Storage of large research data volumes in AFS

(on a very low budget)

Richard Brittain

Dartmouth College

Introduction

- The northstar.dartmouth.edu cell
- Expansion plans
- Problems: financial, historical, cultural
- Some case studies
- Miscellaneous tools

Dartmouth College

- Research Computing support group
- A small cell (by some standards)
- Cell name is a legacy of Project Northstar
- Client mix has changed greatly over time

Cell Statistics

- 3 file servers, 3 DB servers
- 10 TB, 3TB in use (+ 20TB in the mail)
- 2434 volumes
- 701 user homes
- 50 ‘data’ volumes (50-800 GB)
- < 100 clients
- 2 physically separated data centers

Hardware

- IBM x3650, EXP3000 disk vaults, LSI/ServeRAID controllers.
- 750 GB SATA disks
- 1Gbps between servers and central systems
- 100Mbps to all departments and desktops
- RHEL 5 everywhere; ext3 filesystems
- vicepX are 3.4TB (1/2 vault)

Backups

- vos dump | gzip > local disk on each server
- NetBackup picks up compressed images
- 2TB staging space; 2:1 compression
- Large data volumes get replicated instead
- Monthly full / daily incremental
- Work in progress to spread fulls

Growing pains

- Plans for significant expansion, but no committed funding or increased manpower
- Biologists have several TB, will get 100-200 TB in the next couple of years
- Chargeback model needed
- Legacy issues with AFS cell

Flirtation with CIFS

- Engineering School likes CIFS
- Test SAMBA server
- Research SAMBA/AFS integration
- Authentication requirements
- Cost :-(
- Plan B - back to AFS
  - but some of the buzzwords seem to fit again

The Opposition

- USB drives $100/TB
- Buffalo Terastation $150/TB
- Fun with rsync
- Resistance to lots of servers
- NetBackup limits: use AFS replication as backup
  - Explored shadow volumes

Case Study

- Biology: gene sequencers
- 1 TB per “run”
- Typical files 2 MB TIFF
- May be able to compress 4:1 or more
- Store 3 years minimum

Case Study

- Medical School long term study
  - 8 years of aspirin data in SAS datasets
  - Several rounds of hardware and software upgrades
  - Many researchers came and went. ACLs are a mess
  - Data are now frozen

Case Study

- Proteomics research
  - Data acquired on unattended PC off campus
  - Written to AFS with IP ACL
  - Visible to Beowulf Cluster head end
  - High volume; no backups

Case Study

- Auroral Radio Noise research in the Arctic
  - Multiple field sites, but Greenland are the only ones on the internet 24x7
  - High latency; behind NAT; AFS not happy
  - scp daily summaries directly to Dartmouth, into AFS space, visible to web server
  - Researcher happy


SONDRESTROM INCOHERENT
SCATTER RADAR FACILITY
DEDICATED TO SCIENTIFIC STUDIES OF THE
POLAR IONOSPHERE AND UPPER ATMOSPHERE
SRI INTERNATIONAL
MENLO PARK, CALIFORNIA
WITH
THE DANISH
METEOROLOGICAL
INSTITUTE
UNDER
COOPERATIVE
AGREEMENT
WITH
THE NATIONAL
SCIENCE FOUNDATION
WASHINGTON, D.C.
TOURS BY APPOINTMENT ONLY TEL.
INGEN ADGANG FOR UVEDKOMMENDE
AFTALE OM RUNDVISNING KAN TRÆFFES PÅ TELEFON

House Call…

Case Study

- Biology: scanner images
- 650GB stored on Terastation
- Pulled with rsync for several months
- Now use AFS as primary storage
- Replicated volume

Miscellaneous user tools

- afsquota

```
Volume Name       Quota       Used % Used Part.  Available
user.richard      11 GB      10 GB    92% 278%   869 MB
```

- freespace

```
mizar     /vicepa: 1568 GB free out of 3416(54.1% used)
centaurus /vicepa: 2581 GB free out of 3416(24.5% used)
oort      /vicepa: 2806 GB free out of 3416(17.9% used)
```

- listvols

```
users.b.readonly            536979170 RO       2 kB On-line
rc.mizar.a                  536975438 RW       3 kB On-line
…
datad.jhamilton             536956022 RW     115 GB On-line
rep.wibble                  536975374 RW     128 GB On-line
rep.mcpeek                  536967732 RW     603 GB On-line

Total volumes for server mizar:[a] onLine 272; offLine 0; busy 0
```

Miscellaneous tools cont.

- setacl

```
setacl -Rv system:authuser,read publicstuff
```

- moveafsvol

```
moveafsvol dest-server dest-partition [volume-name ...]
```

- klog_wrapper

```
polaris [12:58pm] ~ $klog rbadmin
Running interactive shell with command logging
Enter AFS (rbadmin) Password:
bash-3.2$
bash-3.2$ exit
```

- autoconfigure: `upserver`, `upclient`, and `make`
  (really need to learn how to use puppet)