

AMANDA

The Advanced Maryland Automatic
Network Disk Archiver

- A sophisticated network backup system that can replace scripts used by many sites.
- Can backup all machines to a single server's tape drive.
- Free
- Can run on a variety of UNIX and Linux variants.

- Originally written by James da Silva /U of Maryland.
- Now , supported by a team from around the world.
- www.amanda.org

- Not itself a backup program, but manages other backup software.
- It calls standard backup software, such as dump & restore GNU tar and smbtar (Windows-based hosts).
- Supports a wide variety of tape drives.
- Can compress backup images before the data goes over the N.W

Tape management

- Writes a header on each tape.
- Manages dump levels based on configuration parameters and the fullness of tapes.
- Keeps records of which backups are on which tapes.
- Can print labels showing the contents of each tape.

The Architecture of AMANDA

- Tape drives and holding disks are attached to a central server.
- Config files, logs, databases resides also on the server.
- It Can write one backup image to each tape at a time, but can spool multiple dumps to its holding disks simultaneously and then stream them out to tape.

- Supports multiple configurations with each has its own log and databases.
- AMANDA checks its config files to determine which filesystems need to be backed up, which tape drives are available and how much of the system resources it is allowed to use.
- It asks clients to estimate the size of their backup files.
- Then it schedules backup.

The server

- Its server is a collection of programs.
- Best run on a not-busy fast machine.
- Its holding disk should be at least as large as the largest partition you want to backup.
- It needs some disk space for logs and databases(<75 MB).

Setup

- After unpacking the source code, read the README, docs/system.notes and docs/install files.
- Each device you intend to backup must be readable by Amanda's group, and the /etc/dumpdates must be writeable by this group.
- Run configure with the appropriate parameters.
- Then make and make install to the complete installation.

- Every client needs access to binaries, so install them on each client's local disks.
- **amandad**: handles communication between client & server; runs all other client programs.
- **Selfcheck**: checks the client is setup for amanda.
- **Sendbackup**: performs the backups.
- **Sendsize**: estimates backup sizes at different dump levels.

- **amdump**: does the nightly dumps; usually run by cron.
- **amflush**: flushes the holding disk to tape.
- **amrestore**: handles restore operations from dumps.
- **amlabel**: writes labels on tape.
- **amadmin**: finds the right tape to restore from
- **amcheck**: verifies that you are using the correct tape, enough space,...

- Each configuration is kept in a separate directory and needs the **amanda.conf** and **disklist.amanada.conf**: specifies the general server configurations.
- Disklist: specifies which clients and filesystems to backup.

amanda.conf

- Can be viewed as:
 - Local info
 - Dump strategy
 - Resource parameters
 - Dump type definitions

Amanda.conf example

```
org "Podunk Univ"           //organization name
mailto amanda               // operators
dumpuser "amanada"        //user to run dumps
runtape 1                   // number of tapes
tpchanger "chg-manual"    //tape changer glue script
tapedev "/dev/nst0"       //no rewind
labelStr "^podunk-[0-9][0-9]*$" //label
infofile "usr/bin/anamada/curinfo"
logdir "/usr/adm/amanda/podunk"
```

- Amanda won't use a tape unless it matches the labelstr, therefore tapes must be labeled by **amlabel**.
- The system admin defines the labeling scheme.
- The server rejects any use of a tape before the normal rotation cycle has been completed.(defined in the strategy parameters)

Strategy parameters

dumpcycle 4 weeks // # of days in the dump cycle

bumpdays 2 //minimum days at each level

bubpsize 20Mb //minimum savings to bump level 1->2

runspercycle 20 //# of amdump runs in dumpcycle days

tapecycle 25 //# of tapes in a rotation

- Dumps are not scheduled by calendar date. Instead amanda tries to spread the work out across the entire dump cycle so that tapes are used efficiently and protection is maximized.
- A dump cycle is a period in which level 0 dumps will be done at least once on every filesystem.

- Long cycles increase amanda's scheduling flexibility, but may increase the avg number of tapes that must be read to complete a restore.
- The default value to run dumps is once a day.(otherwise specify **runspercycle**)
- One dump tape will be written per run(default).

Resource parameters

- Specify how much network bandwidth, CPU, disk space to use.
- The type of tape drive to backup to. (*The tape type parameters should never be guessed*)
- Types of client partitions to be backed up.

example

tapetype EXB-8500

inparallel 4

netusage 600 Kbps

etimeout 300

holdingdisk hd1{

directory “/dumps/amanda”

use 8196 Mb}

Define tapetype EXB-8500{

length 4200 filemark 48 Kbspeed 474Kb}

Dump types

```
define dumptype comp-user{  
    compress client fast    priority medium}
```

```
define dumptype comp-root{  
    compress client fast    priority low}
```

```
define dumptype nocomp-user{  
    compress none    priority medium}
```

```
define dumptype clone-user{  
    compress client fast skip-incr    priority medium}
```

- These dump types are all predefined, but you can define your own

The Disklist file

- The `amanda.conf` tells how to do dumps without specifying any clients or filesystems to dump. That info is recorded in the disklist file.

```
// client      partition      dumptype mountpoint  
ocean         sdoa           comp-root  #/  
ocean         sdog          comp-user  #/usr
```

- If you keep amanda logs or other important info on the holdingdisk, you should back it up.

Log files

- It creates 2 log files on the server for each run.
 - *Amdump.n*: contains a description of the scheduling decisions that amanda has made.
 - *Log.date.n*

debugging

- A summary report on the run's activities is mailed to the dumpmeister.
 - Amount of tape used
 - The filesystems that were successfully backed up
 - Any errors encountered

If amanda was unable to write backups onto a tape, it spools them to the holding disk.(use amflush to write them to tape)

restoration

- Use amadmin and amrestore programs.
 1. Find the tapes on which the directory was backedup.(machine, partion, full path,date lost, date last modified).
 2. Run mt fsf to fast-forward the tape(optional)
 3. Use amrestore to actually restore the data.
- Restore must run on the same O.S that generated the original dump.

Other backup products

- Free:-
 - Arkeia, BURT, cddump, star.
- Commercial:-
 - ADSM/TSM, Veritas, Legato.