

Operating in a SAN Environment

March 19, 2002

Chuck Kinne

AT&T Labs

Technology Consultant



Agenda

- Introduction
- What did we do?
- What was the configuration?
- How is it all wired up?
- XP Volume Mapping?
- What was the initial deployment like?
- Did we improve backups?
- How do our upgrades work now?
- Yeah, but can I pay for it?
- What really happened?
- The Future

Introduction

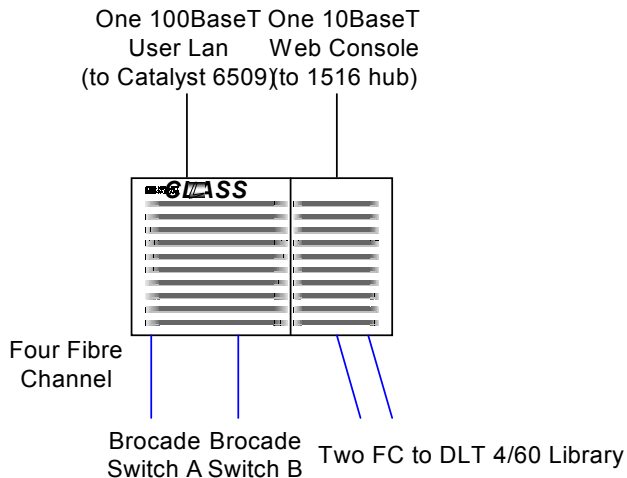
- We needed to replace an aging technology (E, G, H, K Class servers)
400+ HP, 120 Sun, 75 NT
 - **Servers all standalone, internal disk**
 - **Failure rates increasing on internal parts**
 - **Maintenance costs rising**
 - **Needed to support Oracle, Sybase, flat files**
- Needed to improve backup capabilities
 - **Each server was connected to a separate unsecured LAN**
 - **Backup speed limited to speed of the network (100MB)**
- Needed to speed up upgrade process
 - **Used swing servers to migrate to new releases due to uptime requirement**
 - **Very people intensive, manual**
- Needed faster platform for new technology
 - **Products required more RAM, faster processors and disk**

What did we do?

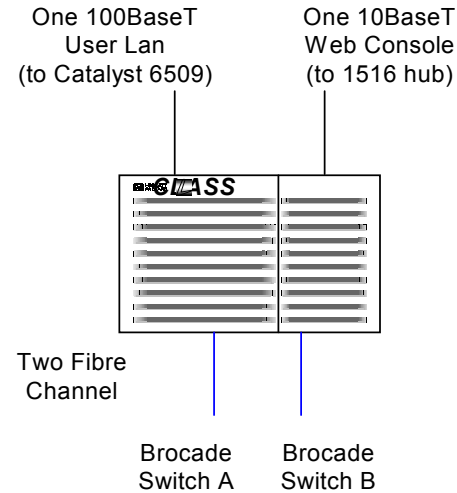
- Replaced the infrastructure with newer HP, SUN, and NT systems
- SAN - Fiber Channel, Brocade switches
- XP disk technology
- No internal disk to the server
- HP Ignite
- Limited Network Connectivity
- POD(SAN) of 80-100 Servers

What was our configuration?

Backup Server Connectivity (L2000)



L-class Server Connectivity



POD Hardware Inventory

- L-class servers
- A500 servers
- two DLT 4/60 tape libraries
- 16 Brocade switches
- one XP256 with two disk frames (space reserved for one disk frame for future expansion)
- two Smart Enclosures with DVD & DAT for A500
- four Cisco 1516 hubs
- one Catalyst 6509 switch


1516 Hubs

- one 10BaseT port for each server(51); used for web console
- eight 10BaseT port for daisy chaining the hubs
- fifty nine 10BaseT ports required; (64) 10BaseT ports total


Catalyst 6509 Switch

- one 100BaseT port for each L-class server (11) and two for each A500 (80)
- one to external LAN

A-class Server Connectivity



How is it all wired up?



Switch #	Switch Port	Server # or XP
FSW_1	0	Rack 1, server 1 PROD
FSW_1	1	Rack 1, server 2 PROD
FSW_1	2	Rack 1, server 3 PROD
FSW_1	3	Rack 1, server 4 PROD
FSW_1	4	Rack 1, server 5 PROD
FSW_1	5	Rack 1, server 6 PROD
FSW_1	6	Rack 1, server 7 PROD
FSW_1	7	Rack 9, server 8 STBY
FSW_1	8	Rack 9, server 9 STBY
FSW_1	9	Rack 9, server 10 STBY
FSW_1	10	Rack 9, server 11 STBY
FSW_1	11	Rack 9, server 12 STBY
FSW_1	12	Rack 9, server 13 STBY
FSW_1	13	Rack 9, server 13 STBY
FSW_1	14	
FSW_1	15	to XP port 1A
Switch #	Switch Port	Server # or XP
FSW_2	0	Rack 9, server 8 STBY
FSW_2	1	Rack 9, server 9 STBY
FSW_2	2	Rack 9, server 9 STBY
FSW_2	3	Rack 9, server 10 STBY
FSW_2	4	Rack 9, server 11 STBY
FSW_2	5	Rack 9, server 12 STBY
FSW_2	6	Rack 9, server 13 STBY
FSW_2	7	Rack 1, server 1 STBY
FSW_2	8	Rack 1, server 2 STBY
FSW_2	9	Rack 1, server 3 STBY
FSW_2	10	Rack 1, server 4 STBY
FSW_2	11	Rack 1, server 5 STBY
FSW_2	12	Rack 1, server 6 STBY
FSW_2	13	Rack 1, server 7 STBY
FSW_2	14	
FSW_2	15	to XP port 1A

XP Volume Mapping

- Volumes dedicated for backup and upgrades.
- In addition to normal RAID capability
- Spread sheet for assigning volumes
- Multiple logical volumes on one physical volume
- Open 9*1 and Open 9*2 (7.2 GB and 14.4 GB effective)

What has the initial deployment been like?

- Hardware has been very stable, very few ELF (early life failures) problems.
 - **In 18 months 21 service calls, 12 on A500s, 4 on the L Classes, 5 on the XP256**
 - **Majority in the first 4 months**
 - **This on (3) XP systems, 48 Brocades, 163 A500s, 63 L Class, 22 N Class**
- Mircode / Software levels required:
 - **HP Server**
 - **HP 11.x (March 2000)**
 - **Fiber Channel PHKL-207 or newer for the Taclite card**
 - **HP Ignite**
 - **Brocade**
 - **XP256/512**
- Fan out rate: 7 to 1
 - FAN out – number of servers using(sharing) a single interface into the array
- Configuration management using EXCEL spread sheets

Did we improve backups?

- **Decreased length of backups (less time)**
 - 57% decrease in the time it took
 - Nightly backups cut in half
 - Fulls once a week, partials every night
- **Increased the amount of data to be backed up (more data)**
 - 2 times the amount of data (1.1 TB to 2.3 TB)
- **Automated scripting with OMNIBack 3.5/4.0 (less people)**
 - Scripts halt DBs and Applications
- **Increased the retention (keep it all longer)**
- **Business Copy feature rather than network backups**

How do our upgrades work?

- **Application or OS changes:**
 - **Business Copy taken**
 - > **Associate Business Copy Volume to system being upgraded**
 - > **Copy made**
 - **Associate Copy to Admin server**
 - > **Apply OS upgrade, Application upgrade, etc**
 - **Re-associate to the original server**
 - **Reverse Business Copy**
 - **No movement of servers or LAN connections**
 - **One person can do it all!**
 - **If not ok , reboot back to the original volume**

How do our upgrades work?

- **XP Changes:**

- Hardware redundant and hot swappable
- Microcode upgraded on the fly

How do our upgrades work?

- **Brocade Changes:**

- Hardware redundant
- Each can be taken off line and reloaded
- Dual pathing to and from each server in the POD.
- 30 second delay as the OS sees the lost link and reacquires through the alternate path

Yeah, but can I pay for it?

- **Things to consider: (see spreadsheet)**
 - **Cost of the hardware**
 - > **Cost / GB with internal vs Cost / GB with external**
 - > **Cost of old backup system vs new, can you reuse the tape subsystems**
 - **Technician's time**
 - > **Can they do more work per person in the new environment?**
 - > **Can they manage more servers in the POD rather than individual standalone servers**
 - **Upgrade costs**
 - > **Can you upgrade more servers with less technicians in less time in the new environment?**
 - > **Cost avoidances by giving a single person more to do**
 - **Technology costs**
 - > **Can you even run the software on the old platform?**
 - > **Support being dropped from the vendor?**

What really happened? (the good, the bad and the ugly)

- **Hardware very reliable**
- **Less footprint**
- **Load balancing with Autopath (or Powerpath from EMC)**
- **Fail over capability very important**
- **Upgrading less cumbersome and much faster**
- **Cable management**
- **Heat dispensation**
- **Software support for configuration management**

The Future

- Tested other vendors
- Integration of products (within vendors)
- Integration of vendors (between vendors)
- Network Concentration (LAN port bottlenecks)
- OVSAM
- Latest ECC