

# ***AMANDA @ TRIUMF***

*Steven McDonald & Konstantin Olchanski  
TRIUMF Network & Computing Services  
[steven.mcdonald@triumf.ca](mailto:steven.mcdonald@triumf.ca)*

# *Perspective*

- Amanda is not in the same class of TSM
  - Make no mistake this is \*not\* a TRIUMF T1 storage solution
- Not something for large experimental data sets
  - Requiring frequent access
- More suitable for group, department sys admins for backing-up user desktops and servers,
  - Infrequent data recovery/access
  - Disaster recovery
  - what it was designed for
- Low cost, easy setup, low maintenance backup of desktops and servers

# *What's on the Menu*

- What type of backups?
- Hardware
- Software configuration
- Our environment
- Typical nightly usage
- Performance
- Advantages & disadvantages
- Conclusions

# *Background Requirements*

- What are we attempting to do?
  - Archive data?
  - Disaster recovery?
  - Careless file deletion and recovery?
  
- Would like something to do all the above with
  - Low cost, economically
  - Minimal man-power resources, monitoring, maintenance
  
- **Amanda** can do this for **us** with some compromise
  
- Obviously this is NOT a T1 Backup solution

# Amanda

**Advanced Maryland Automatic Network Disk Archiver**  
opensource from *<http://www.amanda.org>*

- Amadmin
- Amcheck
- Amcleanup
- Amadd
- Amdump
- Amflush
- Amgetconf
- Amlabel
- Ammnt
- Amverify
- Amoverview
- Amplot
- Amrecover
- Amreport
- Amrestore
- Amrrmtape
- Amstatus
- Amtape
- Amtapetype
- Amtoc

# *Supported Operating Systems*

- Linux (~150)
- Solaris (~6)
- Ultrix
- Digital Unix
- SunOS 4
- HPUX
- Sinix
- IRIX
- SCO
- FreeBSD
- AIX
- Mac OS X
- Microsoft Windows ?

- Amanda system (~\$10k US)
  - SL4.2
  - Dual Opteron 248 2.2 GHz
    - 3U/Dual power
    - 2G Memory
    - Data Storage 16 x400G WD disks ~ 6TB
    - System disks 2x 80G

- Overland Neo 2000 tape library
  - 26 slot capacity

- SDLT 600 drive
  - SDLT2 300GB native per tape

- ~7.5TB native DLT2 tape capacity

- Library & DLT600 drive ~\$12k US



# *Disk I/O Configuration*

- System disks 2x 80GB software RAID1 mirror
- Backup Storage 16x 400GB software RAID5
  - ~6TB useable space
- Disk *I/O* performance from Bonnie<sup>++</sup>
  - 200 MB/sec block writes
  - 280 MB/sec block reads



# *What is backed up*


- ~ 150 machines
  - Mostly linux RH7.3 – SL4.2
  - Some Solaris
- Servers
  - Everything including system and home directories
- Client machines
  - /etc ( ~20MB )
  - /home (no enforced limits, but not data files)
    - 10-60GB is typical

# *Amanda Client & Server Configuration*

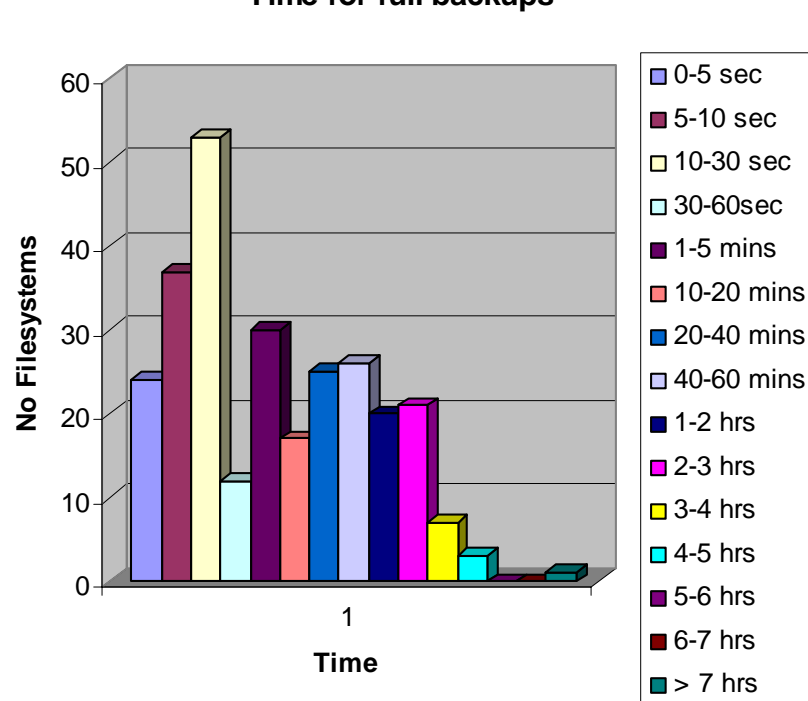
- **On the client:**
  - edit `~amanda/.amandahosts`, add line `"amanda.triumf.ca amanda"`
  - edit `/etc/hosts.allow`, add line `"amandad: amanda.triumf.ca"`
  - `chkconfig amanda on`
  
- **On the server:**
  - edit `~amanda/daily/disklist`, add entry `"client /home1 comp-usr  
#clientusername@triumf.ca"`
  - run `"amcheck -c daily client"`

# *Some statistical distributions*

**Distribution of filesystem sizes**




**Time for full backups**



# How do we use Amanda

**Amanda Clients ~150 in total (2.2TB)**

100Mbit clients  
Typical throughput is  
3-4 MB/sec


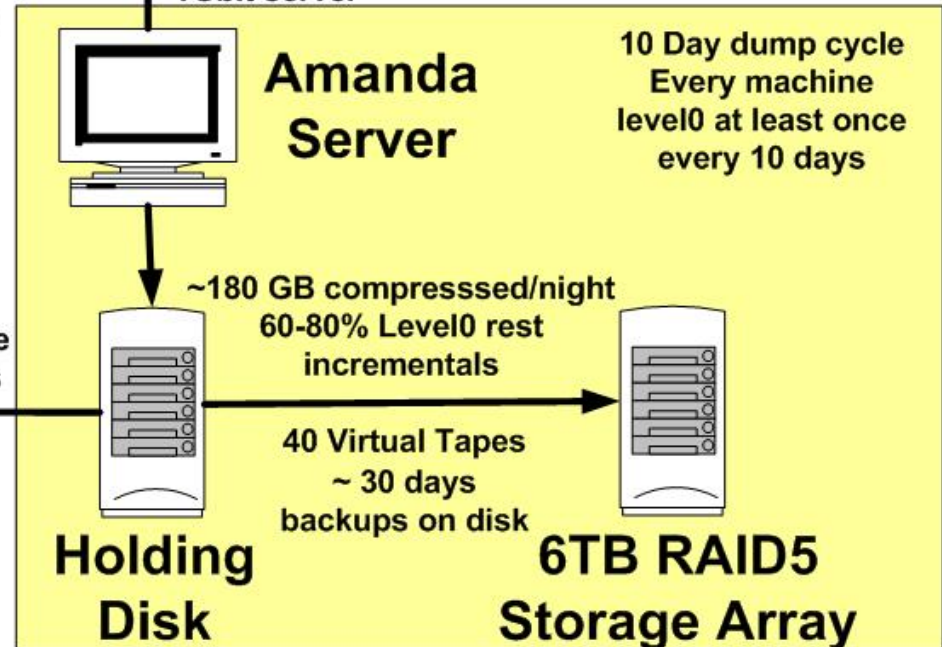


Configured for 10  
Parallel dumpers/  
streams

Compression on  
the clients  
Typically 40-50%


Nightly Dump ~6 hours  
Every clients gets at least an incremental backup  
~20-30 clients a Full backup

1Gbit server




# Amanda Network Usage

Daily



Weekly



# Restoring Data (1)

## Instructions for data recovery (whole disk)

```
ssh root@amanda
mkdir /home1/amanda_recover
chown amanda.root /home1/amanda_recover/
su - amanda
cd /home1/amanda_recover
```

---> we will recover neut04:/etc

```
amadmin daily info neut04 /etc
```

---> prints:

Current info for neut04 /etc:

```
Stats: dump rates (kps), Full: 475.0, -1.0, -1.0
      Incremental: 2.0, -1.0, -1.0
compressed size, Full: 20.8%, -100.0%, -100.0%
      Incremental: 6.3%, 6.4%, 5.2%
```

Dumps:	lev	datestamp	tape	file	origK	compK	secs
	0	20041123	daily0001	5	18320	3802	8
	1	20041130	daily0007	10	650	41	0

---> note the backup levels, the tape numbers and file numbers:

---> level 0: tape daily0001, slot 1, file 5

---> level 1: tape daily0007, slot 7, file 10

ALTERNATIVE1 (extract tarballs on amanda, copy them to client, untar)

```
amrestore -f5 file:/home1/amanda/daily/vtapes/1 neut04 /etc <--- level0 backup
amrestore -f10 file:/home1/amanda/daily/vtapes/7 neut04 /etc <--- level1 backup
---> creates files neut04._etc.20041123.0, neut04._etc.20041130.1
```

```
scp neut04* neut04:
```

```
ssh root@neut04
```

```
mkdir etc
```

```
cd etc
```

```
tar xvBf - ~/neut04._etc.*.0
```

```
tar xvBf - ~/neut04._etc.*.1
```

## Instructions for data recovery (individual files)

```
ssh root@amanda
run amrecover
amrecover> sethost tw04
amrecover> listdisk <----- get list of available disks
amrecover> setdisk /etc
amrecover> ls <---- look at the files
amrecover> add hosts <---- add "hosts" to recovery list
amrecover> extract
...
The following tapes are needed: daily0020 <----- note tape number "20"
...
Restoring files into directory /root
Continue [?/Y/n]? <----- YES
Load tape daily0020 now
Continue [?/Y/n/t]? t <----- answer "t"
New tape device [?]: amanda.triumf.ca:file:/home1/amanda/daily/vtapes/20
./hosts
amrecover>
```

# Amanda

- Advantages

- Low cost (free)
- Pre-installed on RH,SL,FC
- Not resource intensive
- Highly configurable
- Full featured
- Runs on wide range of hardware
- Low maintenance
- CLI highly informative
- Easy restore operation

- Disadvantages

- No nice Web/GUI Administrative Interface
- Had to write scripts to visualize summary data
- No long term data trending
- Difficult to archive data unless you use lot\$ of tapes



# Conclusion

- Overall Amanda provides a very reliable and configurable backup mechanism for our particular environment at TRIUMF
- At a low cost both in
  - Price
  - Man power upkeep and maintenance
- Lots of room for growth.
- More work on importing data into a Database would be nice for long term trending and statistical analysis ... on my todo ...