Xsan and Beyond

David Colville
Key Options Technology

Robert Kite, Ph.D.
SARCOM

Xsan Concepts:

Storage Area Networks

Storage Networking Architectures

- Direct-Attached Storage - DAS
- Network-Attached Storage - NAS
- Storage Area Network - SAN

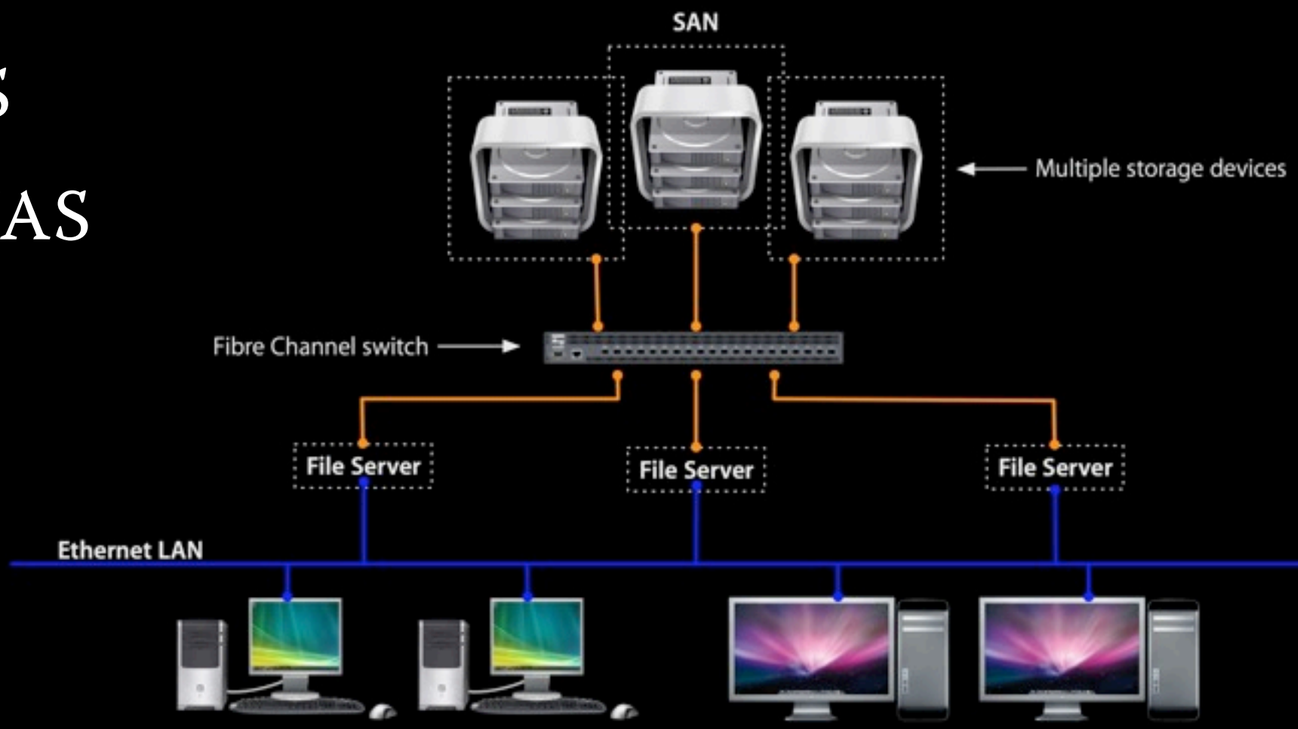

Direct-Attached Storage

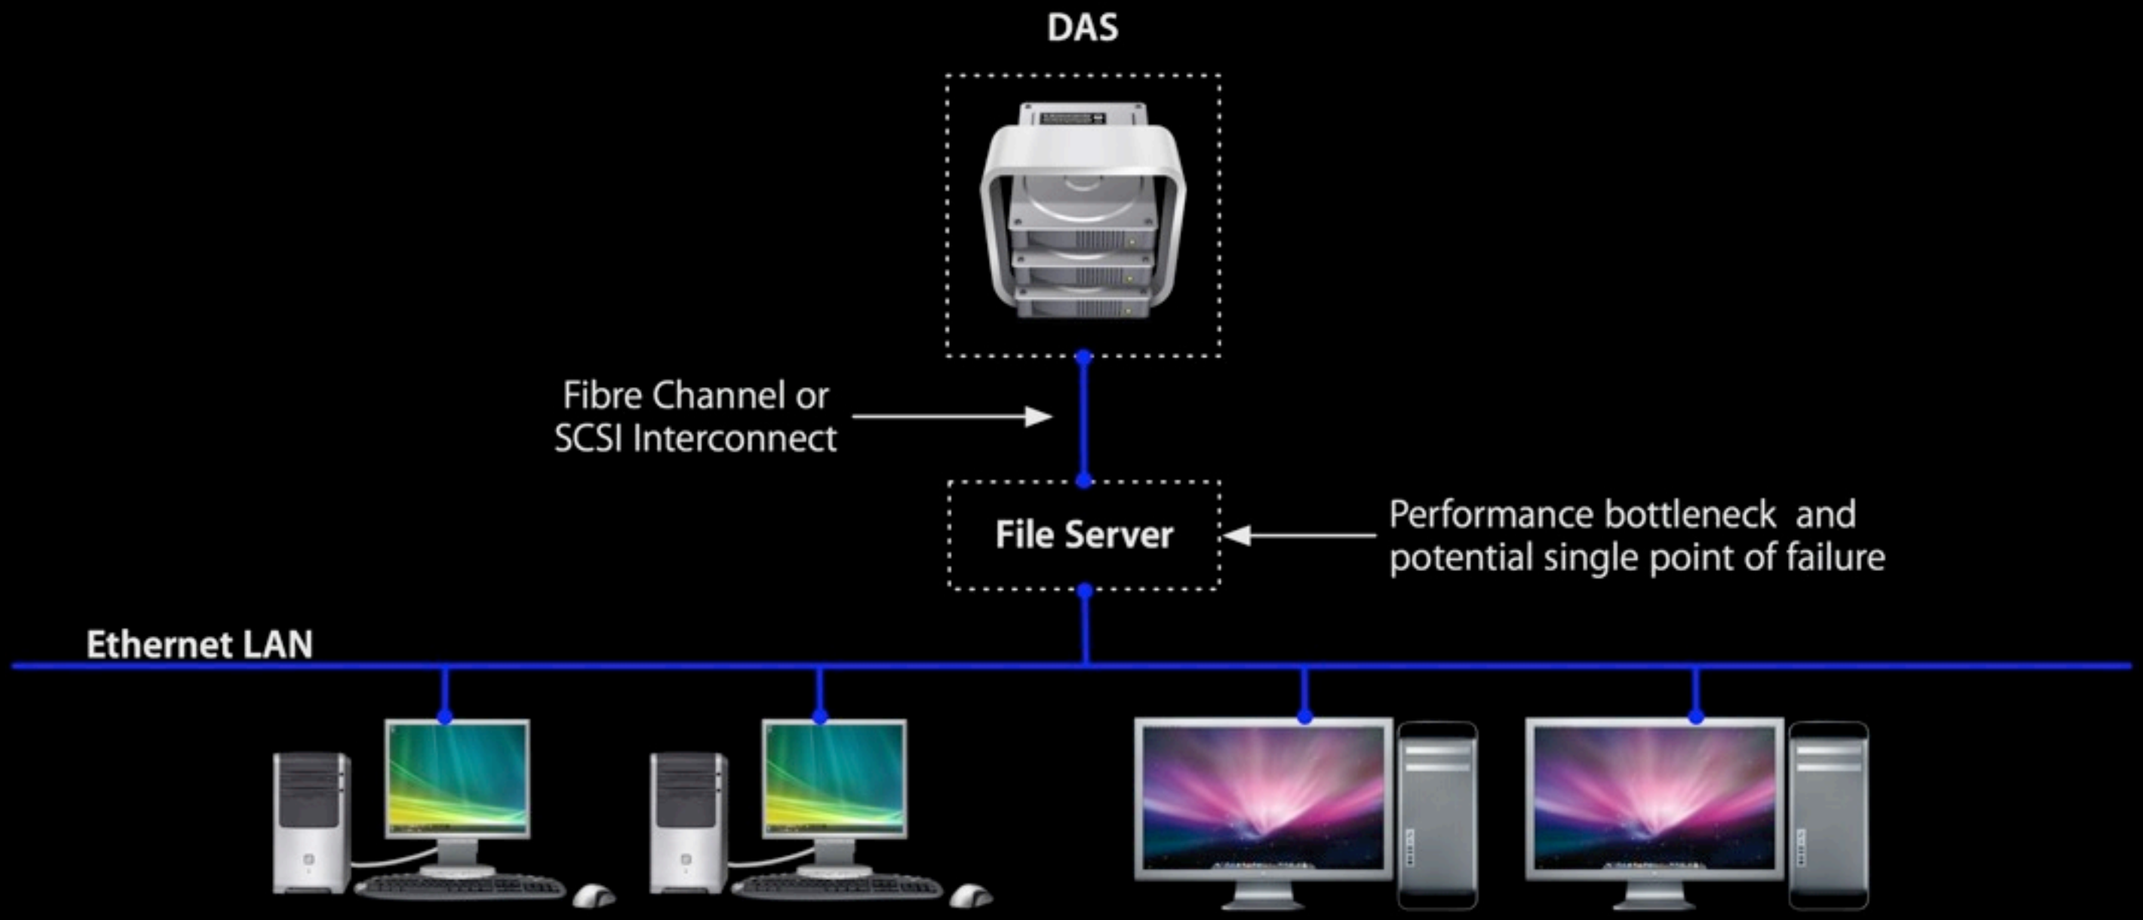

Network-Attached Storage

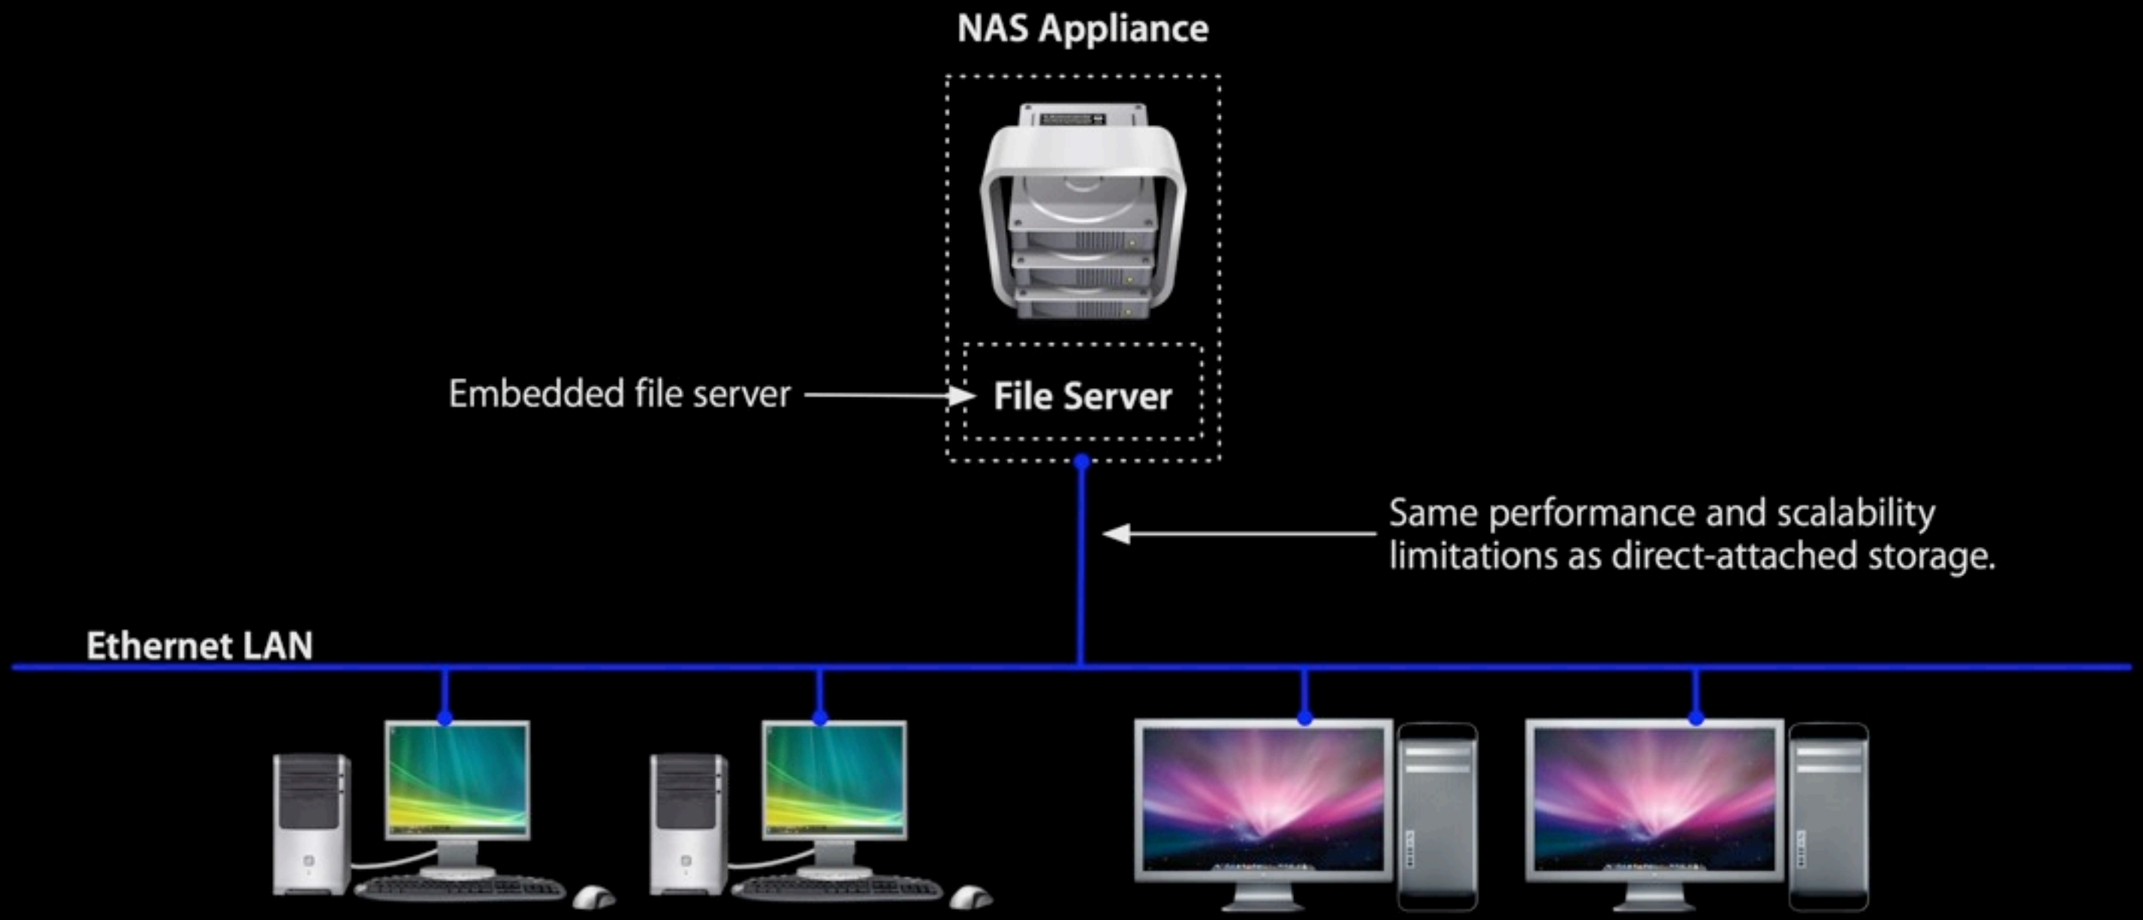

Storage Area Network

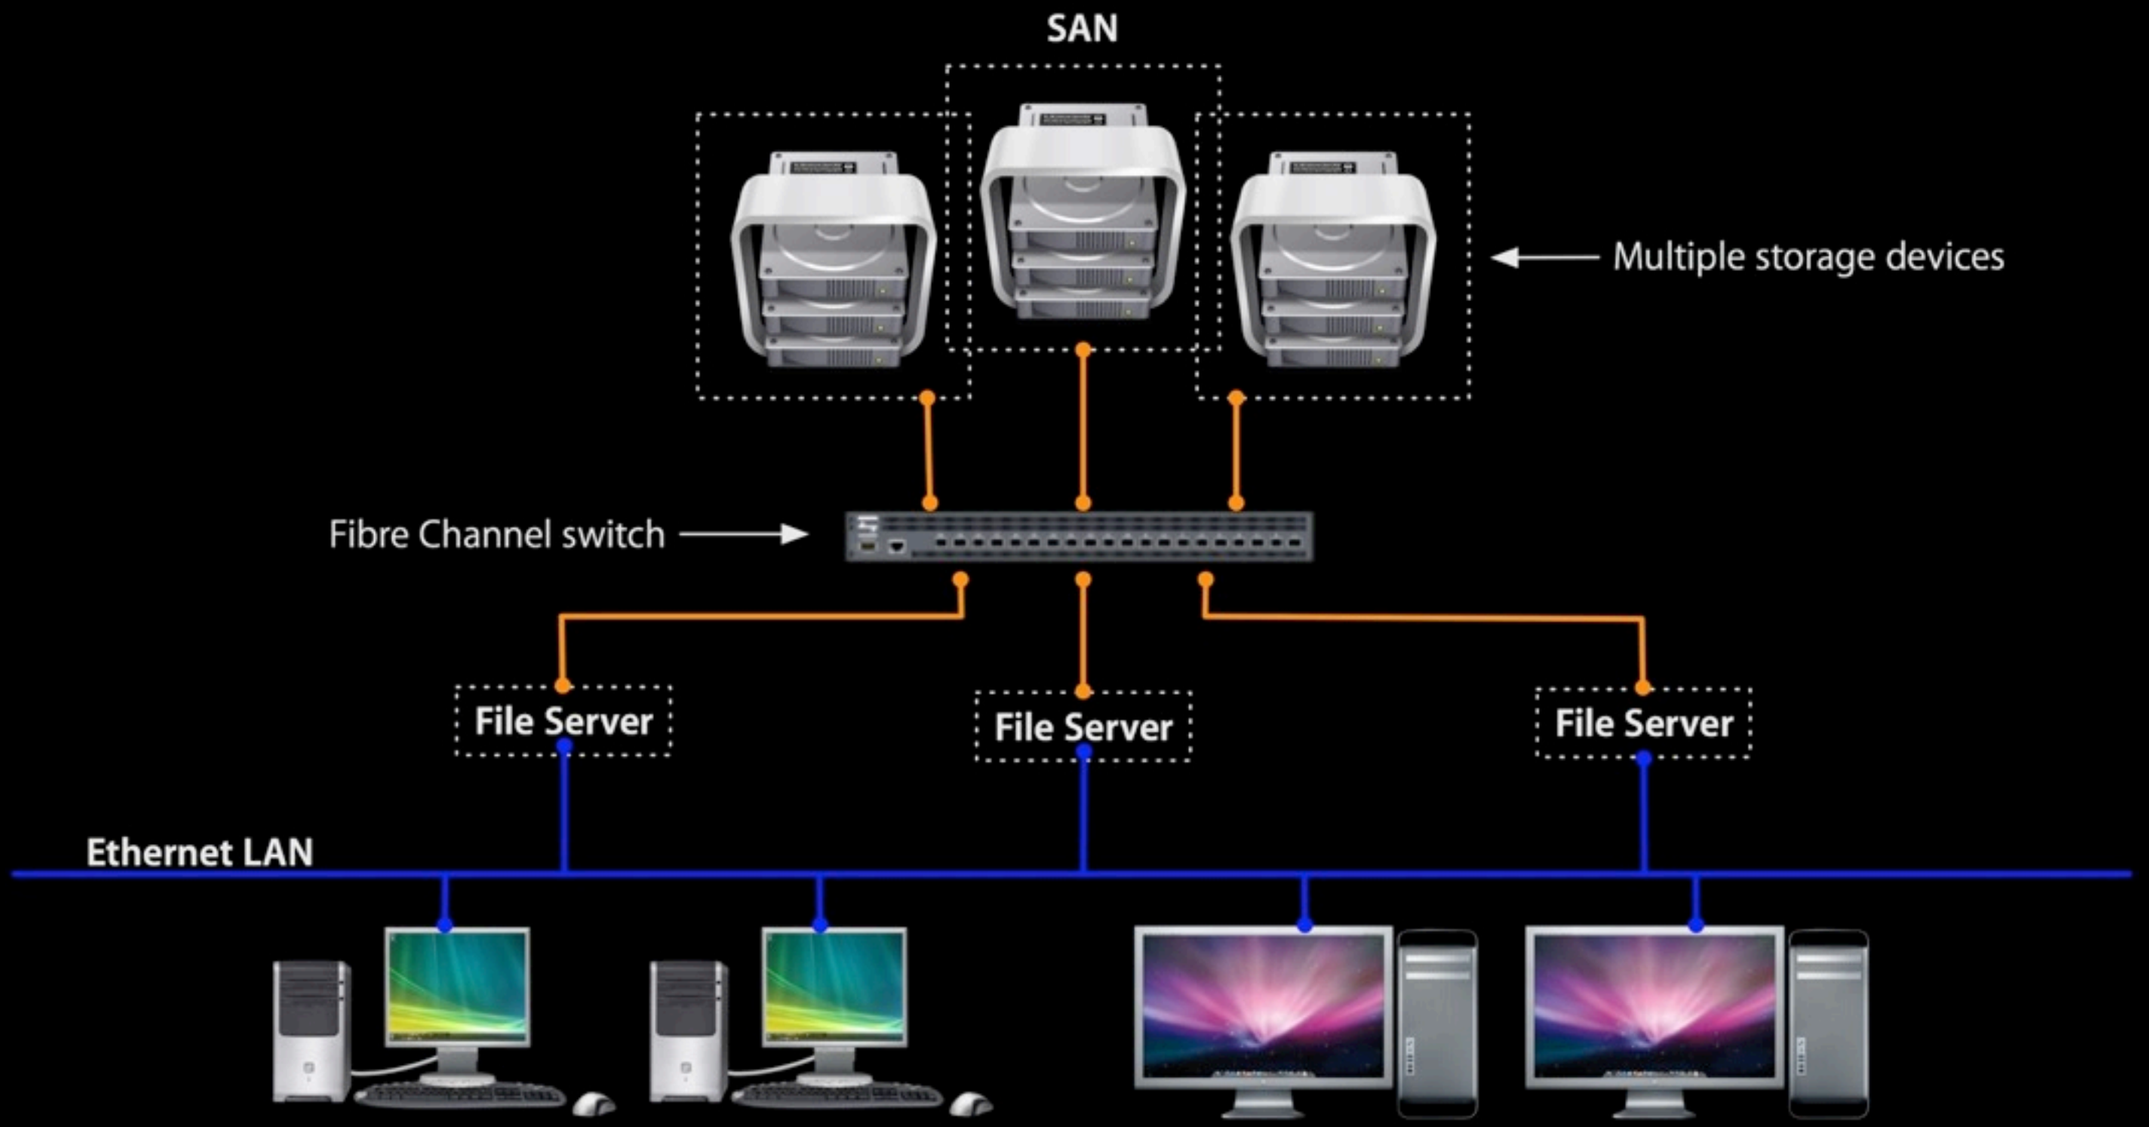

Storage Fundamentals

Direct Attached Storage

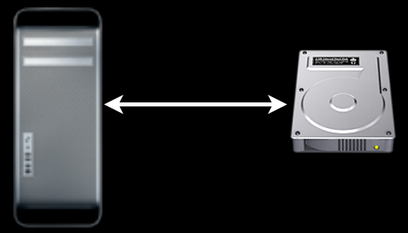

Network Attached Storage

Storage Area Network

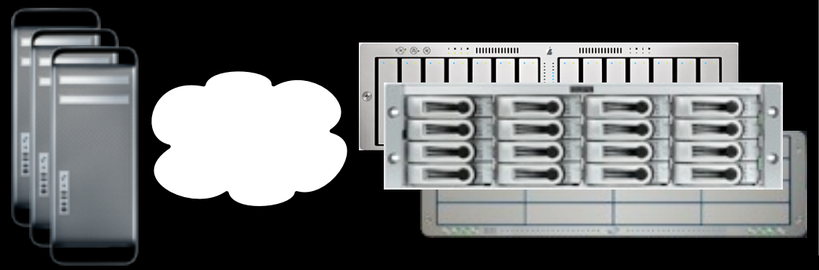

Pro	Con
• Cheap • Fast • Simple	• Limited Expandability • Failure Prone • Single Host
• Inexpensive • Fairly Simple • Well Known	• Limited Speeds • Limited Expandability • Single Point of Failure
• VERY Fast • Expands to PBs • Failure Resistant	• Expensive • Complicated

Benefits of Storage Area Networking

- High-Performance, concurrent file sharing
- Eliminating single points of failure
- Network-based storage management
- Reduced costs through more efficient disk use
- Flexible SAN topology

Fibre Channel Topologies

- Point-to-Point
- Arbitrated Loop
- Switched Fabric

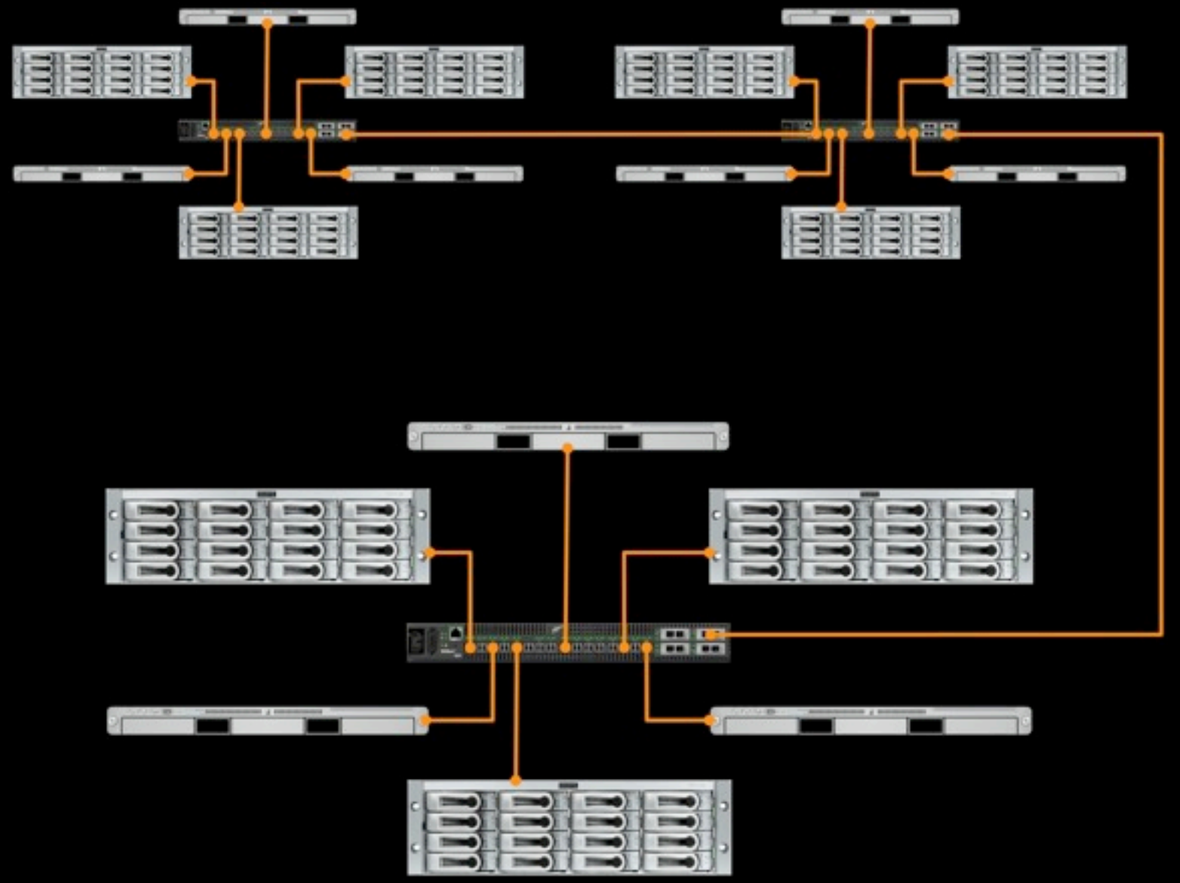

Switched Fabric Topologies

- Cascade
- Cascade Loop
- Mesh Fabric

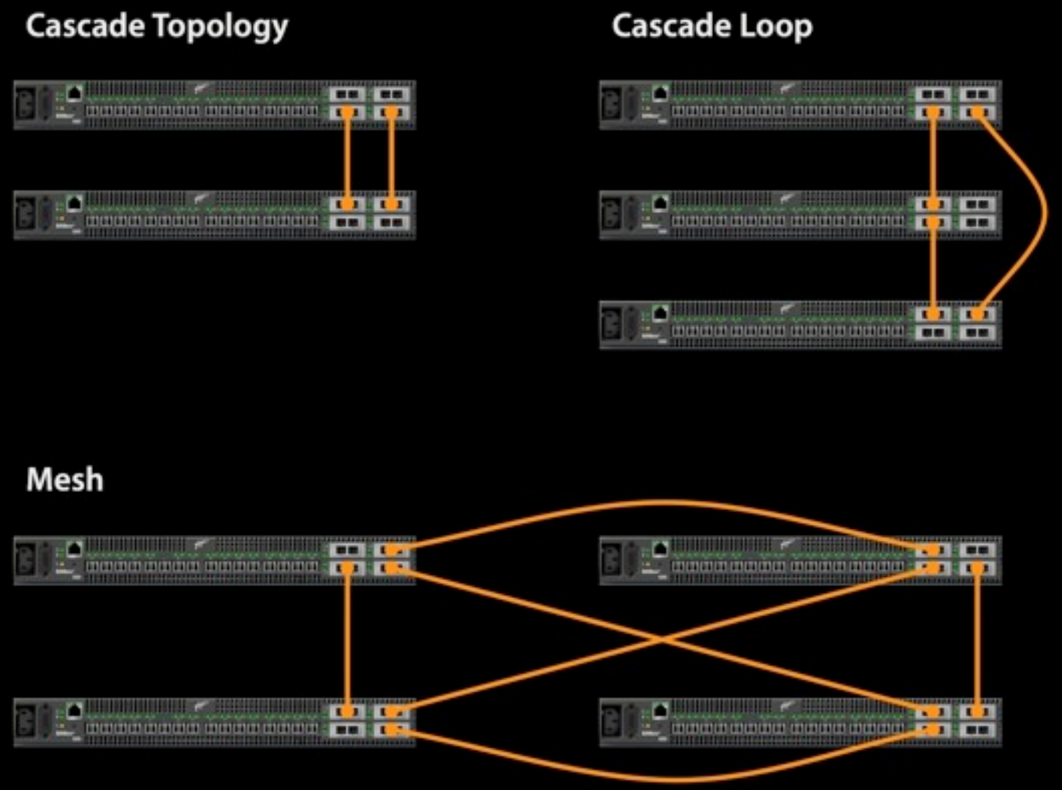

Understanding Xsan

- San Volume
- Fibre Channel Network
- Xsan Metadata Controller(s)
- Xsan Client(s)
- Network Client(s)

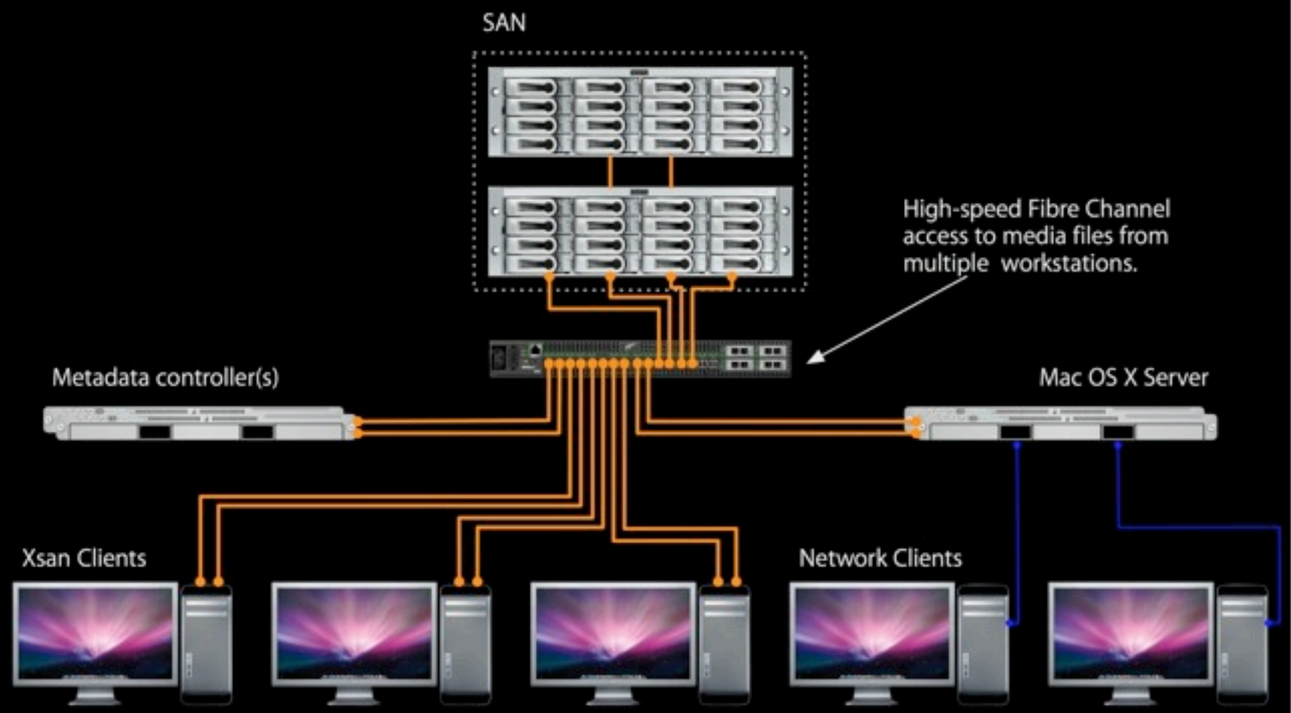

Xsan Design...

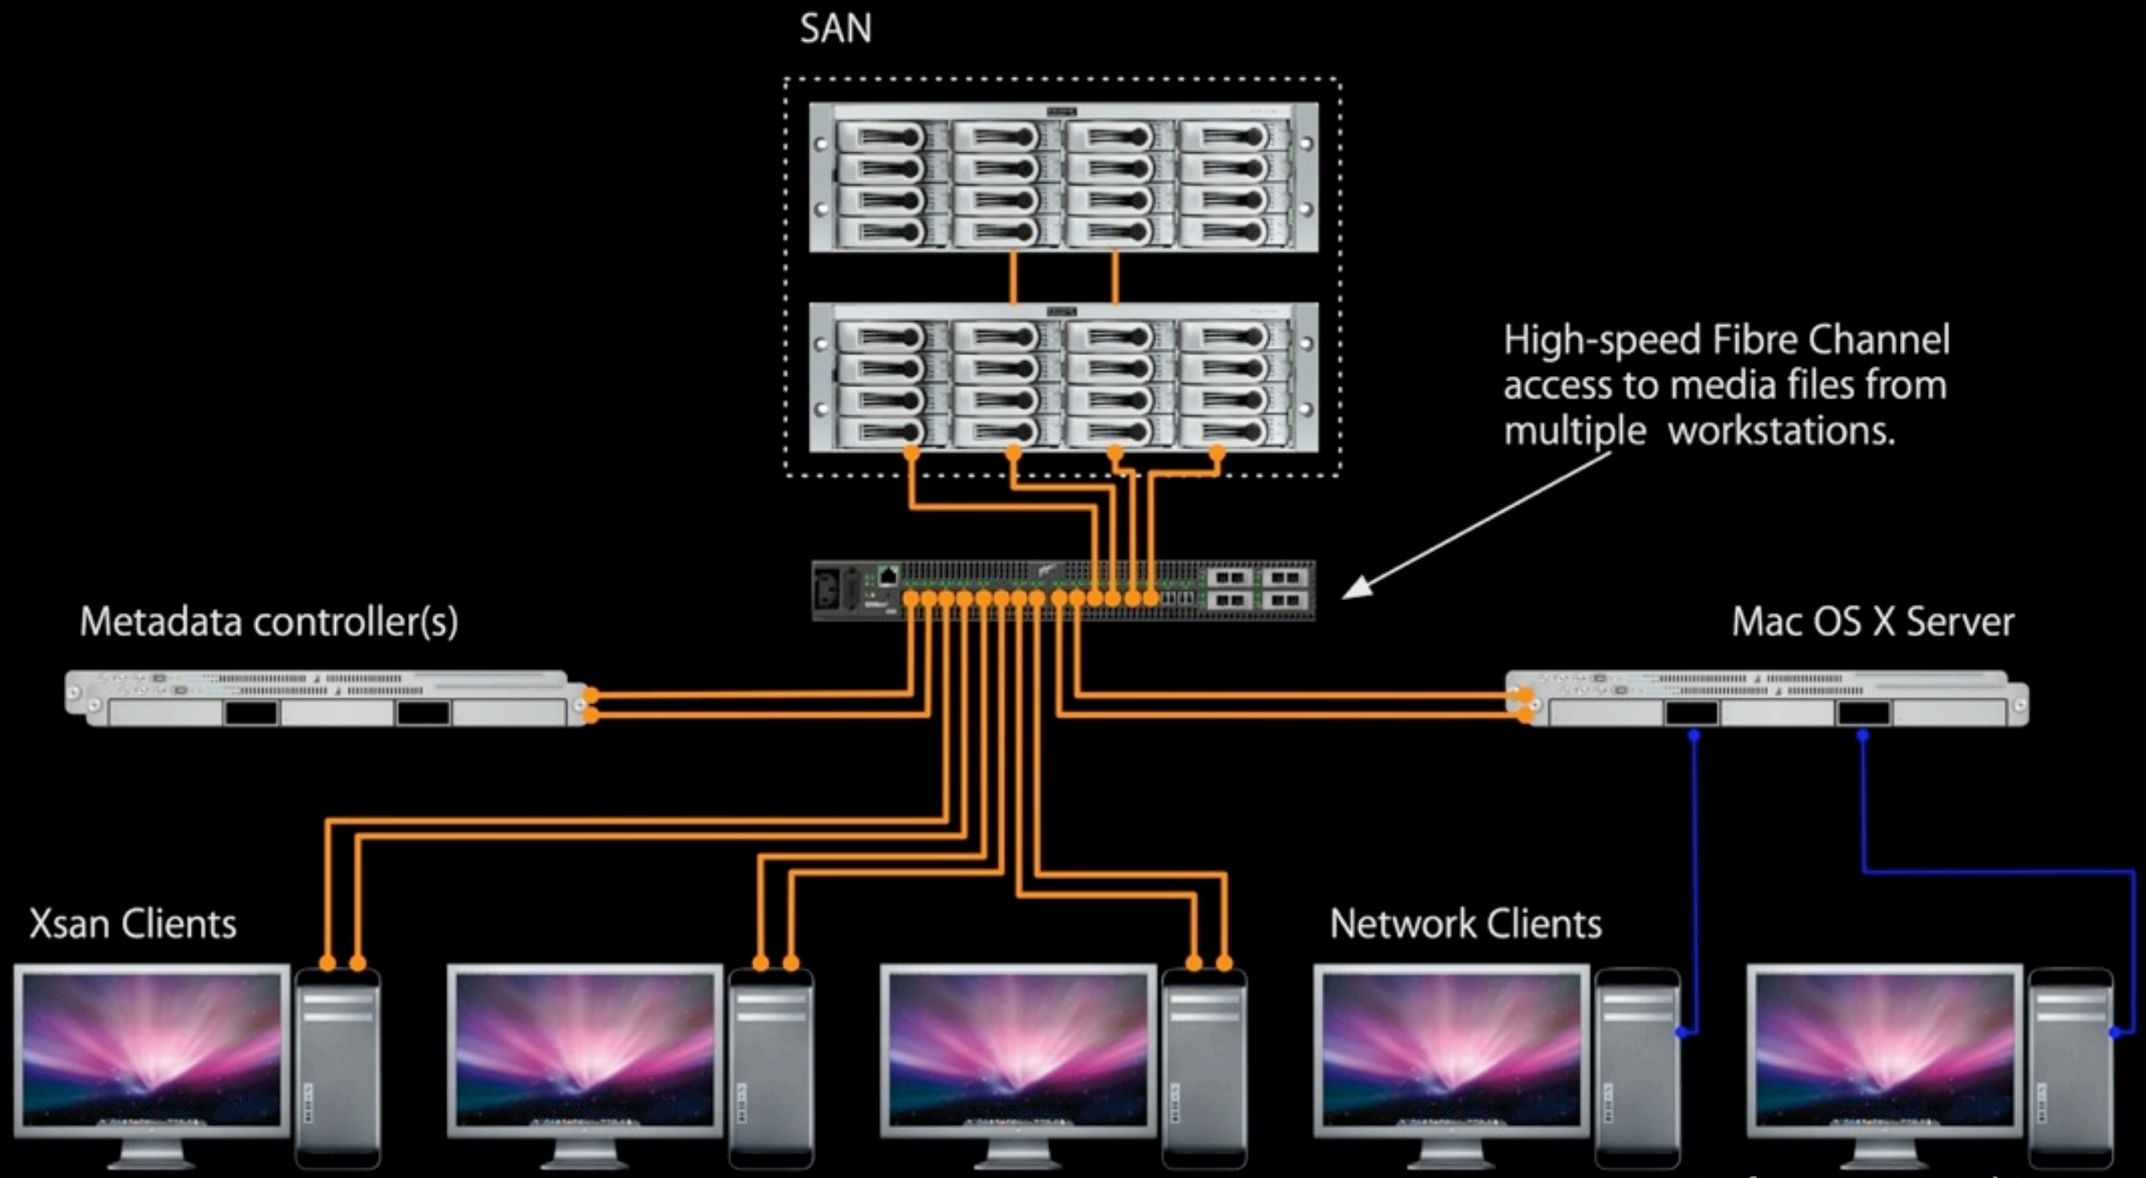
SAN
High-speed Fibre Channel access to media files from multiple workstations.
Metadata controller(s)
Mac OS X Server
Xsan Clients
Network Clients

Xsan 2 Volume Structure

- LUNs
- Storage Pools
- Affinities
- Volumes

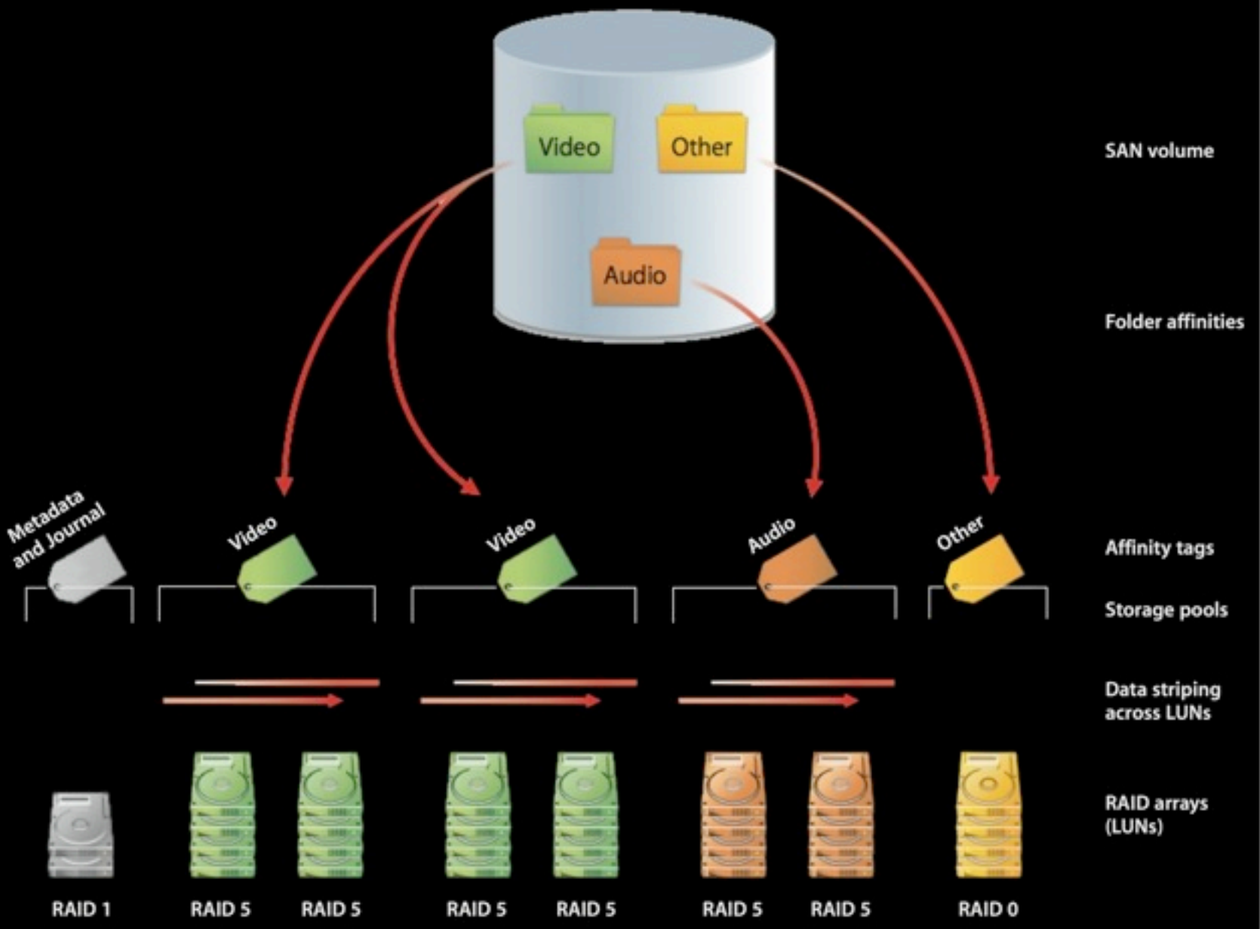

Xsan 2 Volume Structure

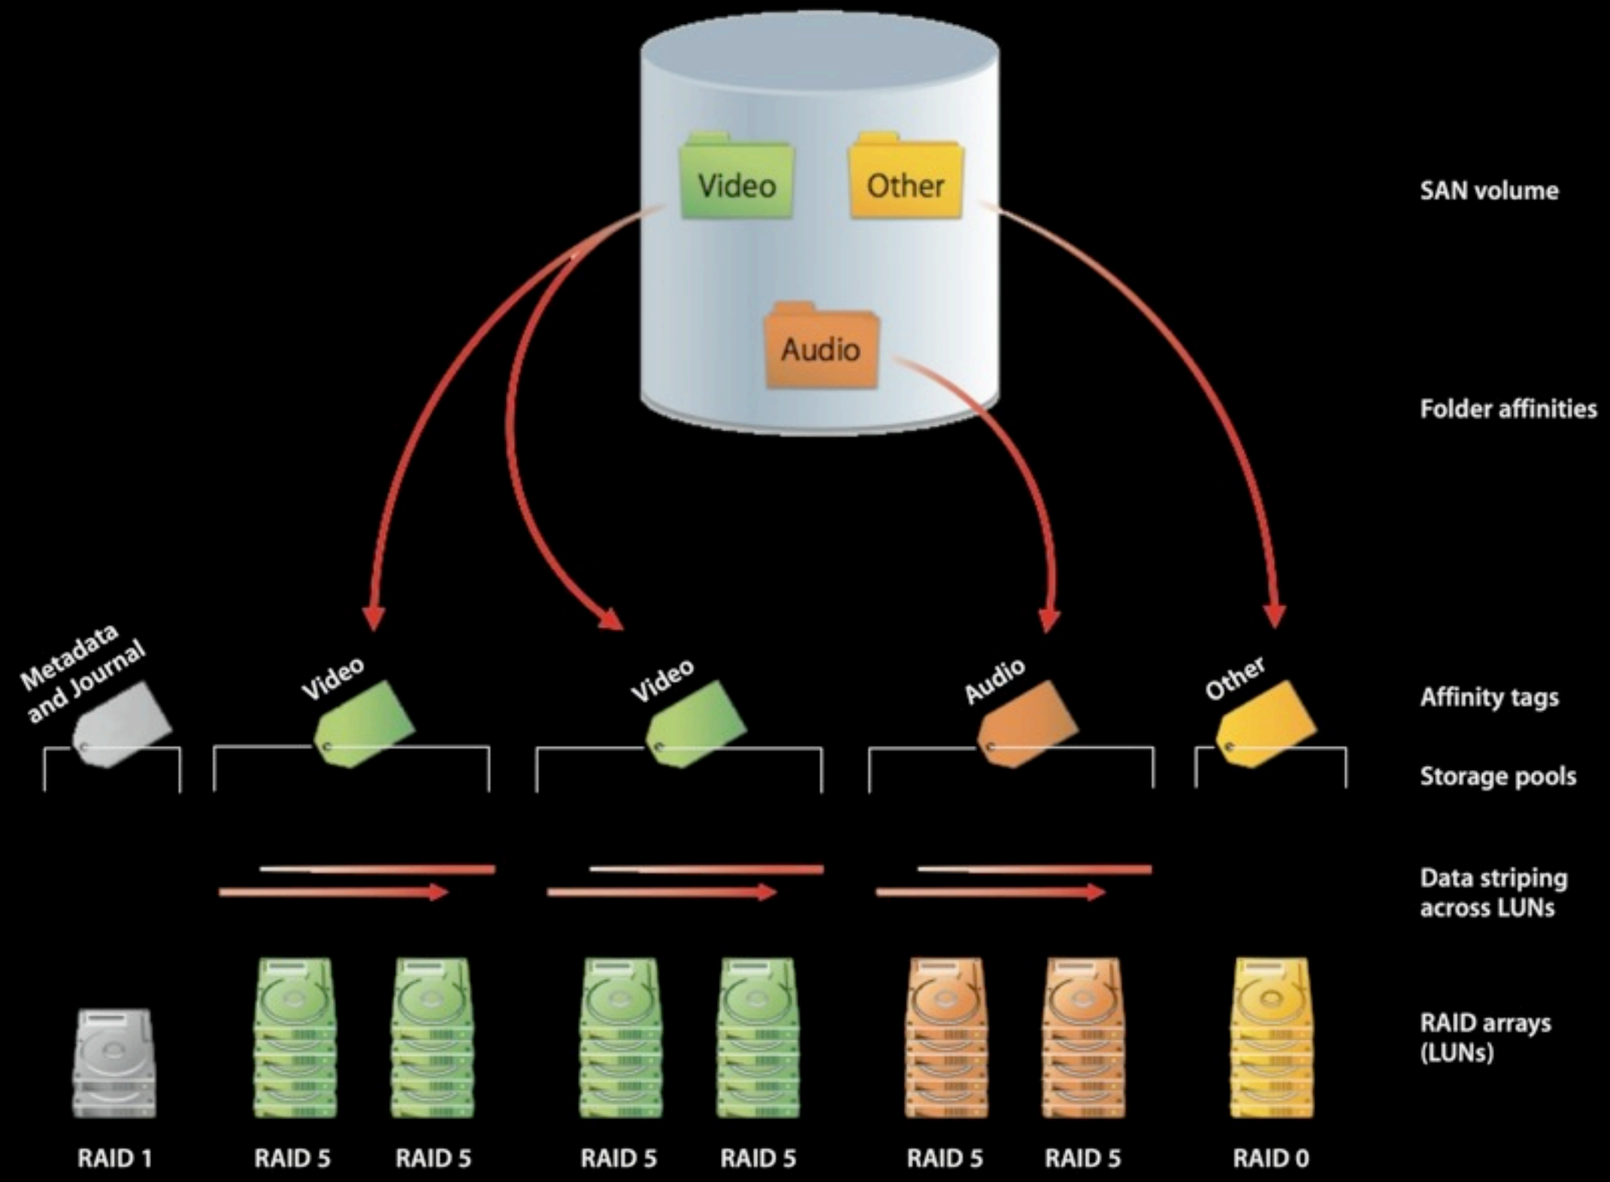

Flexible Volume Management

- Storage pooling
- Dynamic scaling of volumes
- Volume mapping
- Affinities
- Defragmentation

Creating a SAN Volume

Create RAID sets (LUNs)

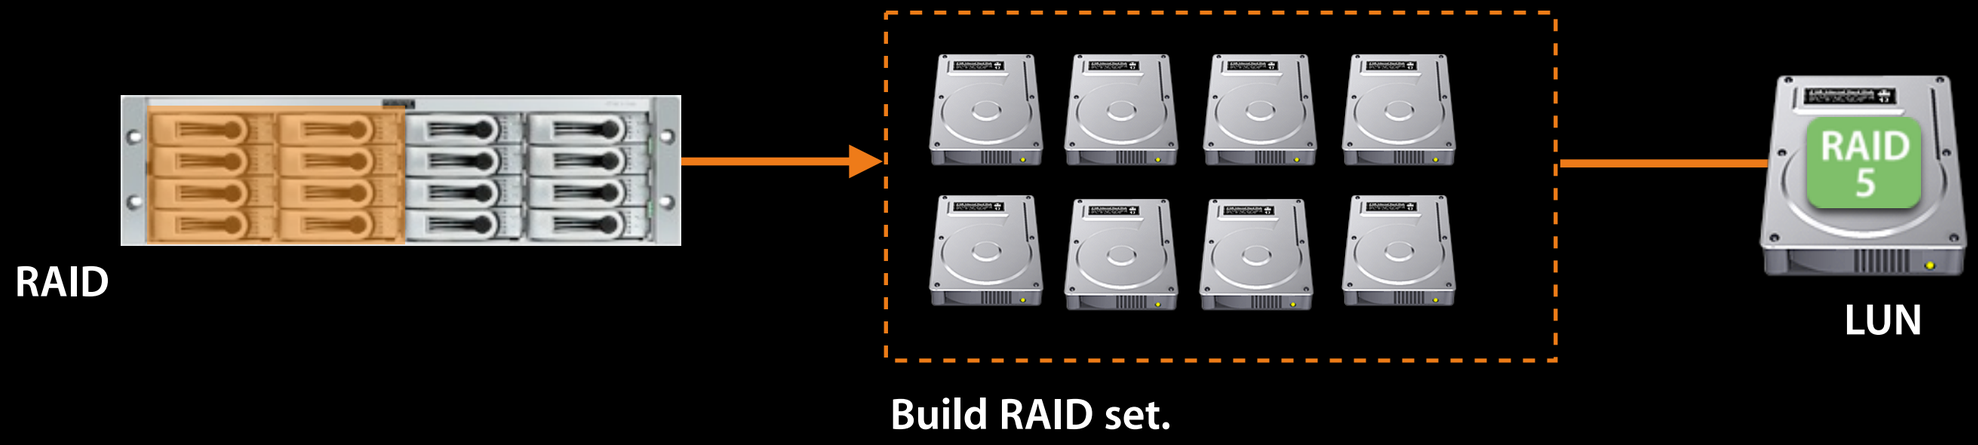

Creating a SAN Volume

Stripe across LUNs to create high-performance storage pool

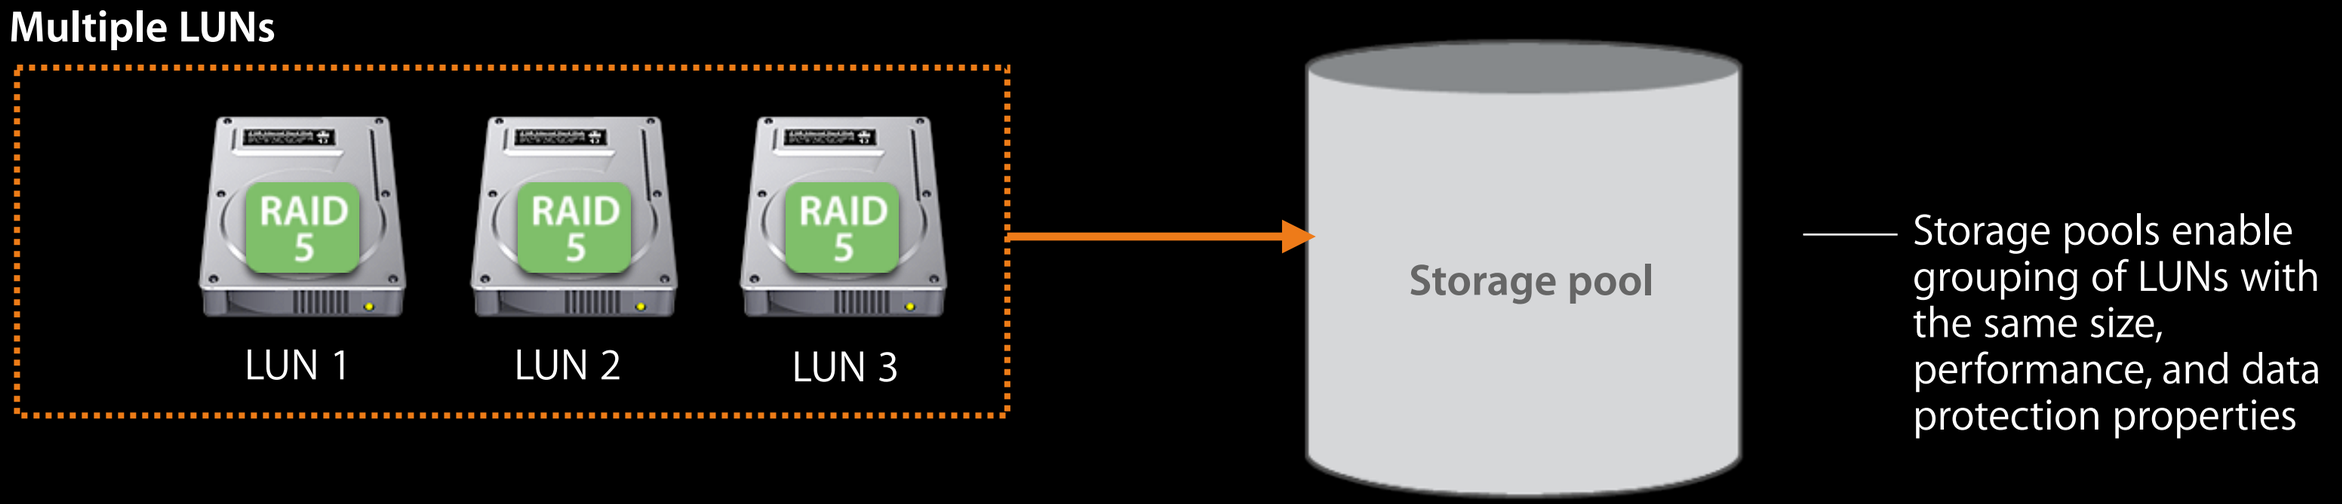

Creating a SAN Volume

Combine storage pools into high-capacity SAN volume

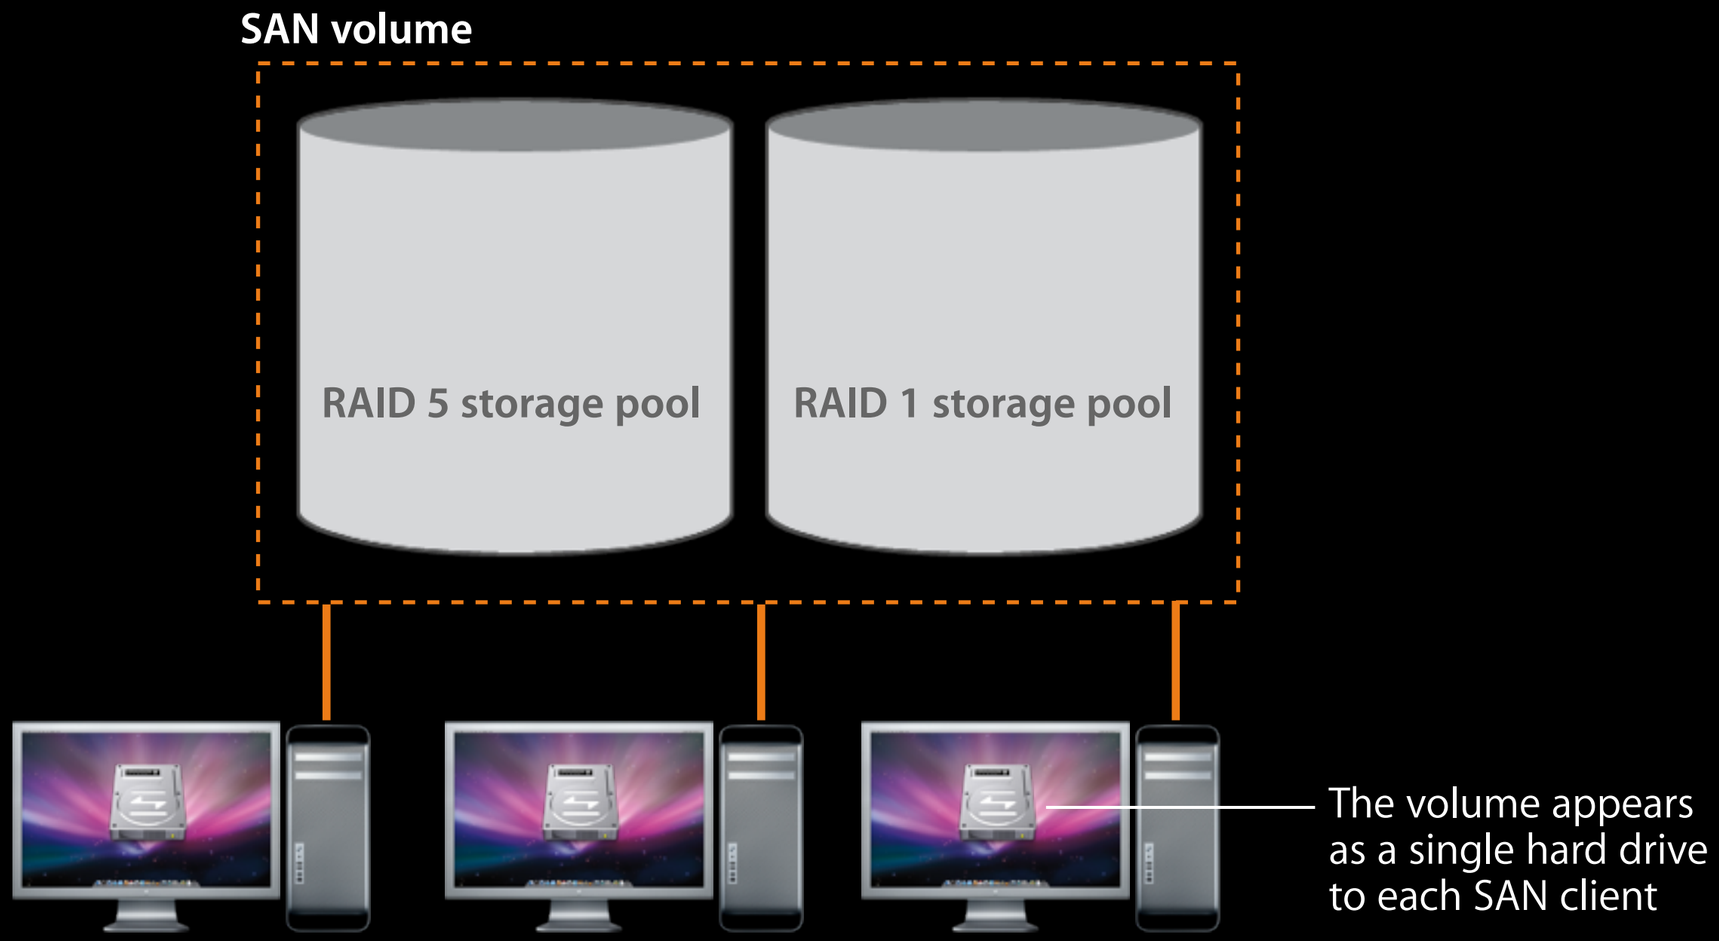

Affinities

Automatically place files in specific storage pools

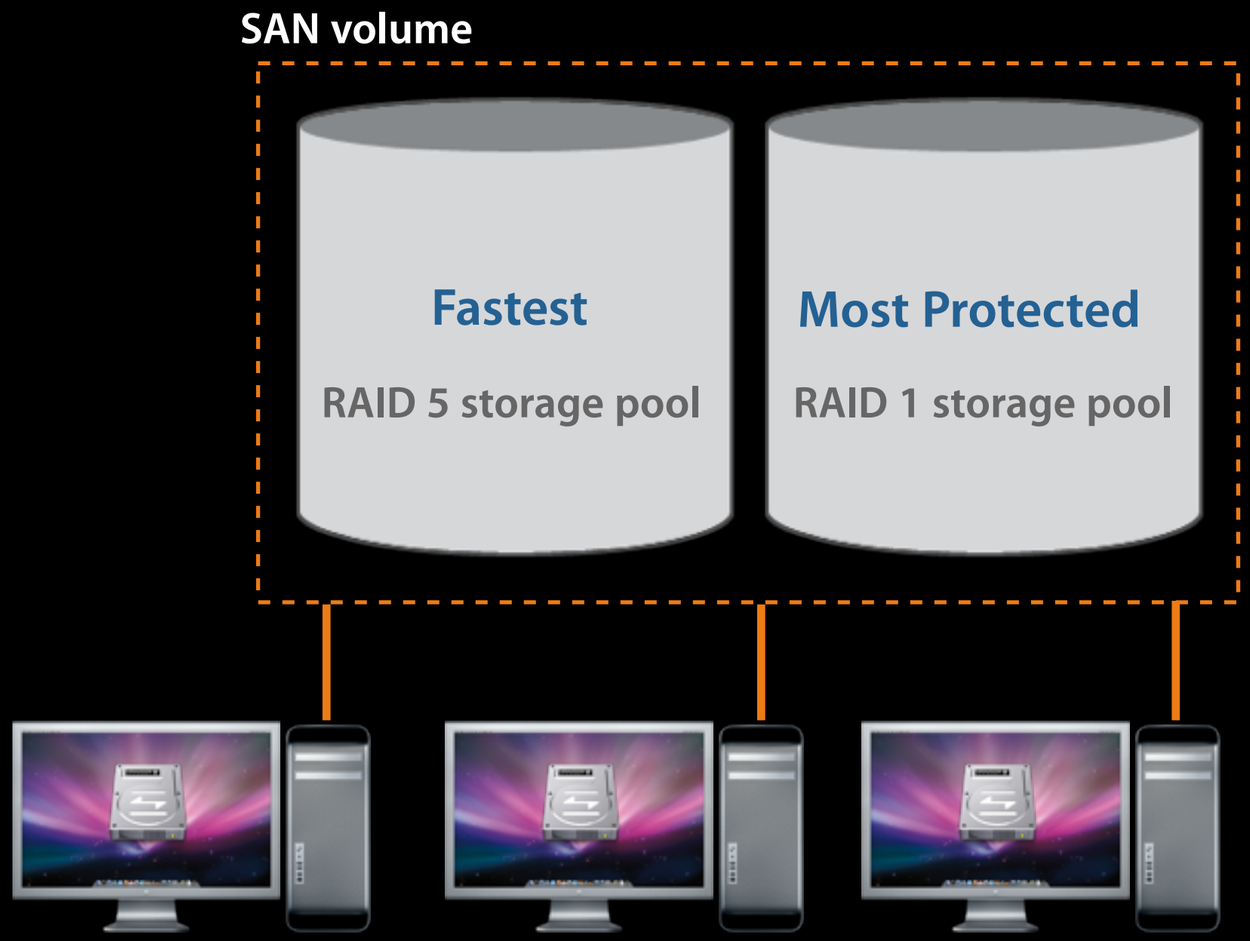

Affinities

Automatically place files in specific storage pools

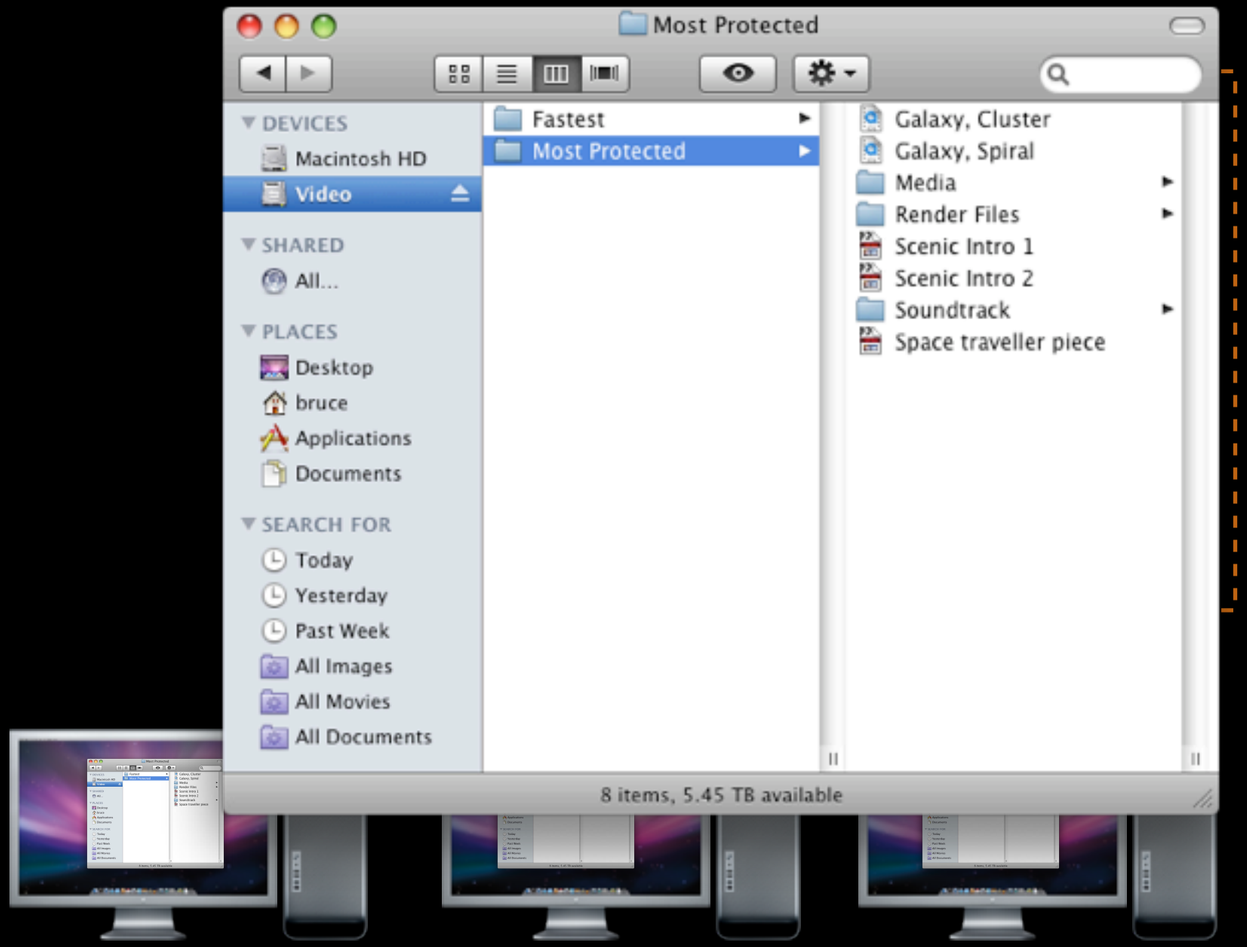

Pools, Affinities, Affinity tags??

Xsan 1.4.x	Xsan 2.1.x
Affinity	Affinity
Pool	Affinity Tag
New terms	Storage Pool/Stripe Group

Stripe Group is a group of Affinity Tags represented as a single affinity..... More soon!

Xsan 2 Volume Structure

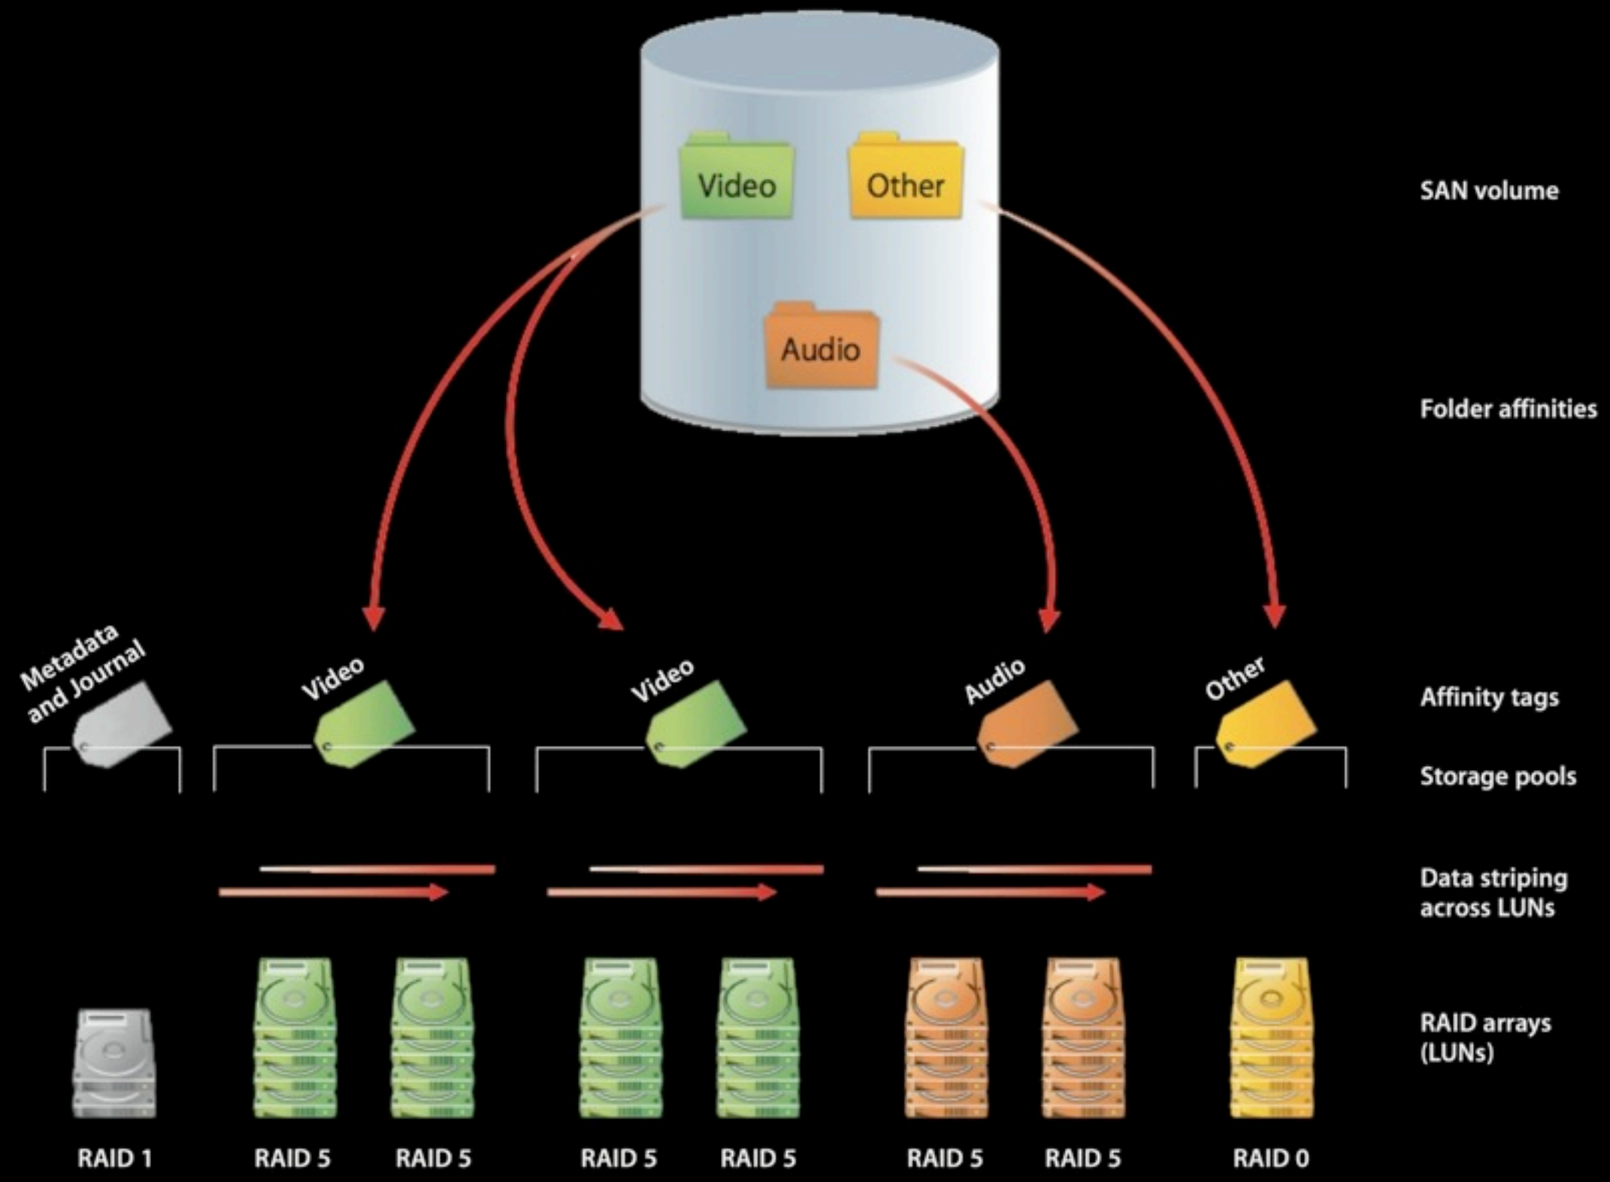

High Availability Architecture

- Metadata controller failover
- File system journaling
- Fibre Channel multipathing

Always Available

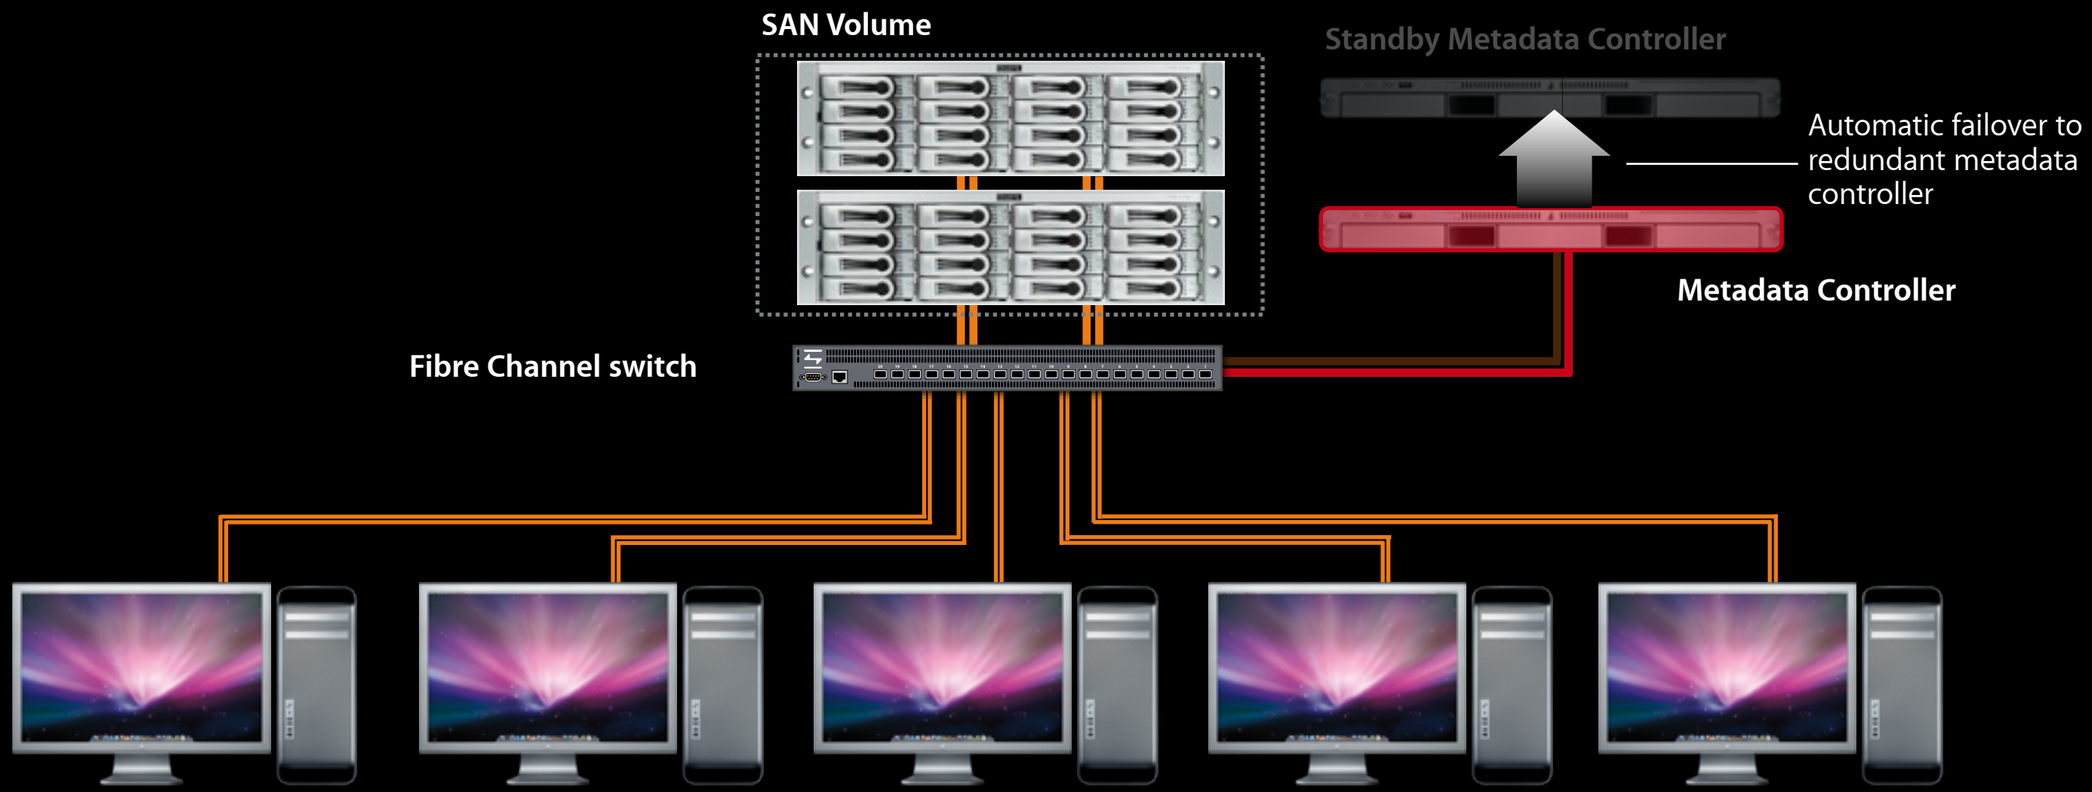

Always Available

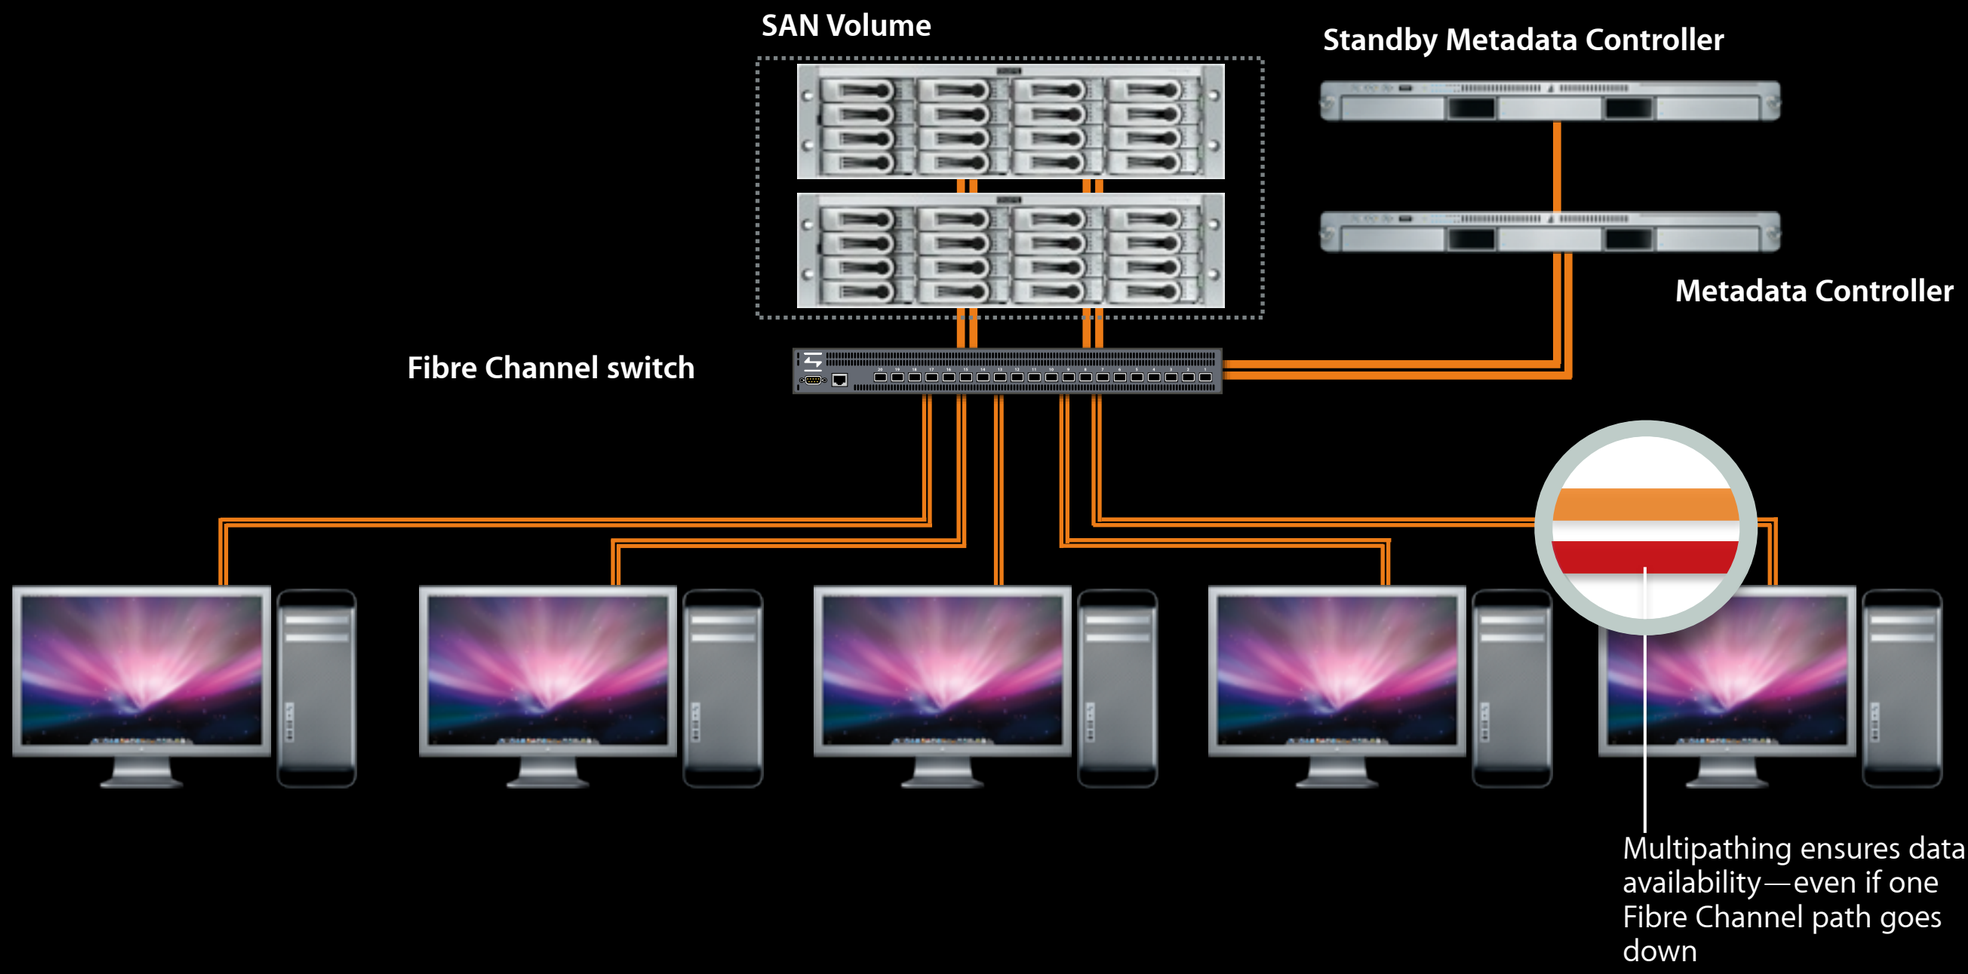

Xsan 2.xx

Server Setup Integration

Xsan Setup Assistant

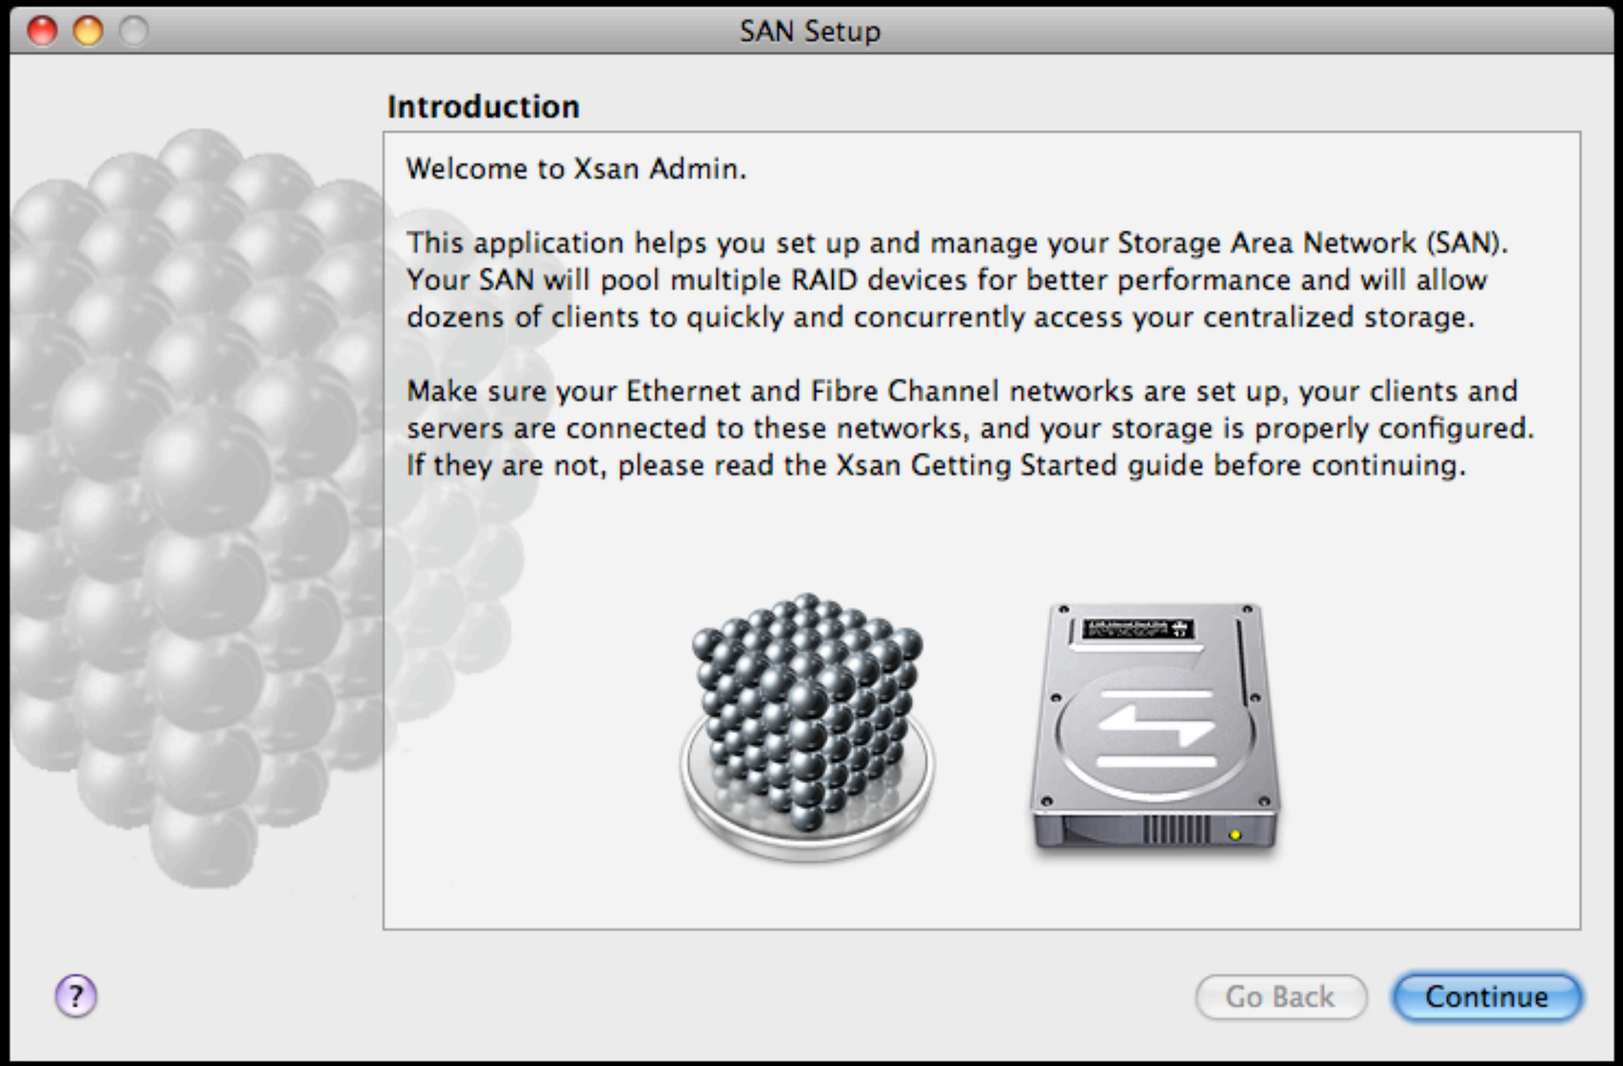

Xsan Setup Assistant

- Volume configuration settings
- LUN labeling
- Storage pool creation
- Computer setup
- Volume mounting
- Directory service configuration

Volume Configuration Settings

- High definition
- Standard definition video
- 2K and 4K video
- General file server
- Home folder server
- Mail cluster
- Podcast Producer cluster
- iCal server cluster
- Final Cut Server

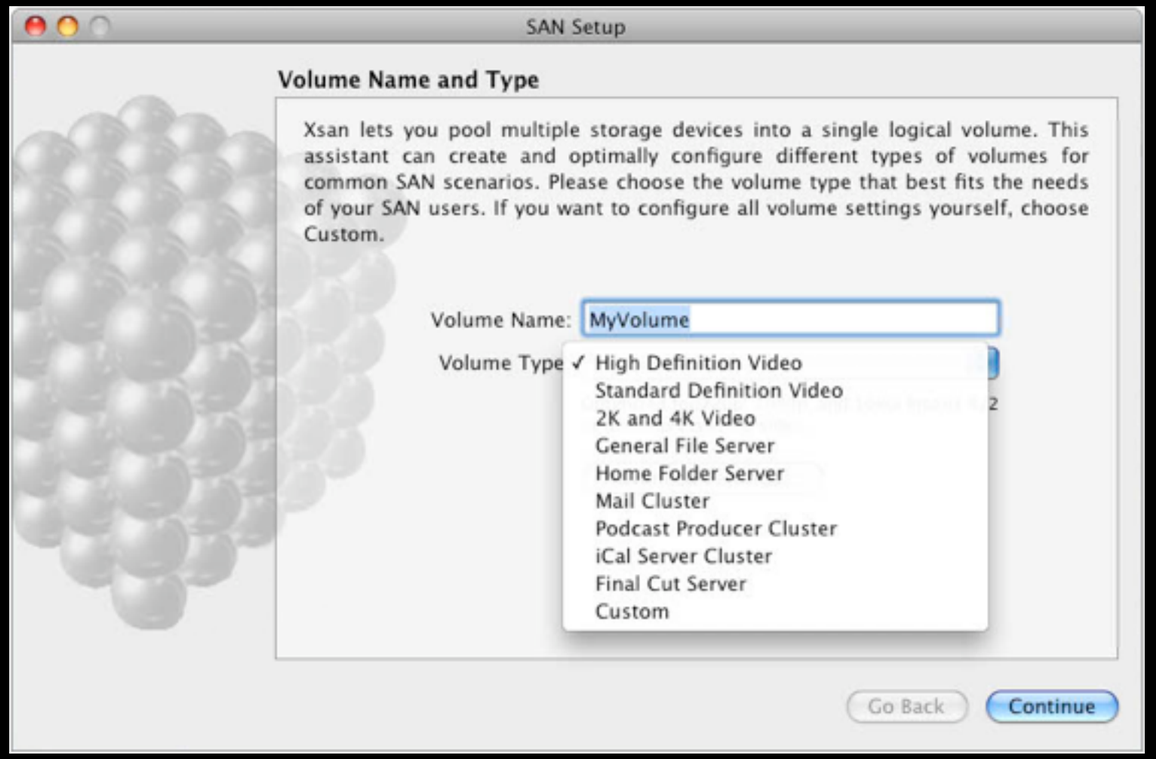

Volume - Allocation Settings

- Block Size
- Allocation Strategy
- ACLs
- Allocation Settings
- Cache Settings

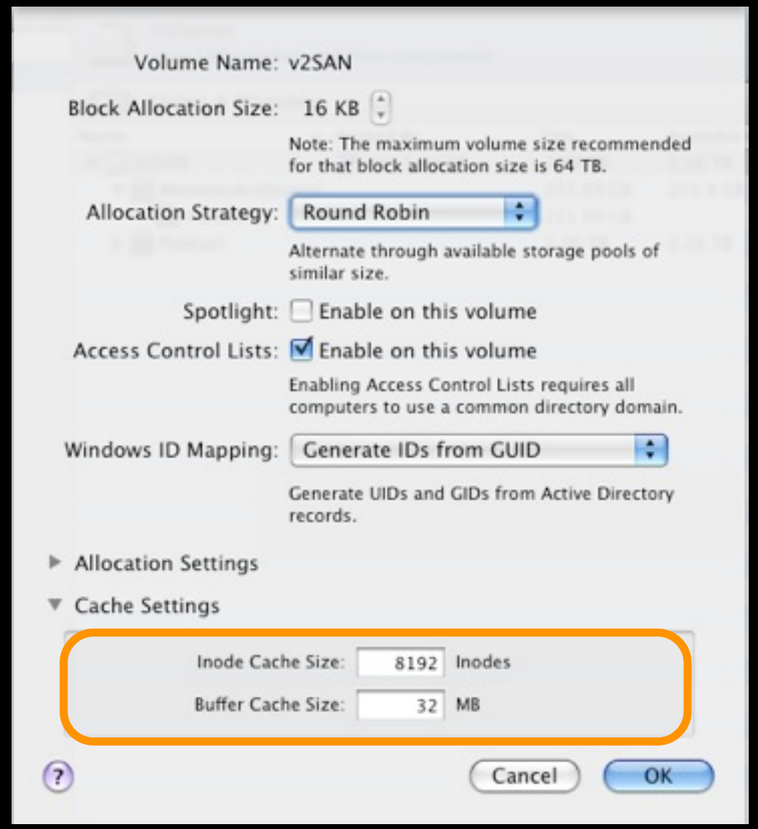

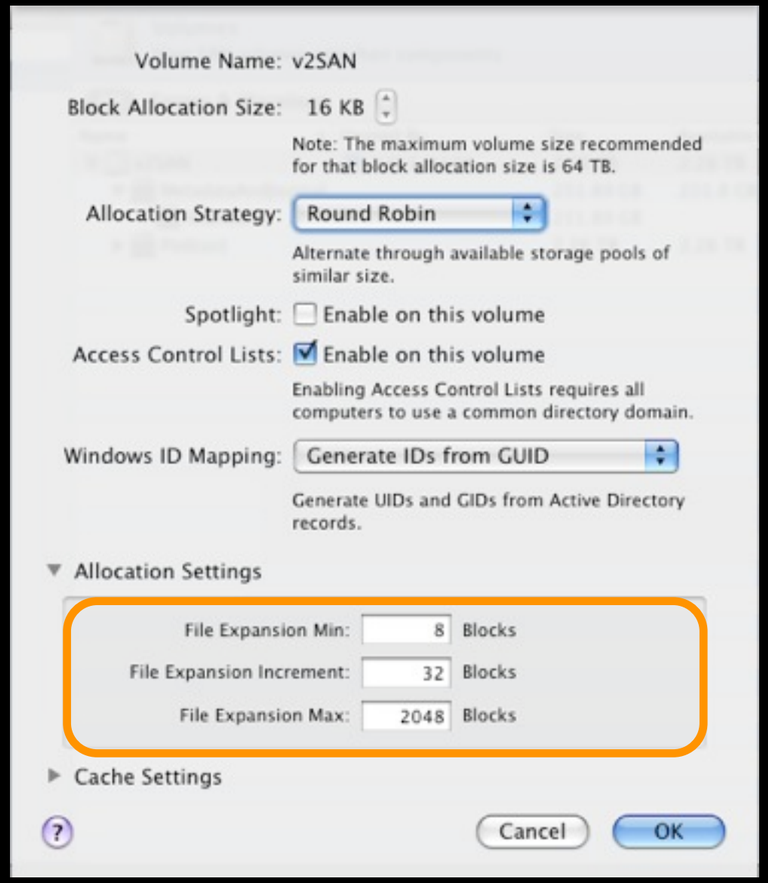

Xsan Admin

- SAN Overview
- Inspector
- LUN identifier
- Volume Expansion
- Per-volume failover priorities
- New file management interface
- Connect using SSH and Screen Sharing

Easy Setup and Remote Administration

- Easy configuration
- Real-time monitoring
- Event notifications
- Usage quotas
- Directory integration
- Flexible file permissions

Performance Optimized

- Volume Configuration Settings
- Data placement and storage pools
- Fibre Channel multipathing
- Multi-port 4Gb Fibre Channel

MultiSAN

- Host volumes on separate servers
- Access multiple volumes simultaneously
- Per-volume failover assignments
- Copy data between SAN volumes over high-speed Fibre Channel

MultiSAN

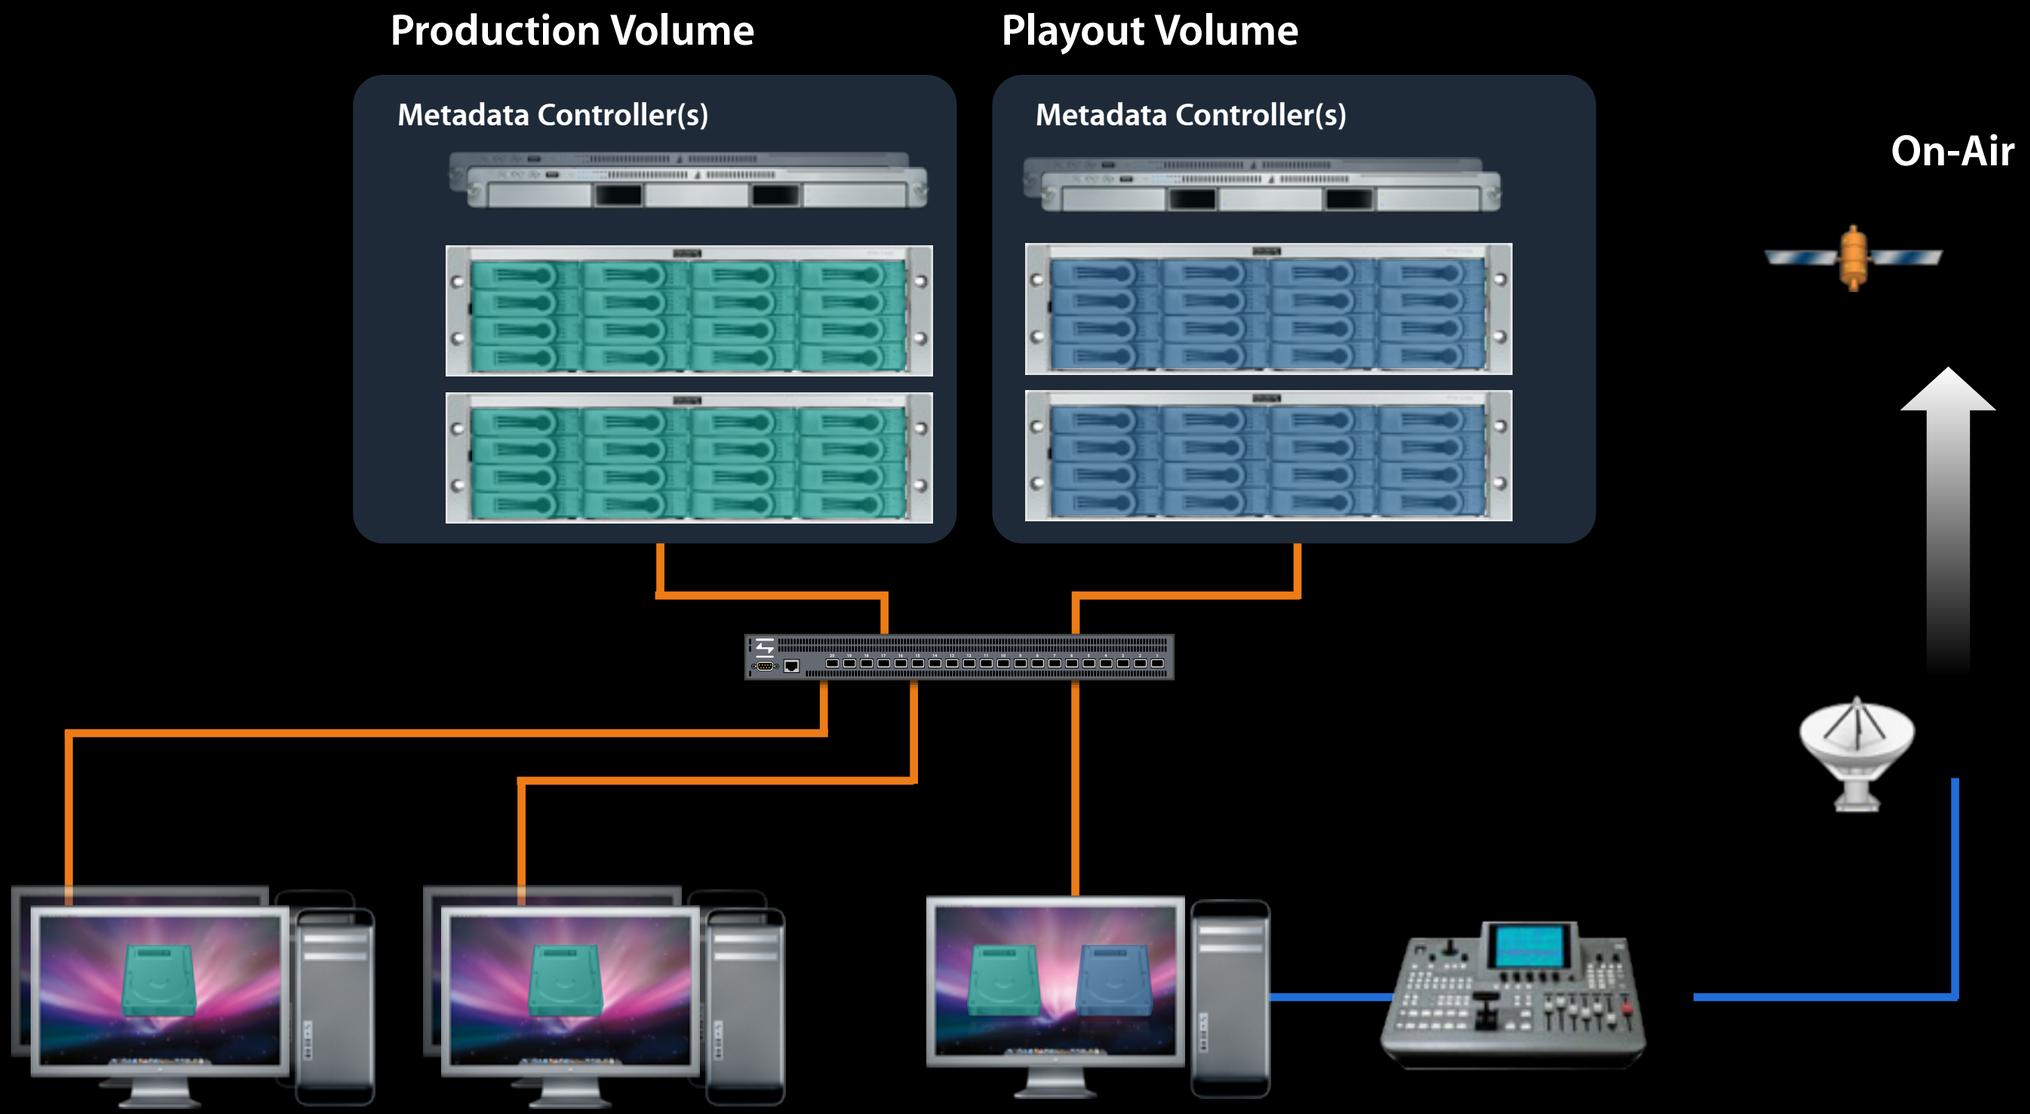
Production Volume
Metadata Controller(s)
Playout Volume
Metadata Controller(s)
On-Air

Per-volume Failover Assignments

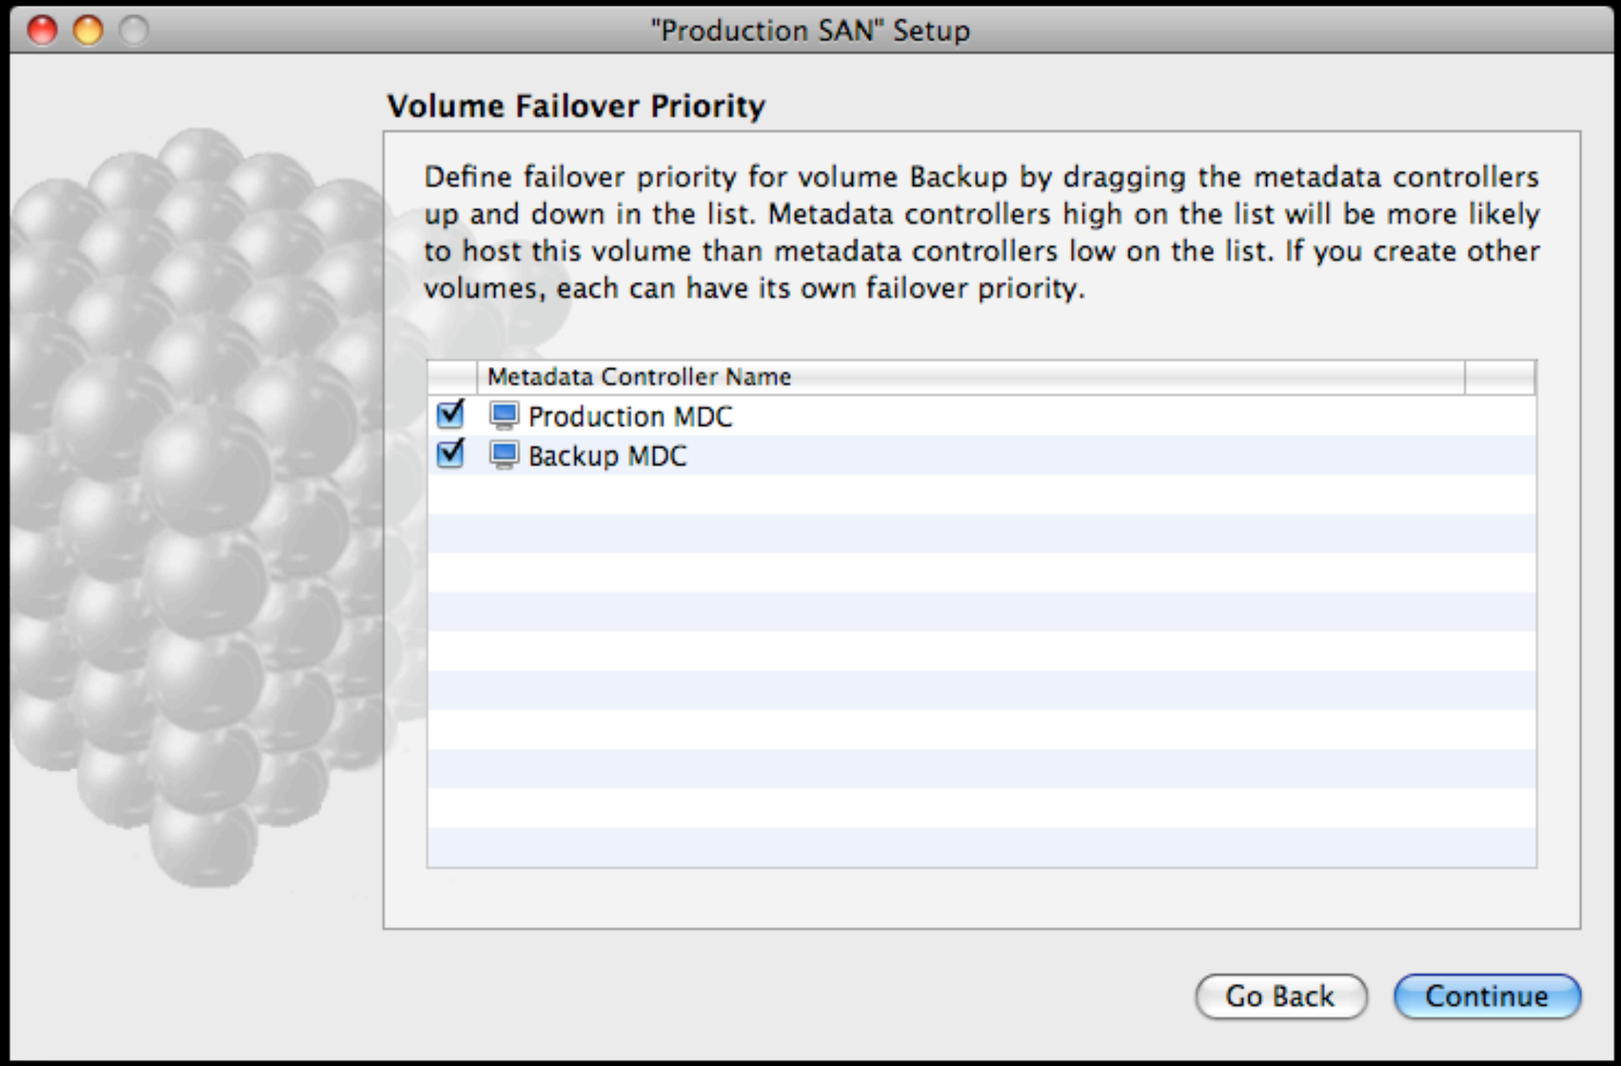

Planning:

Xsan Discovery & Design

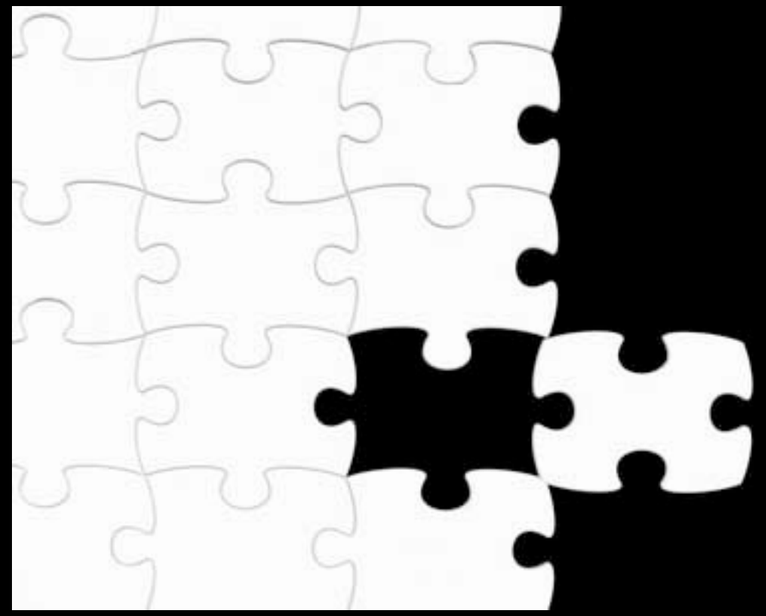

Xsan Design...

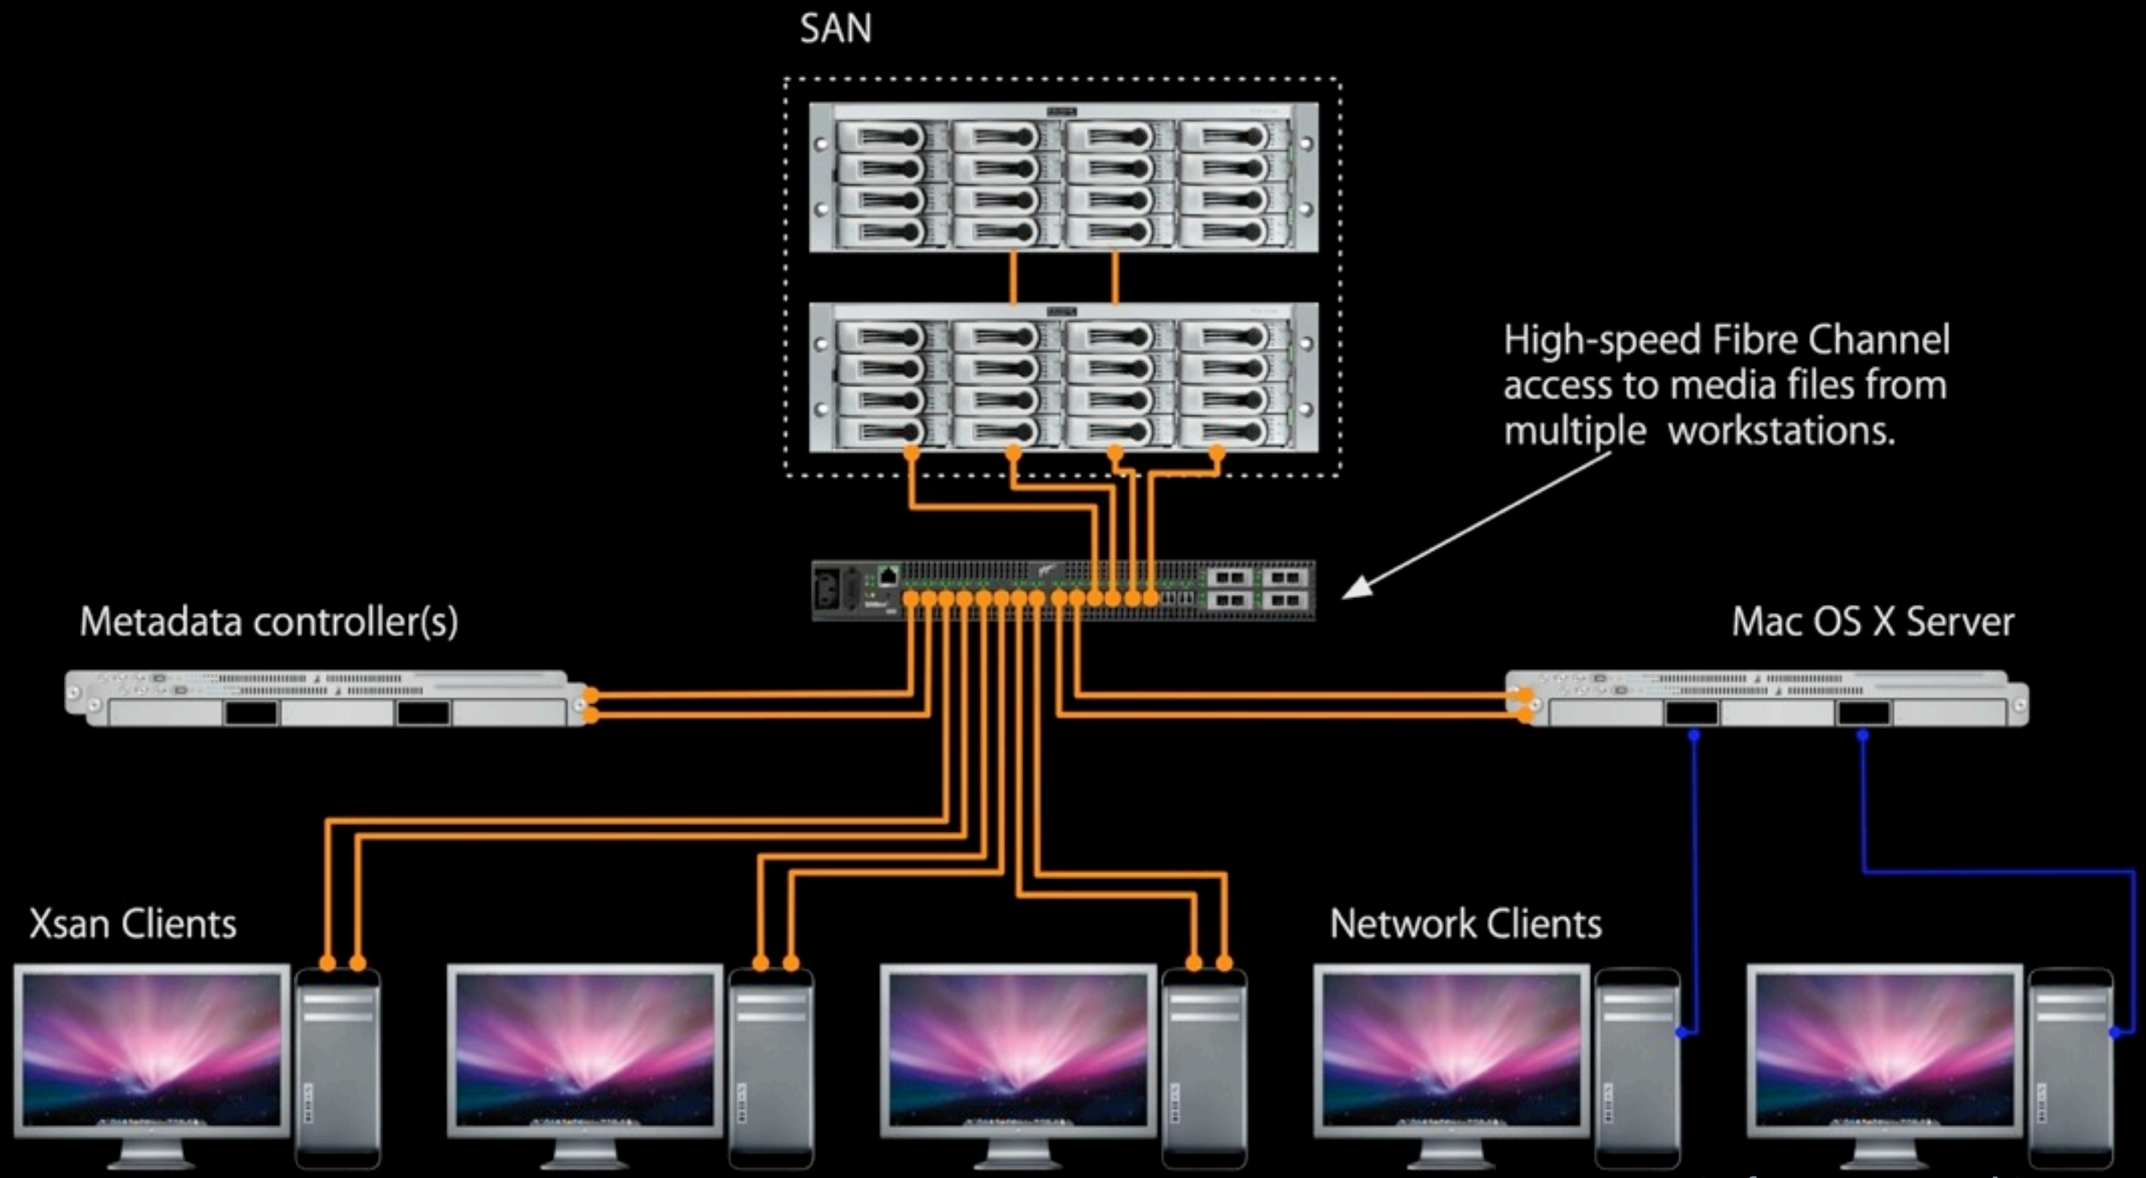
SAN
High-speed Fibre Channel access to media files from multiple workstations.
Metadata controller(s)
Mac OS X Server
Xsan Clients
Network Clients

Choosing Xsan Components

- Supported computers
- Supported storage devices
- Ethernet network
- Fibre Channel fabric, adapters, and switches
- Directory services (Optional?)
- Power and cooling

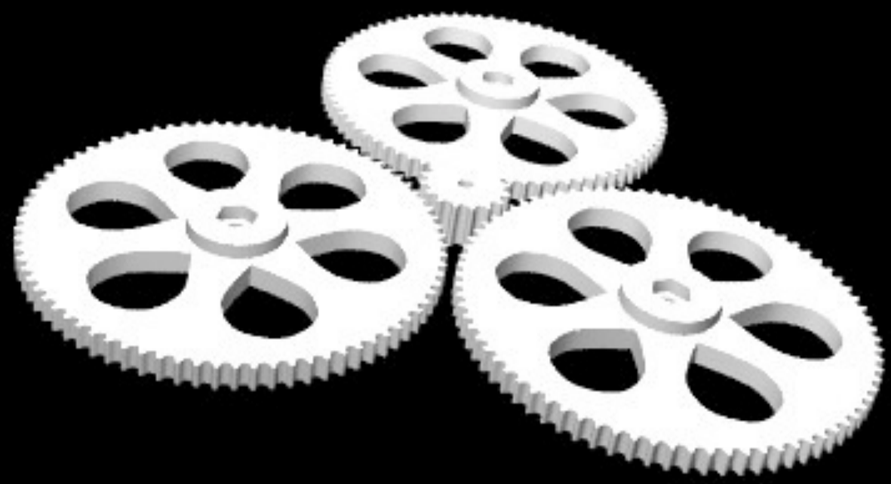

Supported Computers

- Mac Pro, Xserve or Power Mac with an Intel or G5 processor for Xsan 2 and Xserve RAID
- Mac Pro, Xserve with an Intel processor for Xsan 2 and Promise RAID
- Client systems need at least 2 GB of RAM
- Controllers need at least 2 GB of RAM for OS and 2 GB of RAM for each hosted volume

Supported Storage Devices

- Apple Xserve RAID
- Promise VTrak RAID
- Active Storage ???

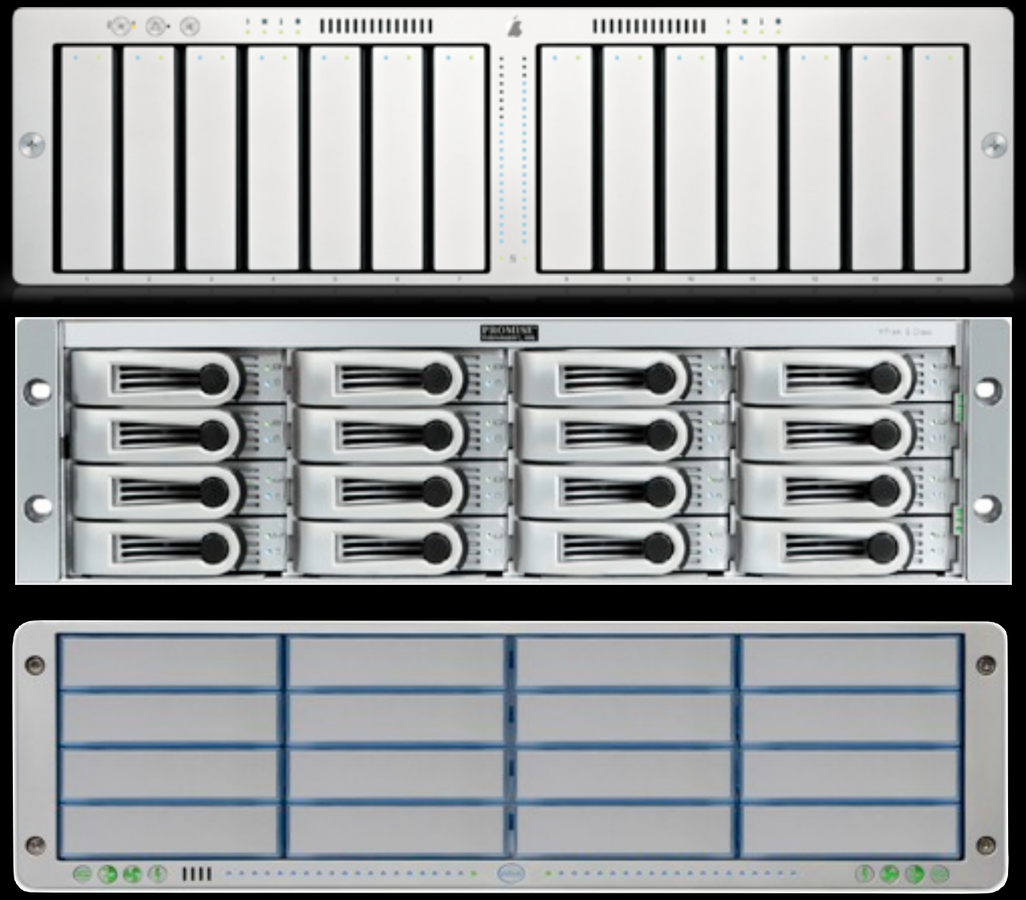

Ethernet Network...

- Xsan clients and metadata controllers use Ethernet to exchange volume metadata
- Xsan clients can use Ethernet for access to networks outside the SAN.
- Xsan metadata controllers can use Ethernet connections for remote management
- RAID systems and other components can use Ethernet for management

Ethernet Network...

- Forward and Reverse DNS for Public and Metadata networks are required for Controllers

Fibre Channel...

- Setting up the Fibre Channel Network
- Copper or Fiber Optic Cables
- Critical Connectors
- Fibre Channel Switches (kbase-HT1769)

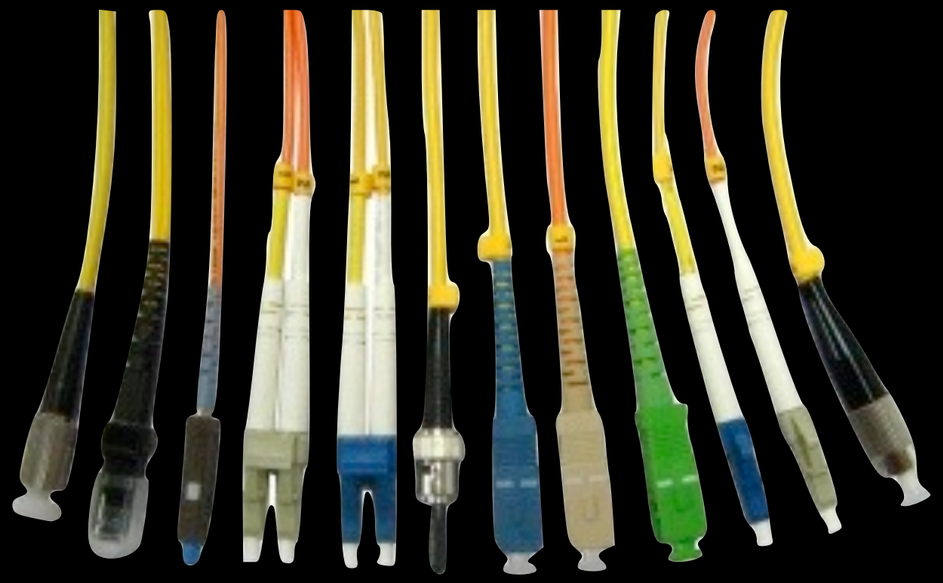

Fibre Channel Network...

- Are you building a new network or expanding an existing one?
- What will be the network's total length?
- How many devices will be networked?

Copper or Fiber Cables...

- Copper is fine for smaller installations
- There is no interference with glass
- Chosen switch may not be ratified to work with copper, many are not!
- Use 50µm OM3 Fiber - Up to 3x the bandwidth of 62.5µm
- Don’t mix 50µm and 62.5µm

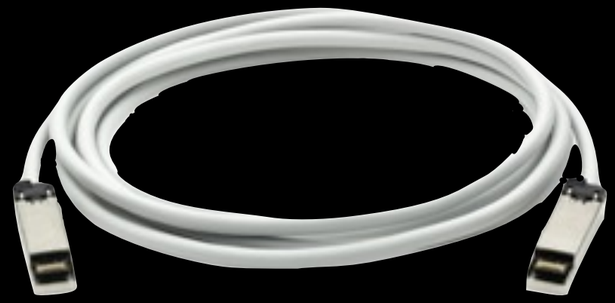

50µm or 60µm?

- Use one or the other when cabling your core Xsan components
- Longer distances don’t matter so much with mixing, but some degradation in performance
- Make sure cables are OM3 standard (laser ratified), can support up to 500m depending on application. Usual rule of thumb is maximum of 300m

50µm or 60µm?... cont

- 62.5µm is being phased out. Originally designed for LED's which modulate at maximum 633 Mbps

Critical Connectors

Duplex SC

SC Optical GBIC Transceiver

SC

LC

LC Optical SFP Transceiver

LC Optical SFF
Board-Mount Transceiver

Fibre Channel Switches

- Port Density
- Expandability
- Speed (Auto-negotiate?)
- Management
- Power (Redundant?)

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

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 8x diameter rule
- Check your termination
- Never skimp on the cabling - you will pay for it later!

Switches and Cabling

- The switch vendor you choose influences cable choices
- Only Qlogic supports Copper
- Cisco and Brocade require optical
- Brocade and Qlogic support Finisar SFP on Apple Store
- Use Cisco SFPs with Cisco switches

Directory Services...

- Allows for the management of users and groups from a central system

Power and Cooling

- 1 Ton AC unit = 1300 BTU/hr

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
Vicom VMirror Engine	460W
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	150 BTU/hr

Deployment:

Physical Installation

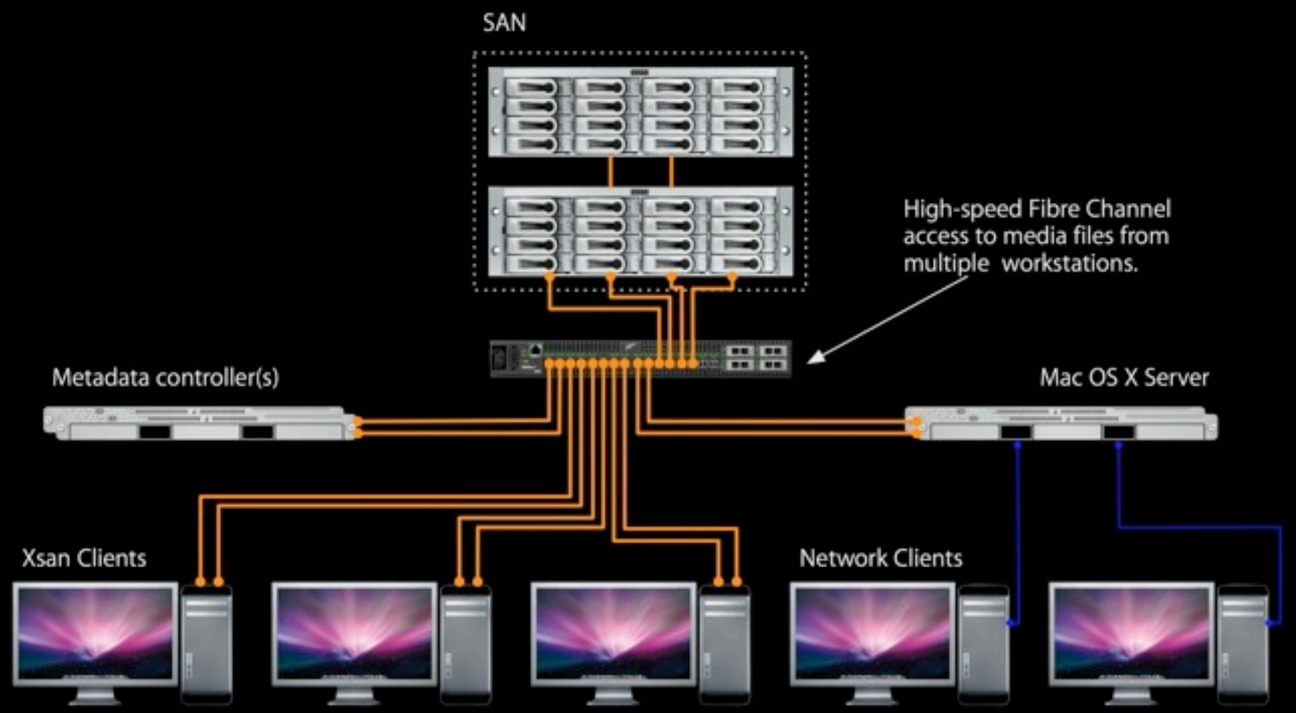

Xsan Design...

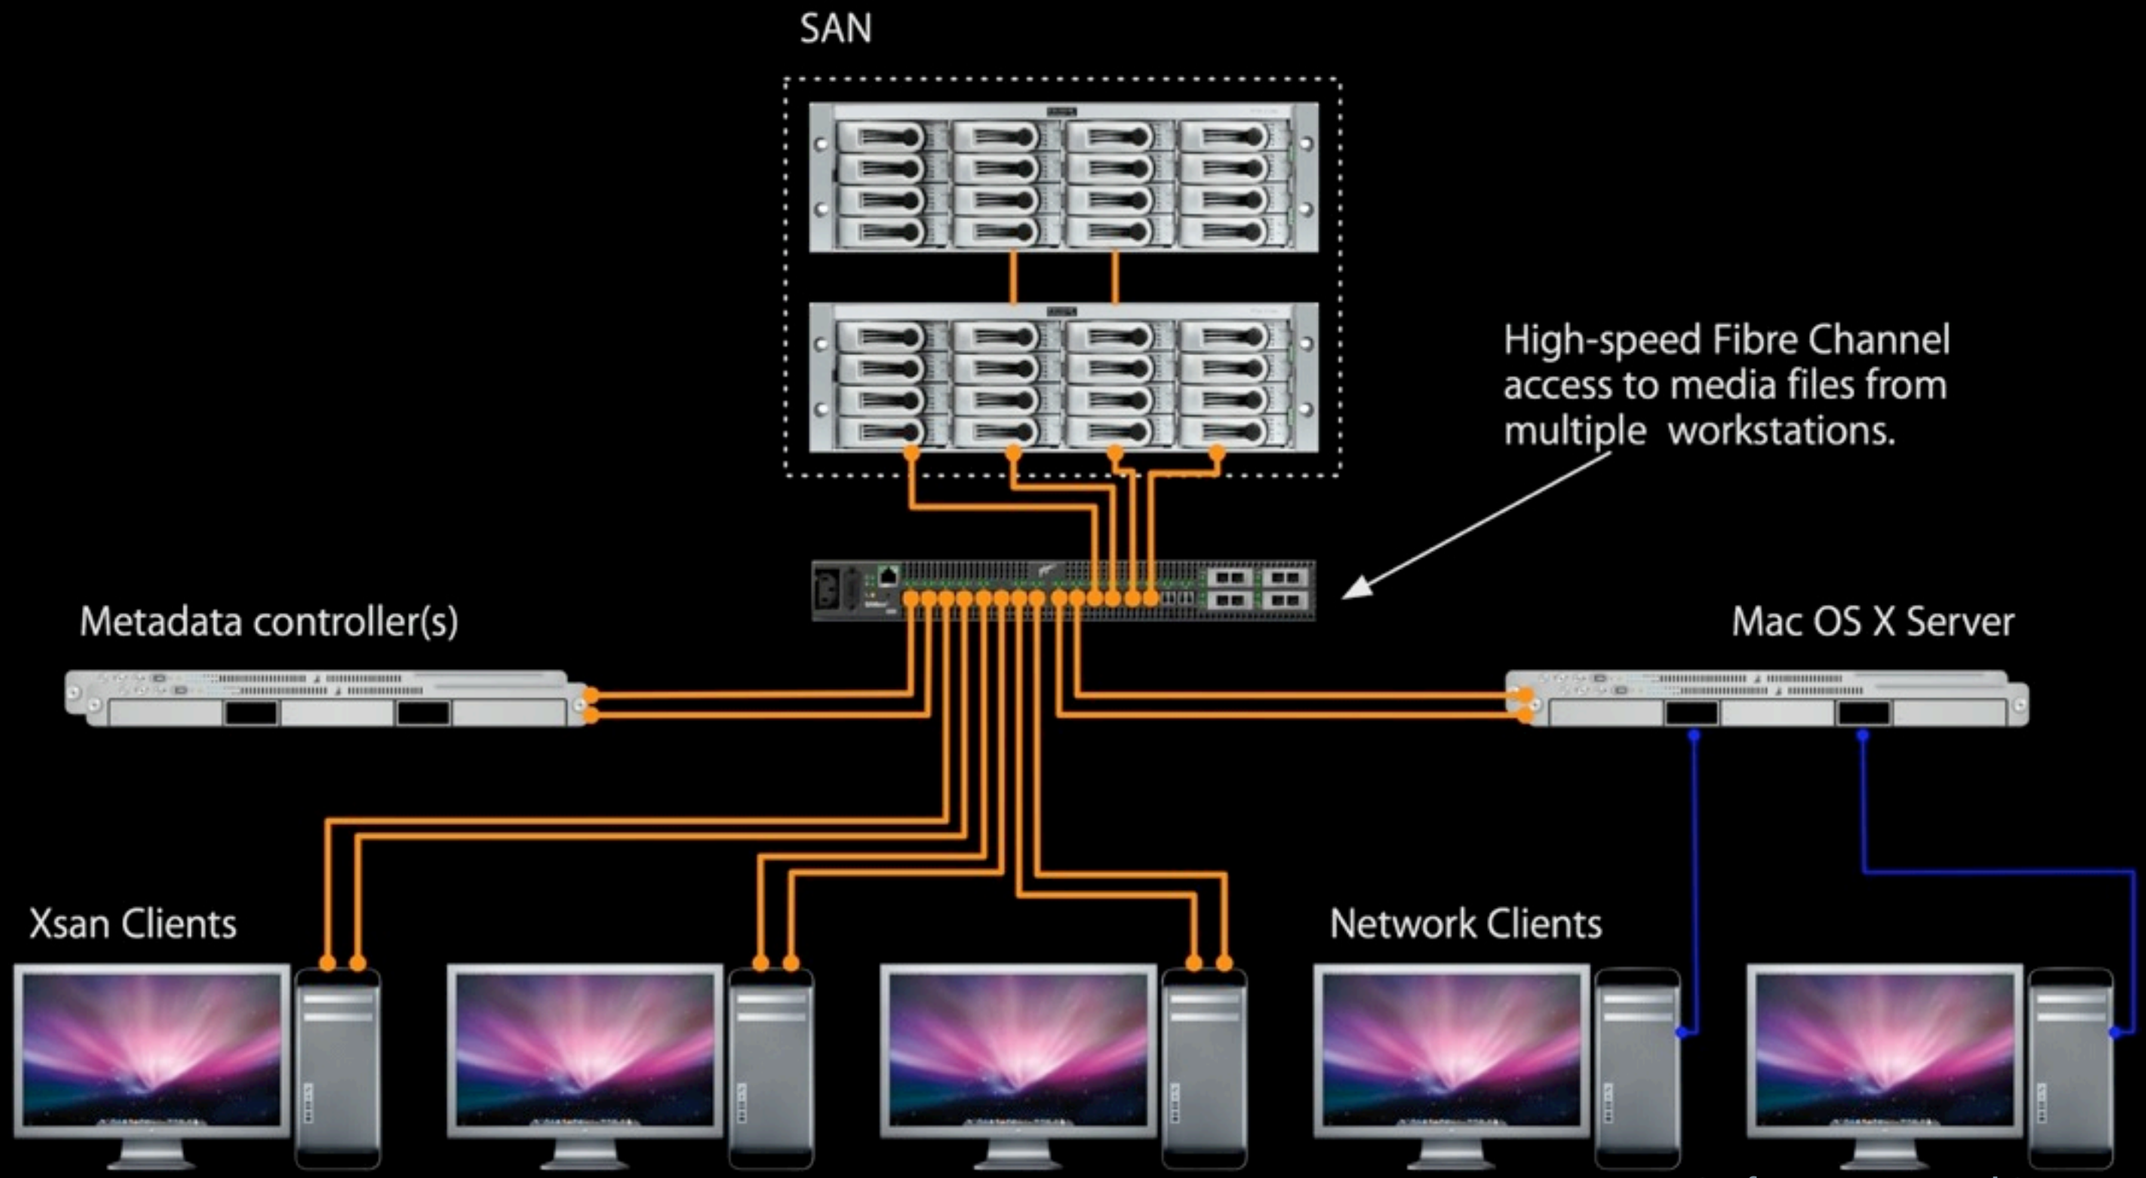
SAN
High-speed Fibre Channel access to media files from multiple workstations.
Metadata controller(s)
Mac OS X Server
Xsan Clients
Network Clients

Deploying Xsan Components

- Ethernet network
- Network services
- Directory services
- Fibre Channel
- Storage
- Xsan

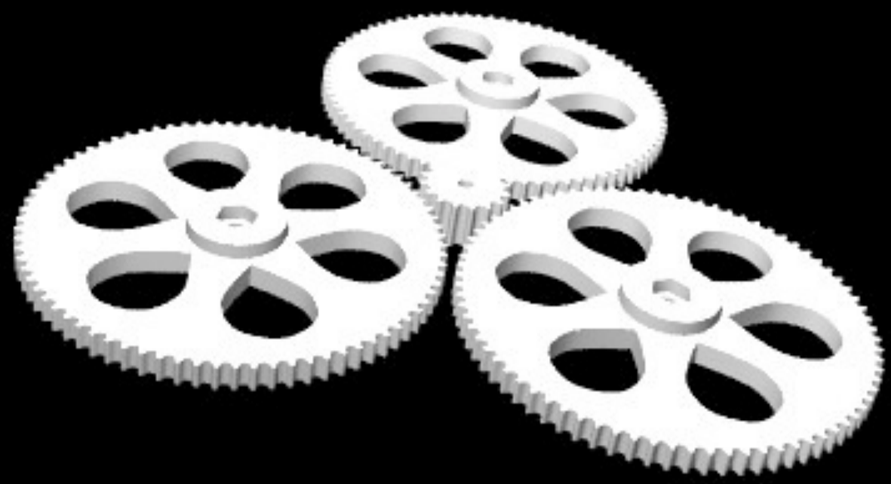

Ethernet Networks

- Place both public and private switches
- Connect port EN0 on each SAN client to public switch
- Connect port EN1 on each SAN client to private switch
- Connect Ethernet ports on Storage and Fibre switches to public switch
- Connect all non-SAN clients to public switch
- DO NOT connect private to public

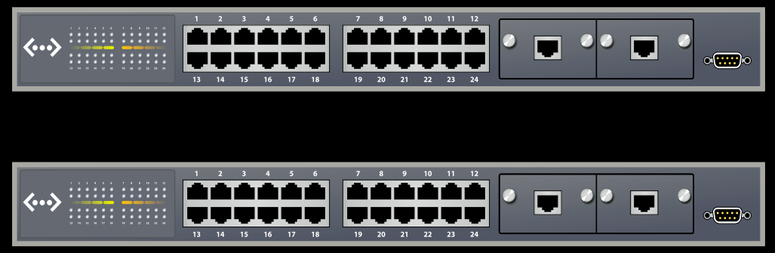

Ethernet Interface Configs

- Disable unused NICs on client and servers
- No need to switch off IPV6 anymore?

Metadata Ethernet Interface

- No router - even though you get an error
- No DNS - even though you have set one up
- No search domain
- Note the primary network interface, that is the default route

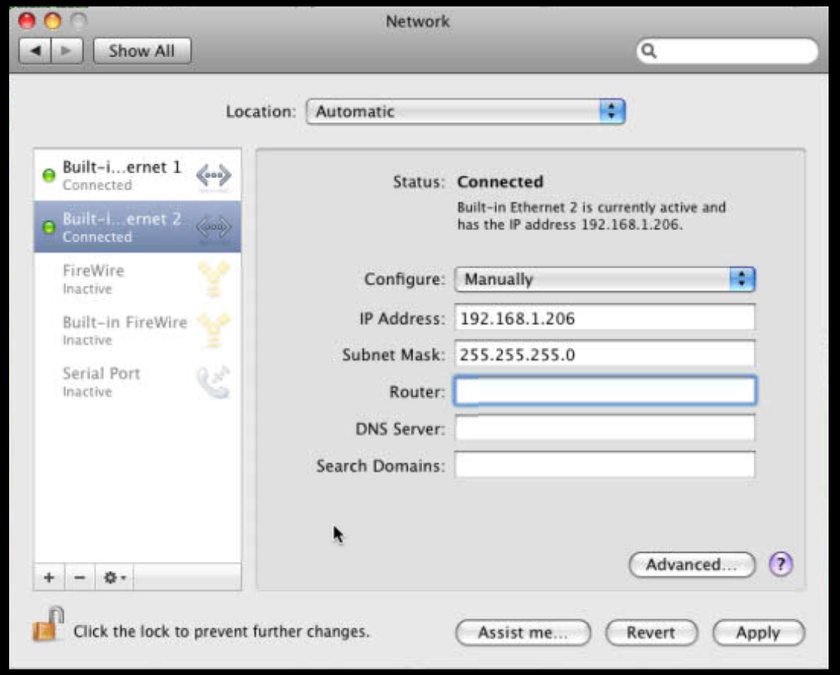

Network Services

- DNS - for both public and private IPs
- NTP - Network time server

DNS Best Practice

- Verify DNS using changeip -checkhostname
- Public and Metadata zones can be on same DNS server
- Use Fully Qualified Domain Names

DNS Configuration...

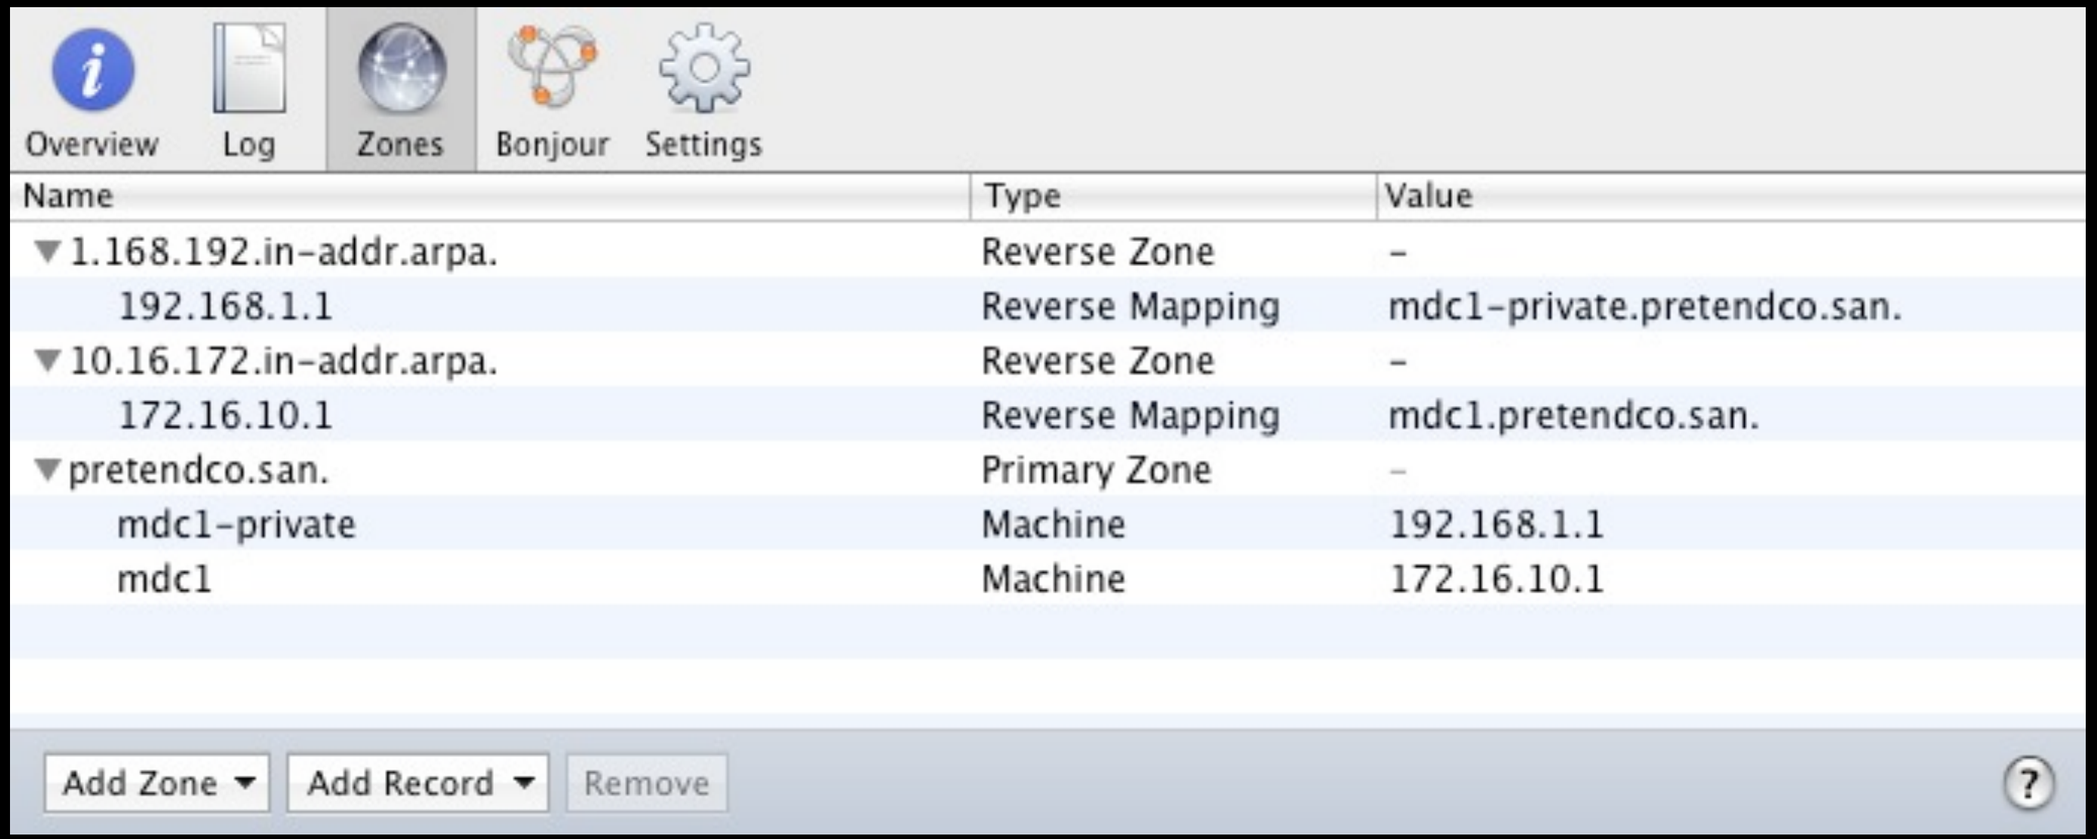

NTP Configuration...

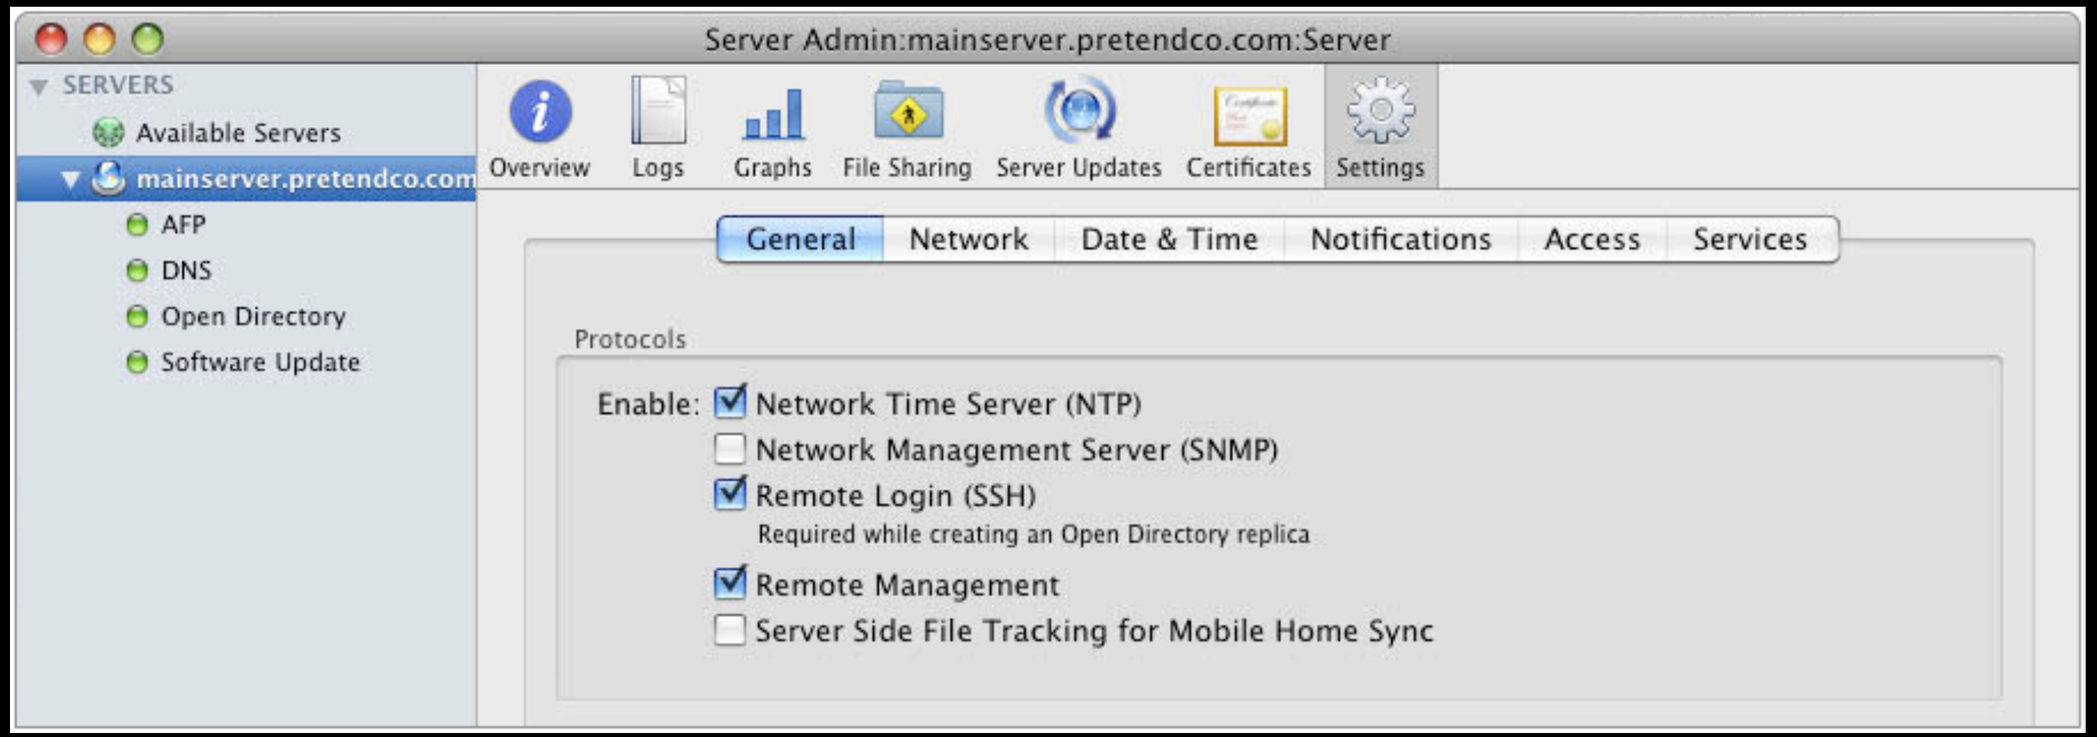

Directory Services and DNS

- Use a centralized Network Directory Service, like OD or AD
- DNS for Controllers is required, recommended for clients
- Forward and Reverse DNS entries are required
- Slow response with Xsan Admin usually indicates a DNS issue

Directory Services

- Create ODM
- Create user accounts
- Create group accounts
- Add users to groups

Fibre Channel

- Have fiber installed and certified
- Place Fibre Channel Switches
- Configure Fibre Channel Switches

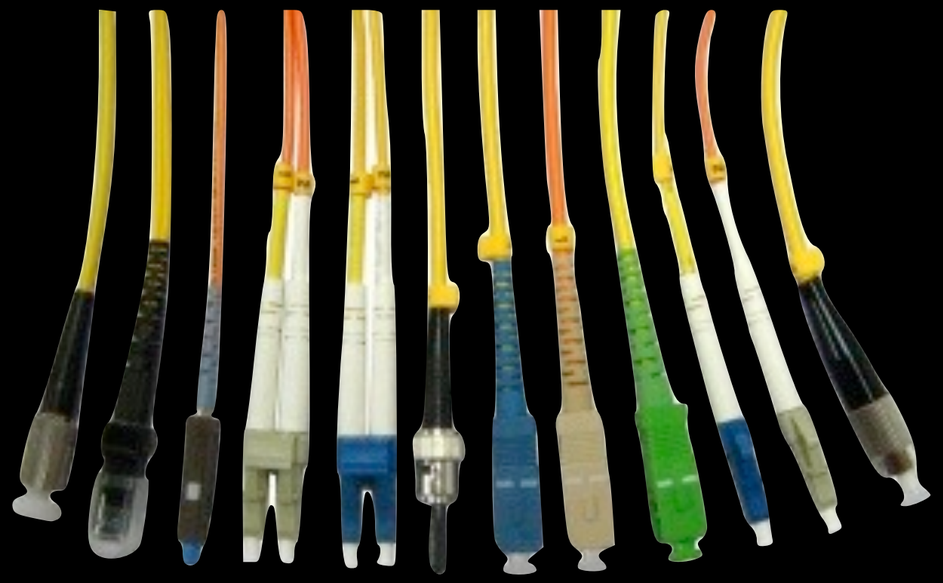

Fibre Switches - Qlogic

- Set a unique Domain ID for each switch
- Device Scan should be enabled on targets and disabled on Initiators
- Stream Guard should be enabled on Initiators and disabled on Targets
- KBase #302135

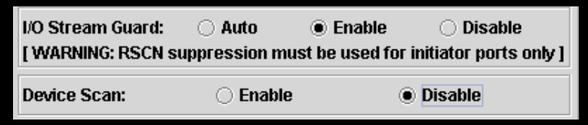

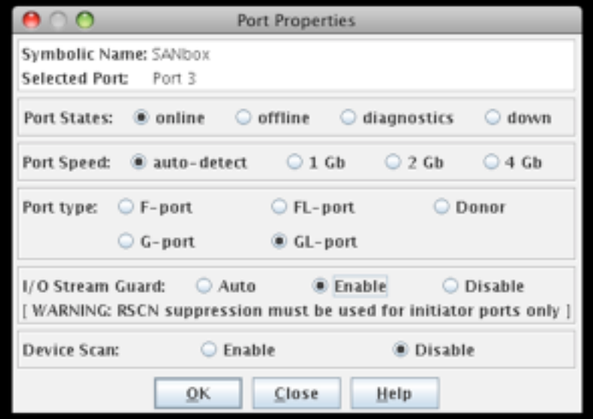

Fibre Switch Zoning

- Always use Zoning
- Set up an Orphan Zone (if there isn't one already)
- Zoning is particularly important for the Backup machine. Bus resets are common and can cause dropped frames
- Nicknames are only on the admin machine, not the switch

Storage

- Choosing RAID levels for LUNs
- Controller to LUN ratio
- Configure Storage

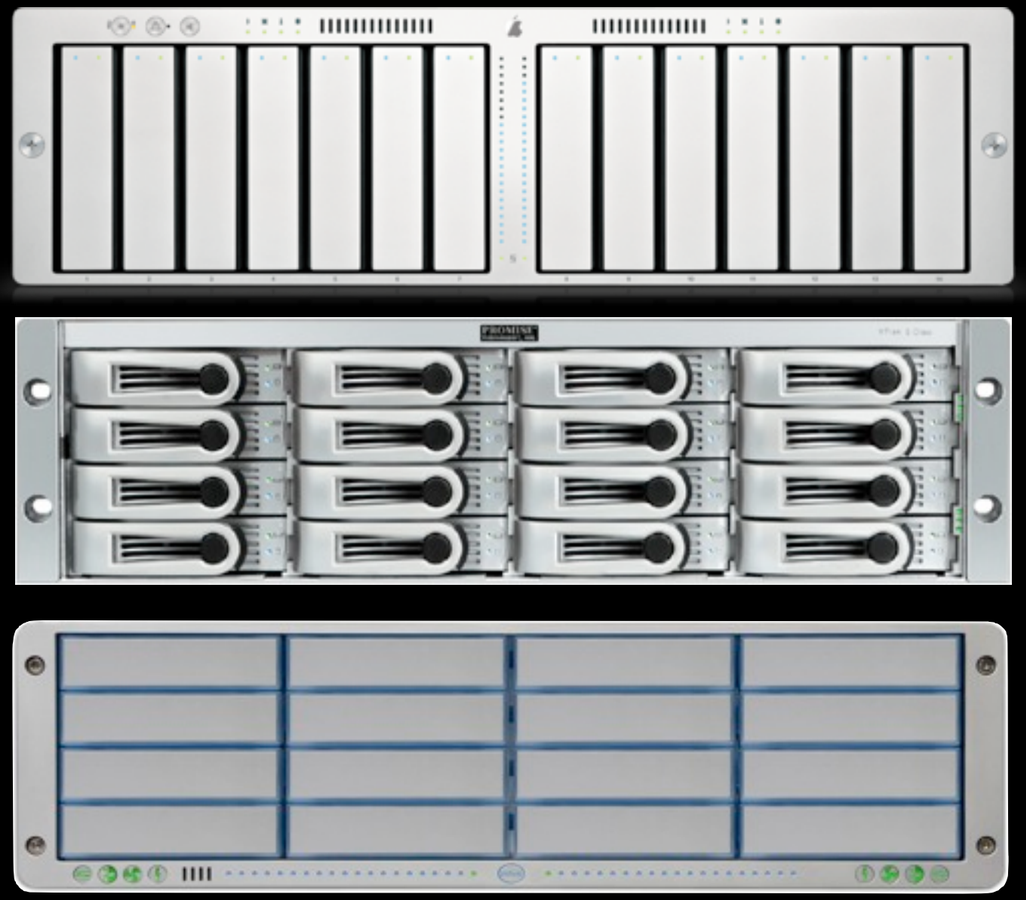

LUN RAID Levels

RAID level	Minimum Number of drives required	Storage efficiency	Read performance	Write performance	Data redundancy
RAID 0	2	Highest	Very high	Highest	No
RAID 1	2	Low	High	Medium	Yes
RAID 3	3	High to very high	Medium	Medium	Yes
RAID 5	3	High to very high	High	High	Yes
RAID 6	4	Medium to high	High	High	Yes, very

Controller to LUN Ratio

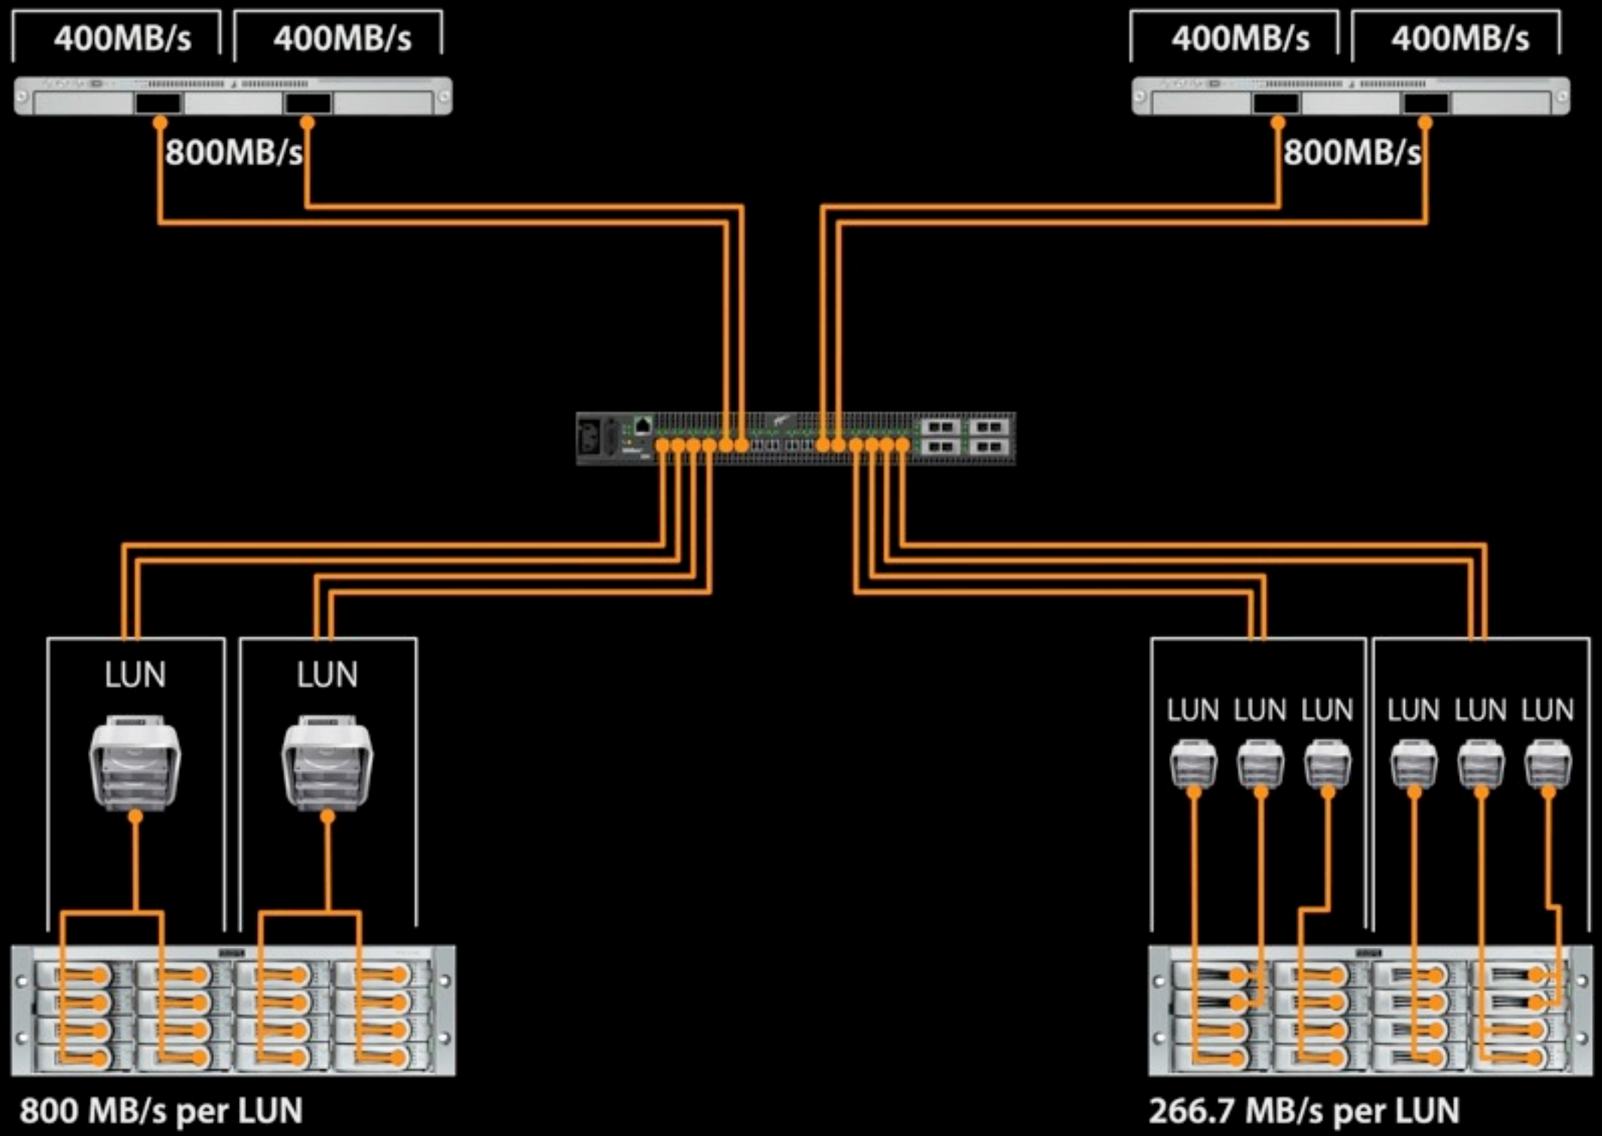

Xsan

- Configure controllers
- Configure SAN
- Configure SAN volume

Xsan Design...

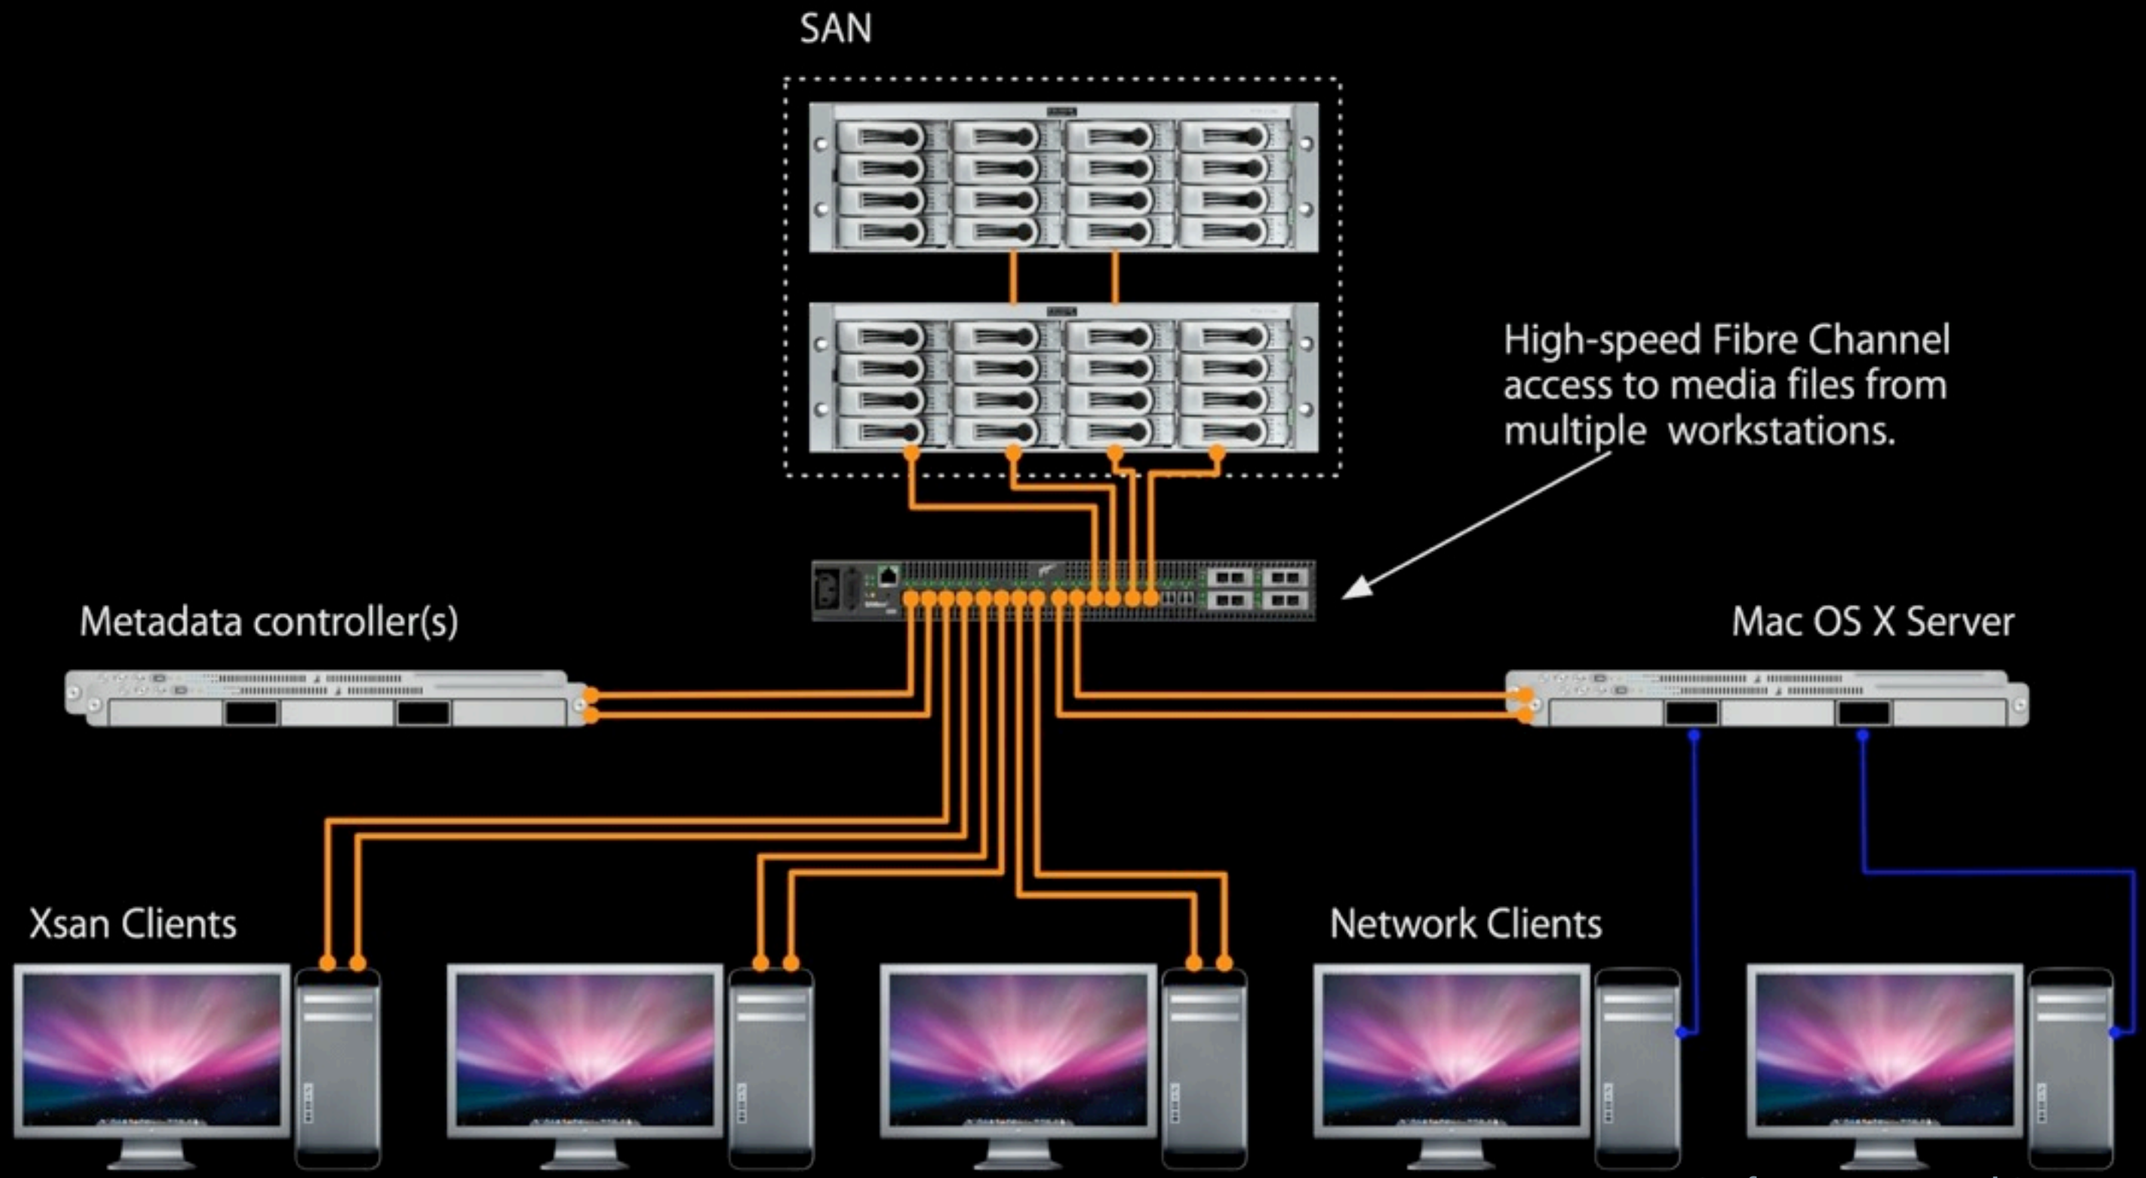
SAN
High-speed Fibre Channel access to media files from multiple workstations.
Metadata controller(s)
Mac OS X Server
Xsan Clients
Network Clients

Q&A