IT802: Xsan2 Planning

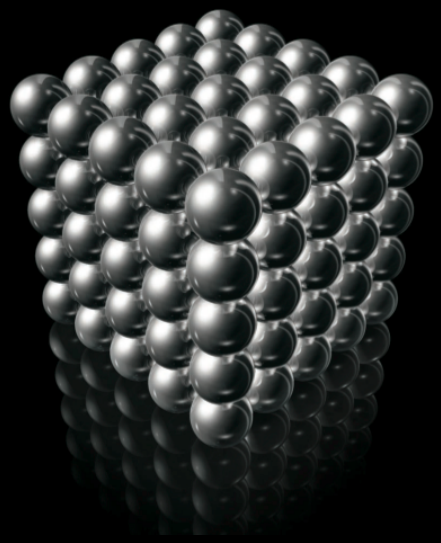

Presenters

John Soward
Rick Wylie

Macworld Conference & Expo

Q&A – MacIT® Conference

We are using Google Moderator to take questions for this session.

1. Go to http://tinyurl.com/633v6e
2. Pick the topic that matches this session
3. Sign in using a Google Account
   User Name: macworldexpo09
   Password: macworld09
4. Submit the questions you want to ask
5. Vote on others' questions you want answered

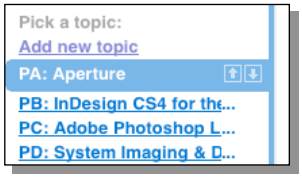

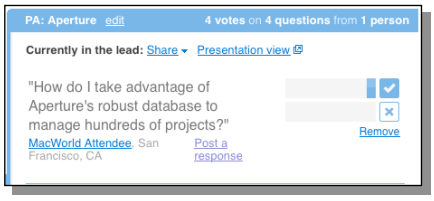

Macworld Conference & Expo

Introductions

- Rick Wylie

  CEO KeyOptions.au

- John Soward

  IT Manager at University of Kentucky, part-time consultant.

- Yourselves....

Syllabus

- ~~Introductions~~
- History
- Fundamentals
- Directory & DNS Integration
- Logical Considerations
- Physical considerations
- Documentation
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

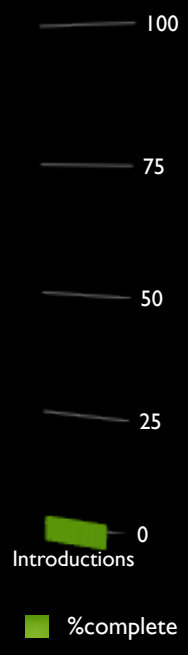

Macworld Conference & Expo

MacIT Conference: Xsan and Storage

History: how did we get here?

1989

Lockheed Martin

MountainGate
(Metrum / CentraVision)

ADIC
StorNext

Apple
Xsan 1.0

Quantum
StorNext/iMover/DLC

Apple
Xsan 2.0

not to scale

2009

Macworld Conference & Expo

Syllabus

- ~~Introductions~~
- ~~History~~
- Fundamentals
- Directory & DNS Integration
- Logical Considerations
- Physical considerations
- Documentation
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

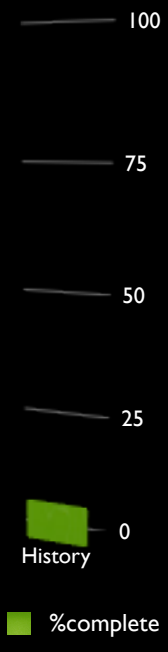

Macworld Conference & Expo

•Xsan Fundamentals

Required Components

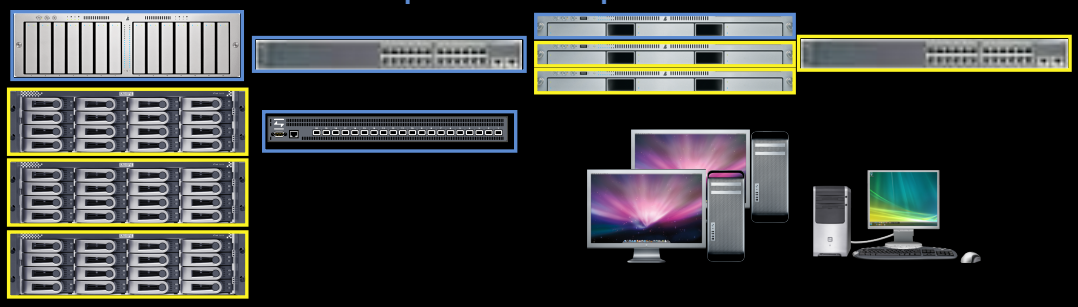

Common Components

- •MDC - MetaData Controller
- •Storage Arrays
- •MetaData LAN
- •FibreChannel Switch
- •Clients
- •General LAN

Macworld Conference & Expo

Storage Fundamentals

Direct Attached Storage

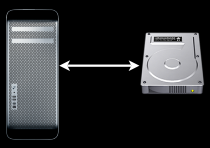

Network Attached Storage

Storage Area Network

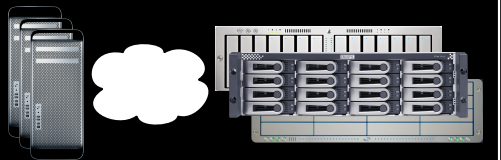

Pro	Con
•Cheap •Fast •Simple	•Limited Expandability •Failure Prone •Single Host
•Inexpensive •Fairly Simple •Well Known	•Limited Speeds •Limited Expandability •Single Point of Failure
•_VERY_ Fast •Expands to PBs •Failure Resistant	•Expensive •Complicated

Macworld Conference & Expo

Storage Fundamentals

Software, filesystems, & protocols

Cluster File System

synchronous, asynchronous, hybrid

Other Examples: QFS (sun/oss),Lustre (sun/oss), OCFS (oracle), GPFS (ibm),VMFS (vmware), GFS (redhat)

Storage Fundamentals

Local FS: HFS+, FFS, FAT32, NTFS...

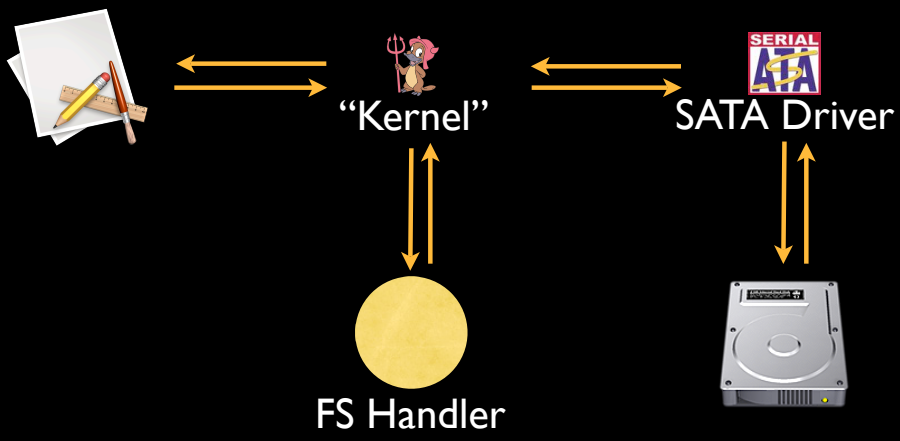

Storage Fundamentals

NAS: AFP, SMB, NFS,...

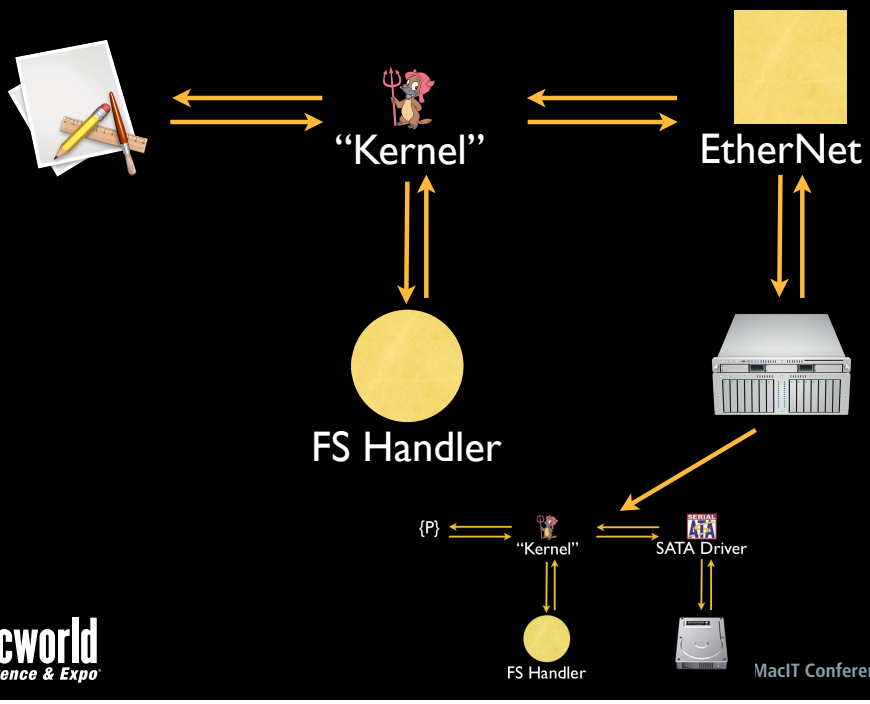

Macworld Conference & Expo

Storage Fundamentals

ClusterFS: Xsan, QFS, Lustre, GFS,...

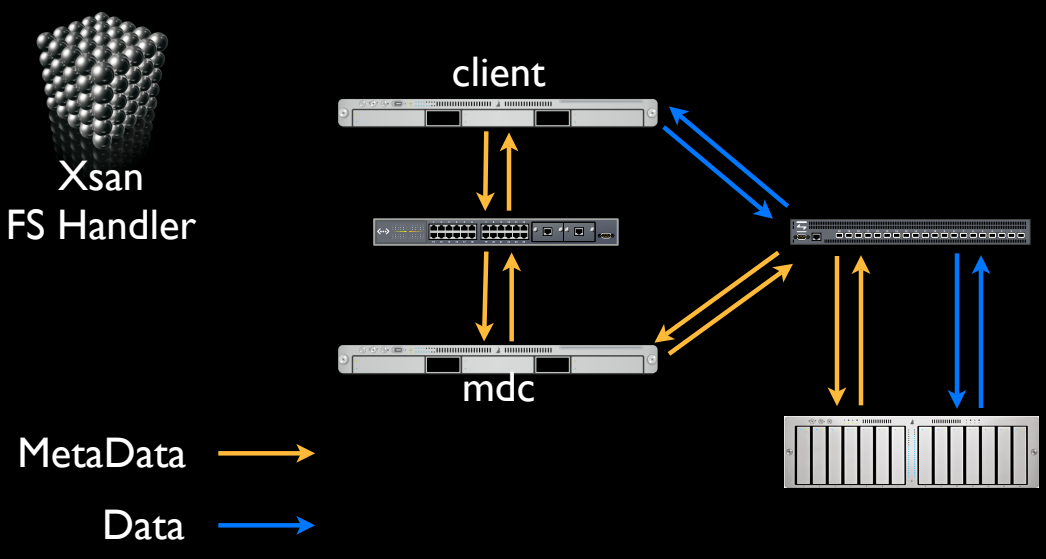

Macworld Conference & Expo

Storage Fundamentals

LUNs

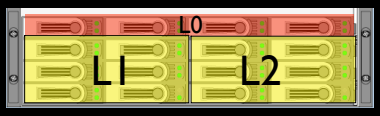

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- Directory & DNS Integration
- Logical Considerations Text
- Physical considerations
- Documentation
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

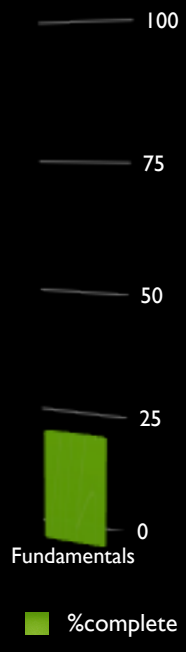

Directory & DNS Intergration

DS, DNS, Network infrastructure.

Do I need to use a centralized Network Directory Service, like OD or AD?

More than likely, YES.

On _all_ the machines?

If you use it, use it everywhere.

What about DNS? for MD?

More than likely, YES.

Forward _and_ Reverse?

DNS must resolve correctly both ways.

DNS Best Practice

- **Make sure you have Forward and reverse DNS records (FQDN) for Public and Private networks - BEFORE YOU EVEN START**
- Sluggishness with the XSAN Admin Tool usually indicates a DNS issue

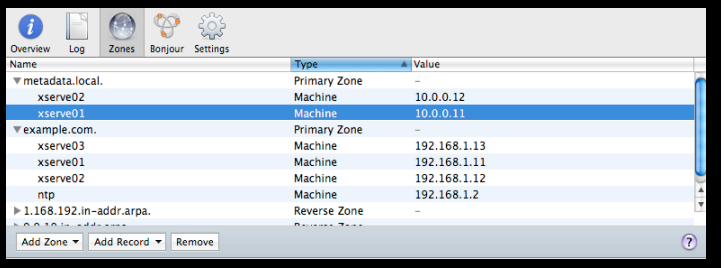

DNS Best Practice

- Verify records using changeip -checkhostname
- Public and Private zones can be on the same DNS server
- Disable unused NICs on client and servers
- No need to Switch off IPV6 anymore (apparently!)
- FQDN

MetaData network Prefs

- No router - even though you get an error
- No DNS - even though you have set one up!
- No Search Domain
- Not the Primary network interface - that's the default route

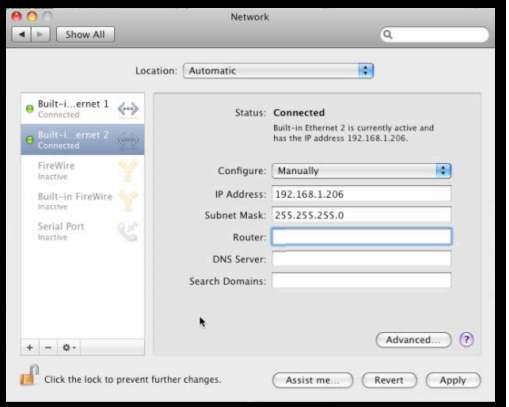

Config and Backup Guidelines

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- ~~Directory & DNS Integration~~
- Logical Considerations
- Physical considerations
- Documentation
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

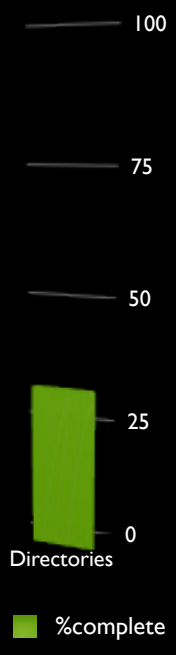

Macworld Conference & Expo
MacIT Conference: Xsan and Storage

Logical Considerations

Storage Capacity
- current, expansion, by type

Performance
- per client node, overall system

Interoperability
- resharing, StorNext clients/MDCs

Availability & Integrity.
- HA, very HA, san mirroring, backups

Application Support
- Case sensitive, 'posix' only.

Maintainability
- custom or stock

Infrastructure Planning

MD network
- Gig-E.
- 1500 MTU
- Dedicated switches
- path/label distinctly

Affinities

Multisan

Logical Considerations

Storage Capacity

Some 'Default' Limitations/Suggestions

Xsan systems in a cluster	64
Xsan Volumes in a cluster	16
Promise LUNS in an array	32
RAM per Volume in MDC	2G

Faster drives are smaller

Recommended /Supported configurations may hold surprises.

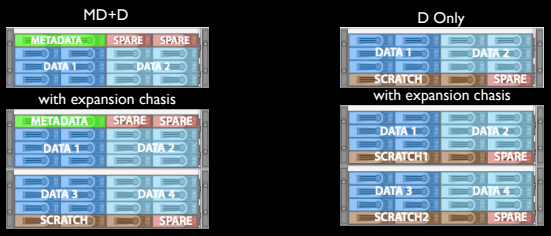

Logical Considerations

Performance: approximate speed calculations

Format	bandwidth
miniDV	3.6MB/s
DVCPRO HD	14MB/s
SD 8bit	20MB/s
720p@24fps	42.19MB/s
8bit 1080 24p	94.92MB/s
10bit 1080i	158.2MB/s

5 clients

1 @ 1080i
2 @ DVCPRO HD
2 @ SD

1	158.2	158.2
2	14	28
2	20	40
Total BW		**226.2**

single 2G FC per client
4G FC to Arrays
1 Array Controller for MD
1 Array Controller
8+ SATA Drives*

*SATA drive 75MB/sec
-20% controller overhead = 50MB/sec
ceil(226.2/50)=5+1(raid5)+2(MD)=8
YMWV.

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- ~~Directory &DNS Integration~~
- ~~Logical Considerations~~
- Physical considerations
- Documentation
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

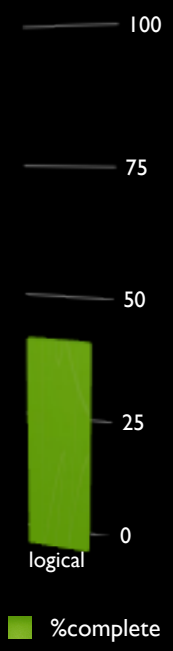

Infrastructure Planning

Location, Location, Location

- Access
- Cabling
- Security
- Expansion
- Sound

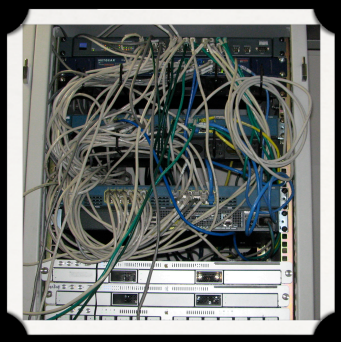

Infrastructure Planning

Power & Cooling & UPS

Xserve (early '08) 2x3.0 GHz	idle 220W, max. 317W
Xserve RAID	av. 300W, max. 400W
Promise VTrak with SAS HDD	524W
3COM 3870 48 ports	220W
Xserve (early '08) 2x3.0 GHz	idle 748 BTU/hr, max. 1078 BTU/hr
Xserve RAID	max. 1365 BTU/hr
Promise VTrak with SAS HDD	max. 1780 BTU/hr
3COM 3870 48 ports	150 BTU/hr

Resources:

www.apc.com/tools/ups_selector

support.apple.com (search for BTU)

Cabling and equipment for Xsan

Copper or Glass?

- Copper is fine for smaller installations
- Glass - No interference
- Check if the switch is ratified for copper, many are not!
- Make sure you use reputable and qualified equipment
- Cabling - Don't skimp, you'll pay for it dearly later!!!

Cabling and Tranceivers

- Qualified components
- Apple Copper cables (4GB rated)
- Fibre-optic cable with matching transceivers (4GB Rated)
- Finisar 850nm SFPs

50µm or 62.5µm?

- Use one or the other if you are cabling your Core XSAN components
- Longer distances don't matter so much with mixing, but some degradation in performance
- 50µm has a higher bandwidth (2 standards of 50µm, OM1 and OM3)
- Make sure cables are OM3 standard (laser ratified), can support up to 550m depending on application. Usual rule of thumb is maximum 300m

50µm or 62.5µm? ... Continued

- 62.5µm cable is being phased out. Originally designed for LED's which modulate at a maximum of 622Mbps
- New 50µm has nearly 3 times the bandwidth capacity than of old 62.5µm cable (500MHz/km vs. 160MHz/km)

Fibre Channel Speed basics

Product Name	Throughput* (MBps)	Rate (Gbaud)	Availability
1GFC	200	1.0625	1997
2GFC	400	2.125	2001
4GFC	800	4.25	2005
8GFC	1,600	8.5	2008
16GFC	3,200	17	2011

*Throughput includes transfer of data in both directions

Fibre considerations

Media Type	Speed (MB/s)	Transmitter	Variant	Distance
Single-Mode Fiber	400	1300 nm Longwave Laser	400-SM-LL-I	2 m - 2 km
	200	1550 nm Longwave Laser	200-SM-LL-V	2 m - >50 km
		1300 nm Longwave Laser	200-SM-LL-I	2 m - 2 km
	100	1550 nm Longwave Laser	100-SM-LL-V	2 m - >50 km
		1300 nm Longwave Laser	100-SM-LL-L	2 m - 10 km
		1300 nm Longwave Laser	100-SM-LL-I	2 m - 2 km
Multimode Fiber (50μm)	400	850 nm Shortwave Laser	400-M5-SN-I	0.5 m - 150m
	200		200-M5-SN-I	0.5 m - 300m
	100		100-M6-SN-I	0.5 m - 300m
			100-M6-SL-I	2 m - 175m

Cabling Do's and Don'ts

- Follow ALL the basic rules of cabling for either copper or glass.
- Don't cable tie glass cables
- Follow the 8 x diameter rule
- Check your termination
- NEVER skimp on the cabling - You will pay for it later!!!!

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- ~~Directory Integration~~
- ~~Logical Considerations~~
- ~~Physical considerations~~
- Documentation
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

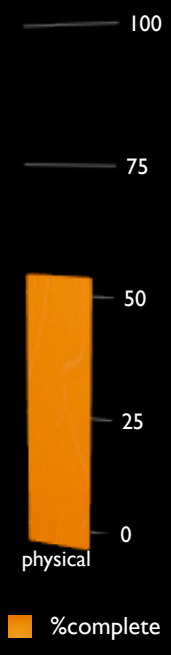

Documentation

What you need. (at least)

- San Map
- Array Map
- Network Map
- Backup info
- Anomalies/Exceptions

Documentation

Items to keep in mind

- Use graphics (OmniGraffle helps)
- Use in-place labels when possible
- Keep a printed version
  (attached to the rack if possible)
- Have someone else review it!
- Build it before the install.
- Keep it up to date.

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- ~~Directory & DNS Integration~~
- ~~Logical Considerations~~
- ~~Physical considerations~~
- ~~Documentation~~
- Upgrading from Xsan1
- Gotchas/Tips-n-Tricks

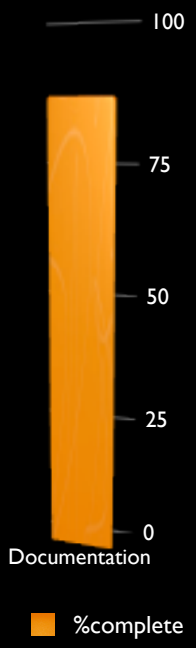

Upgrading from Xsan1

DO YOU REALLY WANT TO?

'ARCHIVE & INSTALL' VS UPGRADE

Upgrading from Xsan1

- Back up the XSAN volumes!
- Unmount and stop the volumes
- Verify the integrity of the volume
- **`cvfsck -vn <volume_name>`**
- Repair the volume if required
- **`cvfsck -vw <volume_name>`**
- Upgrade all Metadata computers to 10.5.x
- Install XSAN 2 on computers, upgrade to XSAN 2.1.1

Upgrading from Xsan1

- Open XSAN Admin and connect to the upgraded MDC
- Add computers to the XSAN
- Add the new licenses
- BE PATIENT!!!!!!!

  The metadata needs to "realign itself" to the new XSAN format. Look for RPL_Upgrade in the syslog on the Primary MDC

  Rule of thumb is 1,000,000 files per hour!!!
- Remount the volume on clients and controllers as required.

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- ~~Directory Integration~~
- ~~Logical Considerations~~
- ~~Physical considerations~~
- ~~Documentation~~
- ~~Upgrading from Xsan I~~
- Gotchas/Tips-n-Tricks

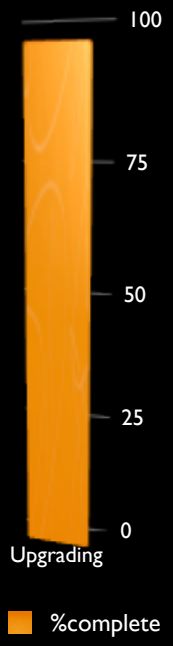

Other 'Gotchas' and Tips

- Buy Promise from Apple (firmware)
- Apple's Promise units do not include sleds
- MUX adapters
- PCI-X risers for new Xserves?
- Tuning
- Time to build Arrays
- Train your users

Syllabus

- ~~Introductions~~
- ~~History~~
- ~~Fundamentals~~
- ~~Directory Integration~~
- ~~Logical Considerations~~
- ~~Physical considerations~~
- ~~Documentation~~
- ~~Upgrading from Xsan1~~
- ~~Gotchas/Tips-n-Tricks~~

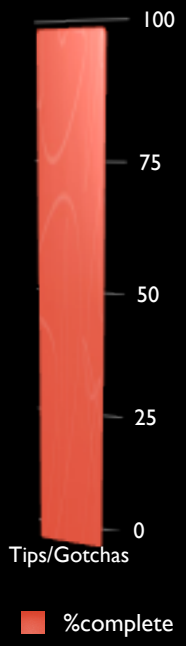

Q&A Friday 3pm. This room.

Google Moderator is being used for the Q&A

1. Go to **http://tinyurl.com/633v6e**
2. Pick the topic that matches this session (IT802)
3. Sign in using a Google Account:
   UserName: macworldexpo09
   Password: macworld09
4. Submit the Questions you want to ask
5. Vote on others' questions you want answered

Next	IT812: Advanced Xsan Configurations	3:30-5:00	121 North
Thursday	IT822: Optimizing Promise RAIDs	10:30-12:00	
Thursday	IT832: Providing Large Scale Backup Solutions	1:00-2:30	
Thursday	IT842: Supporting Final Cut Server	3:00-4:30	
Friday	ITT852: Special Announcement	10:30-12:00	
Friday	IT862: Xsan Case Studies	1:00-2:30	
Friday	IT872: Q&A	3:00-4:00	

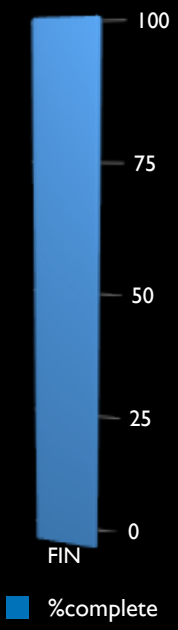

Next session at 3:30

IT812: Advanced Xsan Configurations