed each day.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

-config *config*
Use configuration *config* instead of DailySet1.

-hostwidth *width*
Set **host** field column width to *width* characters instead of eight.

-diskwidth *width*
Set **disk** field column width to *width* characters instead of 20.

-verbose
Amoverview can take a long while on large systems. This option reports intermediate steps while it is working.

EXAMPLE

This shows the */home* file system on *host2* was backed up at level 3 on the 8th, 9th and 10th of December, had a full backup on the 11th, a level 1 on the 12th and a level 2 on the 13th.

```
# amoverview
                         date 12 12 12 12 12 12
host     disk                 08 09 10 11 12 13

host1    /                     0  1  1  1  1  1
host1    /var                  0  1  1  1  1  1
host2    /                     1  1  1  1  1  0
host2    /home                 3  3  3  0  1  2
host2    /opt                  1  1  1  1  1  1
host2    /var                  1  1  0  1  1  1
```

SEE ALSO

amadmin(8), amanda(8)

NAME

amplot − visualize the behavior of Amanda

SYNOPSIS

amplot [**-c**] [**-e**] [**-g**] [**-l**] [**-p**] [**-t** *T*] *amdump_files*

DESCRIPTION

Amplot reads an **amdump** output file that Amanda generates each run (e.g. *amdump.1*) and translates the information into a picture format that may be used to determine how your installation is doing and if any parameters need to be changed. **Amplot** also prints out **amdump** lines that it either does not understand or knows to be warning or error lines and a summary of the start, end and total time for each backup image.

Amplot is a shell script that executes an **awk** program (*amplot.awk*) to scan the **amdump** output file. It then executes a **gnuplot** program (*amplot.g*) to generate the graph. The **awk** program is written in an enhanced version of awk, such as GNU awk (**gawk** version 2.15 or later) or **nawk**.

During execution, **amplot** generates a few temporary files that **gnuplot** uses. These files are deleted at the end of execution.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

-c Compress *amdump_files* after plotting.

-e Extend the X (time) axis if needed.

-g Direct *gnuplot* output directly to the X11 display (default).

-p Direct postscript output to file *YYYYMMDD***.ps** (opposite of **-g**).

-l Generate landscape oriented output.

-t *T* Set the right edge of the plot to be *T* hours.

The *amdump_files* may be in various compressed formats (**compress**, **gzip**, **pact**, **compact**).

INTERPRETATION

The figure is divided into a number of regions. There are titles on the top that show important statistical information about the configuration and from this execution of **amdump**. In the figure, the X axis is time, with 0 being the moment **amdump** was started. The Y axis is divided into 5 regions:

QUEUES: How many backups have not been started, how many are waiting on space in the holding disk and how many have been transferred successfully to tape.

%BANDWIDTH: Percentage of allowed network bandwidth in use.

HOLDING DISK: The higher line depicts space allocated on the holding disk to backups in progress and completed backups waiting to be written to tape. The lower line depicts the fraction of the holding disk containing completed backups waiting to be written to tape including the file currently being written to tape. The scale is percentage of the holding disk.

TAPE: Tape drive usage.

%DUMPERS: Percentage of active dumpers.

The idle period at the left of the graph is time **amdump** is asking the machines how much data they are going to dump. This process can take a while if hosts are down or it takes them a long time to generate estimates.

AUTHOR

Olafur Gudmundsson ogud@tis.com
Trusted Information Systems
formerly at University of Maryland, College Park

BUGS

Reports lines it does not recognize, mainly error cases but some are legitimate lines the program needs to be taught about.

SEE ALSO

amanda(8), amdump(8), gawk(1), nawk(1), awk(1), gnuplot(1), sh(1), compress(1), gzip(1)

NAME

amrecover − Amanda index database browser

SYNOPSIS

amrecover [[**−C**] *config*] [**−s** *index-server*] [**−t** *tape-server*] [**−d** *tape-device*]

DESCRIPTION

Amrecover browses the database of Amanda index files to determine which tapes contain files to recover. Furthermore, it is able to actually recover files backed up with DUMP or GNUTAR, but not (yet) with XFSDUMP, VXDUMP, VDUMP or SMBCLIENT.

In order to restore files in place, you must invoke **amrecover** from the root of the backed up filesystem, or use **lcd** to move into that directory, otherwise a directory tree that resembles the backed up filesystem will be created in the current directory. See the examples below for details.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

[-C] *config*

Amanda configuration (default: DailySet1).

-s *index-server*

Host that runs the index daemon (default: bldsunos26).

-t *tape-server*

Host that runs the tape server daemon (default: bldsunos26).

-d *tape-device*

Tape device to use on the tape server host (default: /dev/null).

COMMANDS

Amrecover connects to the index server and then presents a command line prompt. Usage is similar to an ftp client. The GNU readline library is used to provide command line history and editing if it was built in to **amrecover**.

The purpose of browsing the database is to build up a *restore list* of files to be extracted from the backup system. The following commands are available:

sethost *hostname*

Specifies which host to look at backup files for (default: the local host).

setdate *YYYY-MM-DD*

Set the date (default: today). File listing commands return information only on files from backups for this day, for the day before with the next lower dump level, and so on, until the most recent level 0 backup on or before the specified date is encountered.

For example, if:

1996-07-01 was a level 0 backup
1996-07-02 through 1996-07-05 were level 1 backups
1996-07-06 through 1997-07-08 were level 2 backups

then if 1997-07-08 is the requested date, files from the following days would be used:

1997-07-08
1997-07-05
1997-07-01

Only the most recent version of a file will be presented.

The following abbreviated date specifications are accepted:

--MM-DD

dates in the current year

---DD dates in the current month of the current year

setdisk *diskname* [*mountpoint*]
: Specifies which disk to consider (default: the disk holding the working directory where **amrecover** is started). It can only be set after the host is set with **sethost**. *Diskname* is the device name specified in the *amanda.conf* or *disklist* configuration file. The disk must be local to the host. If *mountpoint* is not specified, all pathnames will be relative to the (unknown) mount point instead of full pathnames.

history Show the backup history of the current disk. Dates and levels of each backup are displayed.

pwd Print the name of the current backup working directory.

cd *dir* Change the backup working directory to *dir.* If the mount point was specified with **setdisk**, this can be a full pathname or it can be relative to the current backup working directory. If the mount point was not specified, then paths are relative to the mount point if they start with "/", otherwise they are relative to the current backup working directory.

lpwd Print the **amrecover** working directory. When you restore files, this directory will be considered the root of the backed up filesystem.

lcd *path*
: Change the directory that will be considered the root of the restored filesystem to *path.*

ls List the contents of the current backup working directory. See the description of the **setdate** command for how the view of the directory is built up. The backup date is shown for each file.

add *path1* [*path2 ...*]
: Add the specified files or directories to the restore list. Each path may have shell style wildcards.

addx *path1* [*path2 ...*]
: Add the specified files or directories to the restore list. Each path may be a regular expression.

delete *path1* [*path2 ...*]
: Delete the specified files or directories from the restore list. Each path may have shell style wildcards.

deletex *path1* [*path2 ...*]
: Delete the specified files or directories from the restore list. Each path may be a regular expression.

list [*file*]
: Display the contents of the restore list. If a file name is specified, the restore list is written to that file. This can be used to manually extract the files from the Amanda tapes with **amrestore**.

clear Clear the restore list.

quit Close the connection to the index server and exit.

exit Close the connection to the index server and exit.

extract Start the extract sequence (see the examples below). Make sure the local working directory is the root of the backed up filesystem, or another directory that will behave like that. Use **lpwd** to print the local working directory, and **lcd** to change it.

help Display a brief list of these commands.

EXAMPLES

The following shows the recovery of an old *syslog* file.

```
# cd /var/log
# ls -l syslog.7
syslog.7: No such file or directory
# amrecover
AMRECOVER Version 1.1. Contacting server on bldsunos26 ...
220 bldsunos26 AMANDA index server (1.0) ready.
```

```
Setting restore date to today (1997-12-09)
200 Working date set to 1997-12-09.
200 Config set to DailySet1.
200 Dump host set to this-host.some.org.
$CWD '/var/log' is on disk '/var' mounted at '/var'.
200 Disk set to /var.
/var/log
WARNING: not on root of selected filesystem, check man-page!
amrecover> ls
1997-12-09 daemon.log
1997-12-09 syslog
1997-12-08 authlog
1997-12-08 sysidconfig.log
1997-12-08 syslog.0
1997-12-08 syslog.1
1997-12-08 syslog.2
1997-12-08 syslog.3
1997-12-08 syslog.4
1997-12-08 syslog.5
1997-12-08 syslog.6
1997-12-08 syslog.7
amrecover> add syslog.7
Added /log/syslog.7
amrecover> lpwd
/var/log
amrecover> lcd ..
/var
amrecover> extract

Extracting files using tape drive /dev/null on host bldsunos26

The following tapes are needed: DMP014

Restoring files into directory /var
Continue? [Y/n]: y

Load tape DMP014 now
Continue? [Y/n]: y
set owner/mode for '.'? [yn] n
amrecover> quit
200 Good bye.
# ls -l syslog.7
total 26
-rw-r--r--   1 root     other      12678 Oct 14 16:36 syslog.7
```

If you do not want to overwrite existing files, create a subdirectory to run **amrestore** from and then move the restored files afterward.

```
# cd /var
# (umask 077 ; mkdir .restore)
# cd .restore
# amrecover
AMRECOVER Version 1.1. Contacting server on bldsunos26 ...
...
amrecover> cd log
/var/log
```

```
amrecover> ls
...
amrecover> add syslog.7
Added /log/syslog.7
amrecover> lpwd
/var/.restore
amrecover> extract

Extracting files using tape drive /dev/null on host bldsunos26
...
amrecover> quit
200 Good bye.
# mv -i log/syslog.7 ../log/syslog.7-restored
# cd ..
# rm -fr .restore
```

If you need to run **amrestore** by hand instead of letting **amrecover** control it, use the **list** command after browsing to display the needed tapes.

```
# cd /var/log
# amrecover
AMRECOVER Version 1.1. Contacting server on bldsunos26 ...
...
amrecover> ls
...
amrecover> add syslog syslog.6 syslog.7
Added /log/syslog
Added /log/syslog.6
Added /log/syslog.7
amrecover> list
TAPE DMP014 LEVEL 0 DATE 1997-12-08
        /log/syslog.7
        /log/syslog.6
TAPE DMP015 LEVEL 1 DATE 1997-12-09
        /log/syslog
amrecover> quit
```

The **history** command shows each tape that has a backup of the current disk along with the date of the backup, the level, the tape label and the file position on the tape. All active tapes are listed, not just back to the most recent full dump.

Tape file position zero is a label. The first backup is file position one.

```
# cd /var/log
# amrecover
AMRECOVER Version 1.1. Contacting server on bldsunos26 ...
...
amrecover> history
200- Dump history for config "DailySet1" host "this-host.some.org" disk "/var"
201- 1997-12-09 1 DMP015 9
201- 1997-12-08 1 DMP014 11
201- 1997-12-07 0 DMP013 22
201- 1997-12-06 1 DMP012 16
201- 1997-12-05 1 DMP011 9
201- 1997-12-04 0 DMP010 11
201- 1997-12-03 1 DMP009 7
201- 1997-12-02 1 DMP008 7
```

```
201- 1997-12-01 1 DMP007 9
201- 1997-11-30 1 DMP006 6
...
amrecover> quit
```

AUTHOR

Alan M. McIvor <alan@kauri.auck.irl.cri.nz>

SEE ALSO

amanda(8), amrestore(8), readline(3)

NAME

amreport − generate a formatted output of statistics for an amdump run

SYNOPSIS

amreport *config* [**-l** *log file*] [**-f** *output file*] [**-p** *postscript file*]

DESCRIPTION

Amreport generates a summary report of an *amanda*(8) backup run. If no configuration name is specified, amanda.conf is read from the current directory, and log files are rotated, i.e., after the report is generated, the log file *log* will be renamed to *log*.*date*.*version* where *date* is the current date in the form YYYYMMDD and *version* is the next higher version of the logfile with the same date stamp starting with 0.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

config Name of the configuration to process.

-l *log file*

Name of the log file to parse to generate the report. If a log file is not specified, then it defaults to the file:

logdir/log,

logdir is the log directory defined in amanda.conf.

-f *output file*

Normally, **amreport** sends the report via e-mail to the *mailto* user as defined in the amanda.conf file. If *output file* is specified, then the report is put in *output file.*

-p *postscript file*

Send the postscript output to the file *postscript file* instead of to the *lpr*(1) command. This option has an effect only if the *lbl-templ* directive is specified in amanda.conf. See **LABEL PRINTING** below.

LABEL PRINTING

Amanda has the ability to print postscript labels which describe the contents of the tape being written to. The labels are designed to be folded and inserted into the tape case along with the tape or hole punched and put in a 3-ring binder. Various label templates have been provided which format data in ways that are ideal for each tape size. The information provided varies slightly from one label template to the next, due to size constraints. The labels contain one line for each host/file-system pair and can also contain the file number on the tape, the level of the dump, the original size of the dump and the size of the (possibly compressed) file written to the tape. The ability to print these labels is achieved by adding the lbl-templ field to the tapetype definition in amanda.conf. If you don’t add this line to your tapetype definition, then **amreport** will not print a tape label.

You can also specify an alternate "printer" definition to print the label to other than the system default printer.

Labels Provided

The authors have provided label templates for the following tape types. These are pretty generic labels, and should be easy to customize for other tape types. Others are encouraged to do so.

* exabyte 8mm tapes
* dat 4mm tapes
* DLT tapes
* 8.5x11 3-ring binder

The 8.5x11 3-ring binder type is the most generic. It can be used to make a hardcopy log of the tapes which can be kept in a 3-ring binder.

SEE ALSO

amanda(8), amflush(8)

BUGS

The function of rotating logs probably should not be part of amreport since they are 2 different funtions and really have nothing to do with each other.

NAME

amrestore − extract files from an Amanda tape

SYNOPSIS

amrestore [**−r** | **−c**] [**−p**] [**−h**] *tapedevice* [*hostname* [*diskname* [*datestamp* [*hostname* [*diskname* [*datestamp* ...]]]]]]

DESCRIPTION

Amrestore extracts files from the tape mounted on *tapedevice* that match *hostname*, *diskname* and *datestamp* patterns given on the command line. The tape must be in a format written by the **amdump** or **amflush** program.

If *diskname* is not specified, all backups on the tape for the previous *hostname* are candidates. If *datestamp* is not specified, all backups on the tape for the previous *hostname* and *diskname* are candidates. If no *hostname*, *diskname* or *datestamp* are specified, every backup on the tape is a candidate.

Hostname, *diskname* and *datestamp* are regular expressions that may match more than one backup. For example, if *diskname* is "rz[23]a", it would match disks **rz2a** and **rz3a**. *Hostname*, *diskname* or *datestamp* may be an empty string ("") to match every host, disk or date (this is just an abbreviation for ".*").

Datestamp is useful if you have multiple **amflush** runs on the tape.

Unless **-p** is used, candidate backup files are extracted to files in the current directory named:

hostname.diskname.datestamp.dumplevel

OPTIONS

-p Pipe output. The first matching backup file is sent to standard output, which is normally a pipe to **restore** or **tar**, then **amrestore** quits. It may be run again to continue selecting backups to process. Make sure you specify the no-rewind *tapedevice* when doing this.

Note: **restore** may report "short read" errors when reading from a pipe. Most versions of **restore** support a blocking factor option to let you set the read block size, and you should set it to 2. See the example below.

-c Compress output. **Amrestore** normally writes output files in a format understood by **restore** or **tar**, even if the backups on the tape are compressed. With the **-c** option, **amrestore** writes all files in compressed format, even if the backups on the tape are not compressed. Output file names will have a *.Z* extension. This option is useful when the current directory disk is small.

-r Raw output. Backup files are output exactly as they are on the tape, including the **amdump** file headers. Output file names will have a *.RAW* extension. This option is only useful for debugging and other strange circumstances.

-h Header output. The tape header block is output at the beginning of each file. This is like **−r** except **−c** may also be used to compress the result. **Amrecover** uses the header to determine the restore program to use.

EXAMPLES

The following does an interactive restore of disk *rz3g* from host *seine*, to restore particular files. Note the use of the **b** option to **restore**, which causes it to read in units of two 512-byte blocks (1 Kbyte) at a time. This keeps it from complaining about short reads.

```
% amrestore −p /dev/nrmt9 seine rz3g | restore ivbf 2 −
```

The next example extracts all backup files for host *seine*. This is the usual way to extract all data for a host after a disk crash.

```
% amrestore /dev/nrmt9 seine
```

If the backup datestamp in the above example is *19910125* and *seine* has level 0 backups of disks *rz1a* and *rz1g* on the tape, these files will be created in the current directory:

```
seine.rz1a.19910125.0
seine.rz1g.19910125.0
```

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8), amdump(8), amflush(8), restore(8)

NAME

amrmtape − remove a tape from the Amanda database

SYNOPSIS

amrmtape [**-n**] [**-v**] [**-q**] [**-d**] *config label*

DESCRIPTION

Amrmtape invalidates the contents of an existing backup tape in the configuration database. This is meant as a recovery mechanism when a good backup is damaged either by faulty hardware or user error, e.g. the tape is eaten by the drive or is overwritten.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

-n Generate new *tapelist* and database files with *label* removed, but leave them in */tmp* and do not update the original copies.

-v List backups of hosts and disks that are being discarded. Enabled by default.

-q Opposite of **-v .**

-d Run in debugging mode so all executed commands are shown.

EXAMPLE

Remove tape labeled *DailySet1.034* from the *DailySet1* configuration.

```
# amrmtape DailySet1 DailySet1.034
```

AUTHOR

Adrian T. Filipi-Martin <atf3r@cs.virginia.edu>

SEE ALSO

amadmin(8), amanda(8)

NAME

amstatus − give the current state of an amanda run

SYNOPSIS

amstatus [--config] *config* [--file *amdumpfile*] [--summary] [--dumping] [--waitdumping] [--waittaper] [--dumpingtape] [--writingtape] [--finished] [--failed] [--estimate] [--gestimate]

DESCRIPTION

Amstatus gives the current state of the amanda run specified by the *config* configuration.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

All options may be abbreviated to the shortest non-ambiguous sub-string. If no options are given then everything is printed.

[--config] config
: Specify the amanda configuration you want to print the state for.

--file amdumpfile
: Specify an alternate file instead of the *amdump* file.

--summary
: Print a summary of the state of the run.

--dumping
: Print all partitions that are dumping.

--waitdumping|wdumping
: Print all partitions that are waiting to be dumped.

--waittaper|wtaper
: Print all partitions dumped that are waiting to be written to tape.

--dumpingtape|dtape
: Print all partitions that are dumping directly to tape.

--writingtape|wtape
: Print all partitions that are writing to tape.

--finished
: Print all partitions that are dumped and written to tape.

--failed|error
: Print all partitions that failed.

--estimate
: Print all partitions whose estimate is finished. Works only during the estimate phase.

--gestimate|gettingestimate
: Print all partitions whose estimate is not finished. Works only during the estimate phase.

SEE ALSO

amanda(8), amcheck(8), amdump(8), amrestore(8), amadmin(8)

NAME

amtape − user interface to Amanda tape changer controls

SYNOPSIS

amtape *config command* [*command options*]

DESCRIPTION

Amtape performs tape changer control operations. It uses the underlying tape changer script defined by the **tpchanger** option for a particular Amanda configuration as specified by the *config* argument.

Tape changers maintain a notion of the *current* and *next* slot for each configuration. These may or may not correspond to an actual physical state of the device, but do tend to minimize searching through the tape rack. If the desired tape is in the current slot, it is likely the next tape needed is in the next slot rather than at some random position in the rack.

See the *amanda*(8) man page for more details about Amanda.

COMMANDS

reset Reset the tape changer to a known state. It changes the *current* slot to the *first* slot and may have other device-specific side effects. Some gravity stackers need to be reset to the top position by hand. This command notifies Amanda the stacker is back in that position.

eject If a tape is loaded in the drive, it is ejected and returned to the slot from which it was loaded.

show Show the contents of all slots in the rack. This can be slow.

label *label*

Search for and load an Amanda tape with label *label*.

taper Perform the **taper** scan algorithm. Load the next tape in the configuration's tape sequence, or a fresh tape with a suitable label.

device Print the name of the current tape device to *stdout*.

slot *slot*

Eject any tape in the drive and put it away, then load the tape from slot *slot* and reset *current*.

slot current

Eject any tape in the drive and put it away, then load the tape from the current slot.

slot prev

Eject any tape in the drive and put it away, then load the tape from the previous slot and reset *current*.

slot next

Eject any tape in the drive and put it away, then load the tape from the next slot and reset *current*.

slot first

Eject any tape in the drive and put it away, then load the tape from the first slot and reset *current*.

slot last

Eject any tape in the drive and put it away, then load the tape from the last slot and reset *current*.

slot advance

Eject any tape in the drive and put it away. Advance *current* to the next tape, but do not load it.

This is useful with non-gravity stackers to unload the last tape used and set up Amanda for the next run. If you just use **eject**, the current tape will need to be mounted again in the next run, where it will be rejected as being in use, ejected and the next tape requested. Using **slot next** followed by **eject** does an unnecessary mount.

Note: most changers optimize the **slot** command to not eject the loaded tape if it is the one being requested.

AUTHOR

James da Silva <jds@cs.umd.edu>
University of Maryland, College Park

SEE ALSO

amanda(8)

NAME

amtoc − generate tape table of contents for an Amanda run

SYNOPSIS

amtoc [**-c**] *config* [**-l** *label*] [**-h** *host*] [**-v**] [**-i**] [**-t**] [**-f**]

DESCRIPTION

Amtoc generates a tape table of contents (TOC) for an Amanda run. It's a perl script, you need to install perl first.

See the *amanda*(8) man page for more details about Amanda.

OPTIONS

[**-c**] *config*
Name of the configuration to process.

-l *label* Label of the tape (default: last Amanda run).

-h *host* Only process the specified host (default: all hosts).

-v Be verbose. This output is not really a TOC, but a full dump of the database. It may be restricted by the **-l** and **-h** options.

-i Add an informational text header to the beginning of the TOC containing a reminder about how to restore files.

-t Use tab rather than spaces between the output fields.

-f Write output to a file rather than standard output. The file is named:

logdir/*label*.toc,

where *logdir* is the log directory defined in amanda.conf, and *label* is the tape label (any '/' characters are converted to '_').

EXAMPLE

Amtoc is normally run right after **amflush** or **amdump**.

```
# amdump DailySet1
# amtoc DailySet1 −f
```

AUTHOR

Nicolas MAYENCOURT <nicolas@cui.unige.ch>
University of Geneva/Switzerland

SEE ALSO

amanda(8), amflush(8), perl(1)

BUGS

Amtoc uses the Amanda database, which only maintains information about the last tape used at a given backup level. Generating a TOC for the most recent run should be OK, but generating one for an older run is usually only partially correct.

NAME

amverify − check an Amanda tape for errors

SYNOPSIS

amverify *config* [*slot*]

DESCRIPTION

Amverify reads an Amanda format tape and makes sure each backup image can be processed by **amrestore** and, if possible, the appropriate system restore program (e.g. **tar**).

Amverify runs **amrestore** on each file of the tape and pipes the output to the system restore program with an option to create a catalogue of the backup. The catalogue itself is discarded. Only the success or failure of the operation itself is reported.

If the dump image cannot be processed by the system restore program, e.g. if it was written on a different operating system, the image is still processed by **amrestore** but thrown away. This still determines if the tape is readable, but does not do any internal consistency check on the image.

If *config* is set up to use a tape changer, the *slot* argument may be used to choose the first tape to process. Otherwise, the **current** slot is used.

The **runtapes** configuration parameter determines how many tapes are processed.

See the *amanda*(8) man page for more details about Amanda.

AUTHOR

Axel Zinser <fifi@icem.de>

SEE ALSO

amrestore(8), amanda(8)

Amanda at Fermilab

- Intro/Terms
- Dump data flow
- Recover data flow
- Client install
- Server Configuration

Terms

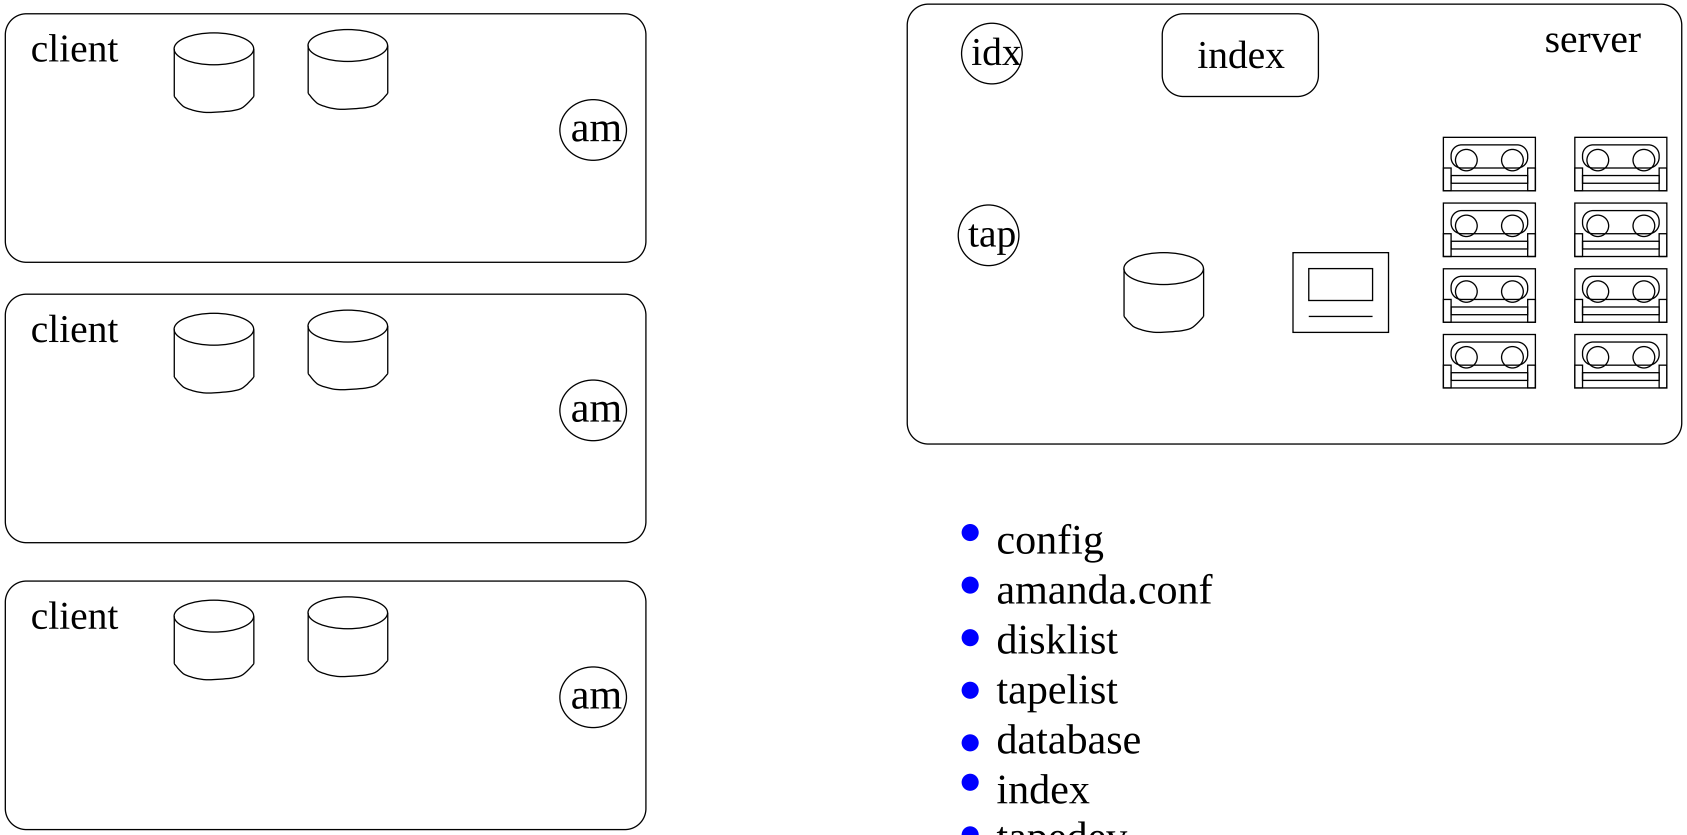

- config
- amanda.conf
- disklist
- tapelist
- database
- index
- tapedev
- holdingdisk

Dump Data Flow

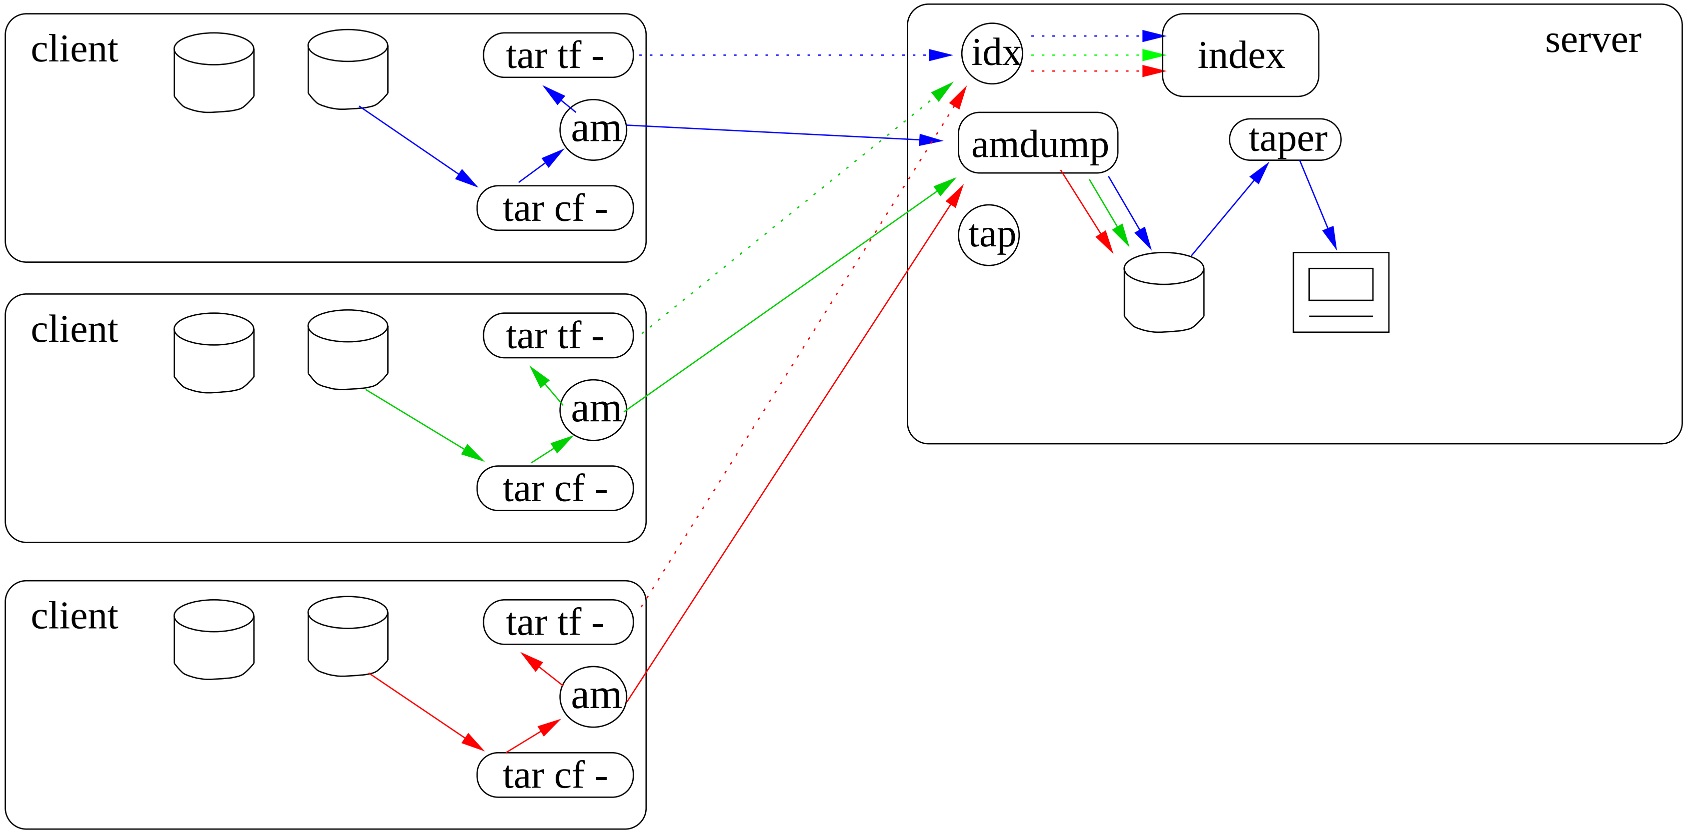

Recover Data Flow

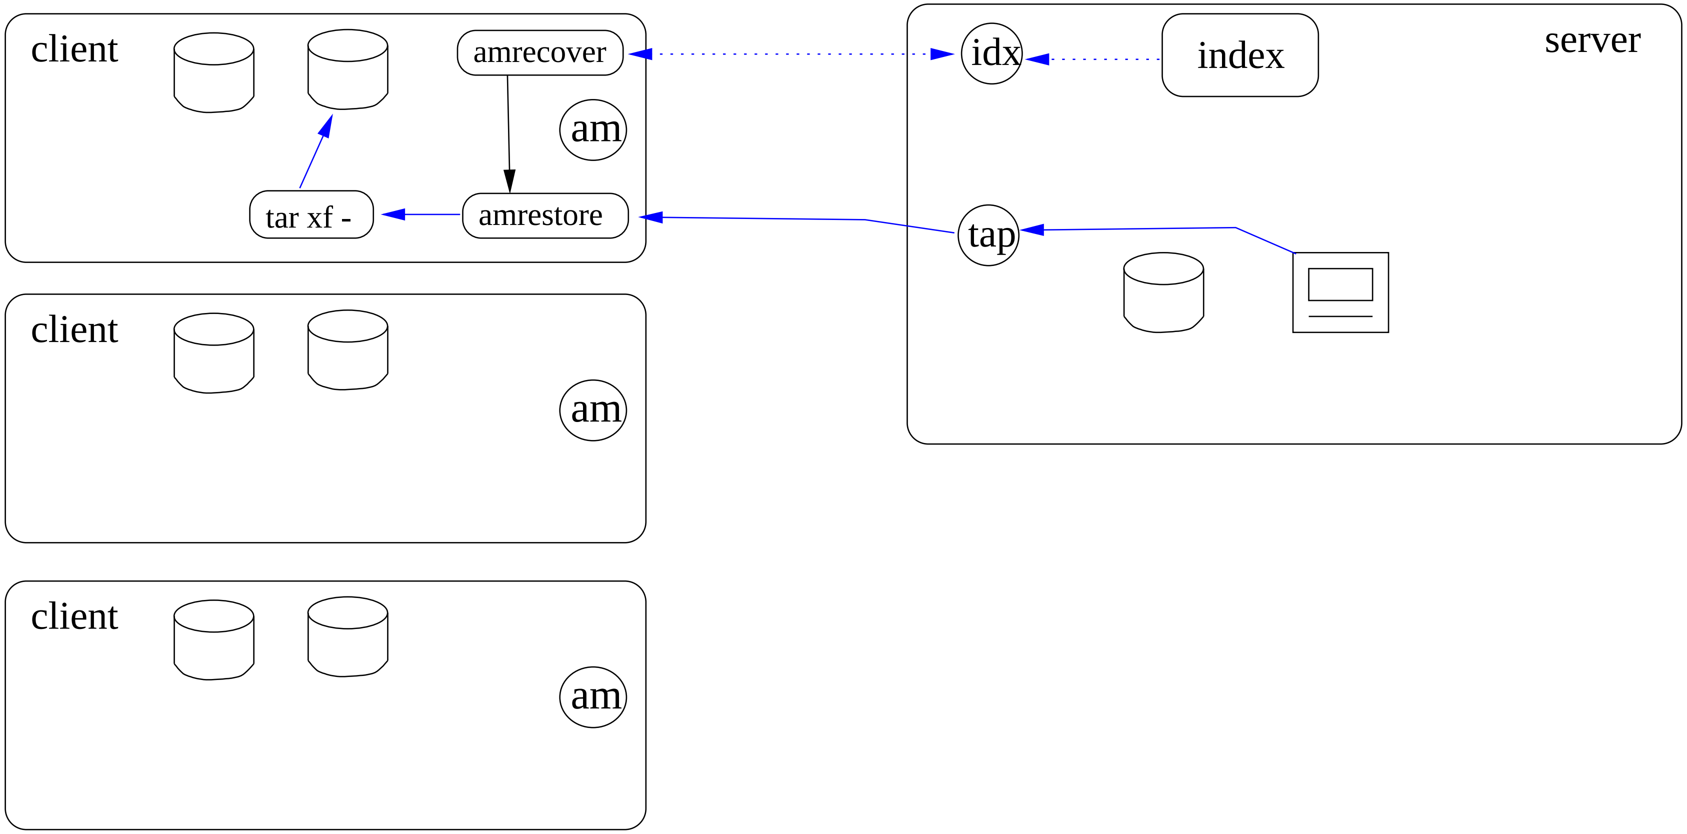

Client Installation

- upd install amanda
- ups InstallAsRoot amanda
- ~daemon/.amandahosts

Server Configuration

- **Client Config first**
- **~daemon/.amandahosts**
- **amanda.conf**
- **disklist**
- **label tapes**