



DSS

Data & Storage Services

CERN
IT
Department

Backup Infrastructure at CERN

Alex Iribarren
CERN IT/DSS

24 April 2012 – HEPiX Spring 2012 - Prague



- Description of the service
- Current infrastructure
- Issues
- The AFS Model
- The Implementation
- New Infrastructure




DSS Backup Service

- The Backup Service provides backup, archiving and restore capabilities for servers containing critical data in the computer centre.
- We back up:
 - Network filesystems (67'000 AFS, 1'500 DFS volumes)
 - Email (18'000 mailboxes)
 - Web sites (12'000 websites)
 - Databases (120 DB servers)
 - Servers (1'000 Linux and Windows servers)
 - Virtual Machines (120 hypervisors)
- We don't back up:
 - Physics data (using CASTOR for this)
 - User PCs (already backing up home AFS/DFS directories)



- 5.4 Petabytes on tape
 - 38 % growth over the last year
- 1245 clients
- 1'600 million files
- ~57 TB daily traffic



20 production servers
- 2 library managers

- TSM 5.5 & TSM 6.3
- RHEL 4 & RHEL 5



190 TB of
disk storage



- 2 IBM TS3500 libraries
- 48 IBM drives
- 6'000 IBM 3592 cartridges






DSS Issues

- Complex architecture:
 - At least 3 disk arrays per server
 - Disk array shared by two servers
- Fault tolerant architecture not so great
 - Creates more problems than it solves
- Expensive!
- Difficult to maintain
 - Equipment from different procurement cycles, different warranties, etc.
 - Hard to keep track of who uses what
- “Special” hardware just for us



- Storage Units:
 - Two head nodes, two SAS expander disk arrays
 - Two TSM servers
 - Each TSM server can failover to the other head node
- 2-disk Software RAID1 over both disk arrays





DSS The Hardware

- Head node:
 - Supermicro CSE-825TQ-R740LPB, 2U
 - 1x Intel E3-1260L, 4 cores, 2.4GHz
 - 32 GB DDR3 RAM
 - 2x WD 250 GB SATA drives
 - Dual-port LSI 9200-8e SAS HBA
 - Dual-port Qlogic QLE-2562 8Gb/s Fibre Channel HBA
 - Dual-port 10Gb Ethernet
- Disk array:
 - Promise J830sD, 24U
 - 4x 300GB, 15 kRPM SAS drives
 - 20x 2TB, 7.2 kRPM SATA drives





- RHEL 5.8 and TSM 6.3.1.0
- Two TSM servers, one per head node
 - Running as different users
- Two IP services
- No fancy high availability framework
 - Too complex for our needs
 - Manual failover
 - Homemade startup script + sync mechanism




DSS Benefits

- Pros:
 - Lower cost per GB (~ 2/3 previous price)
 - Simpler maintenance
 - Faster procurement
 - Same hardware as other services in DSS
 - AFS (and maybe EOS & CASTOR soon)
- Cons:
 - Less flexible
 - Untested waters
 - Operational procedures need to be clarified



DSS Good results


- Failover time: ~10 minutes
 - Good enough
- Good performance:
 - One TSM storage volume per RAID volume
 - Limits the number of concurrent backups ...
 - ... but it also eliminates disk contention
 - Potential for dual-path access to tape
 - Tested at 640 MB/s





DSS New Infrastructure

- Reduced amount of hardware
- Cheaper
- Easier management and maintenance
- Improved failure recovery
- Better performance
- Faster procurement cycle
 - Aligned with other services





DSS

Thank you, Questions?

Alex.Iribarren@cern.ch

<http://www.cern.ch/backup>