




Basic Concepts Quickstart

Table of Contents

- Where to Get Help ?
- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Where to Get Help ?

- 
- NetApp web site
 - <http://now.netapp.com/>
 - NetApp phone numbers:
 - +32 2 416 32 90
 - +31 235 679 601
 - Uptime ServiceDesk
 - +32 3 451 23 74
 - servicedesk@uptime.be

Note: Keep filer system ID/serial number handy !

Note: a filer has both a serial number and a system ID

serial number = linked to hardware, changes when hardware is replaced

system ID = set in software, doesn't change when hardware is replaced

Table of Contents

- **NetApp Products**
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



NetApp Products

Filer

- FAS20x0, FAS30x0, FAS60x0, ...
- FC disks, SATA disks, SAS disks
- FC & SAS disks faster (random access), SATA slower (archiving)



NearStore (Nearline Storage)

- R100, R150, R200: (S)ATA disks
- Bigger, slower disks
- Typically used as a backup solution in tapeless environments



Gateway products

- To other vendor's storage solutions: V3000, V6000

NetCache (sold)

- Web proxy product

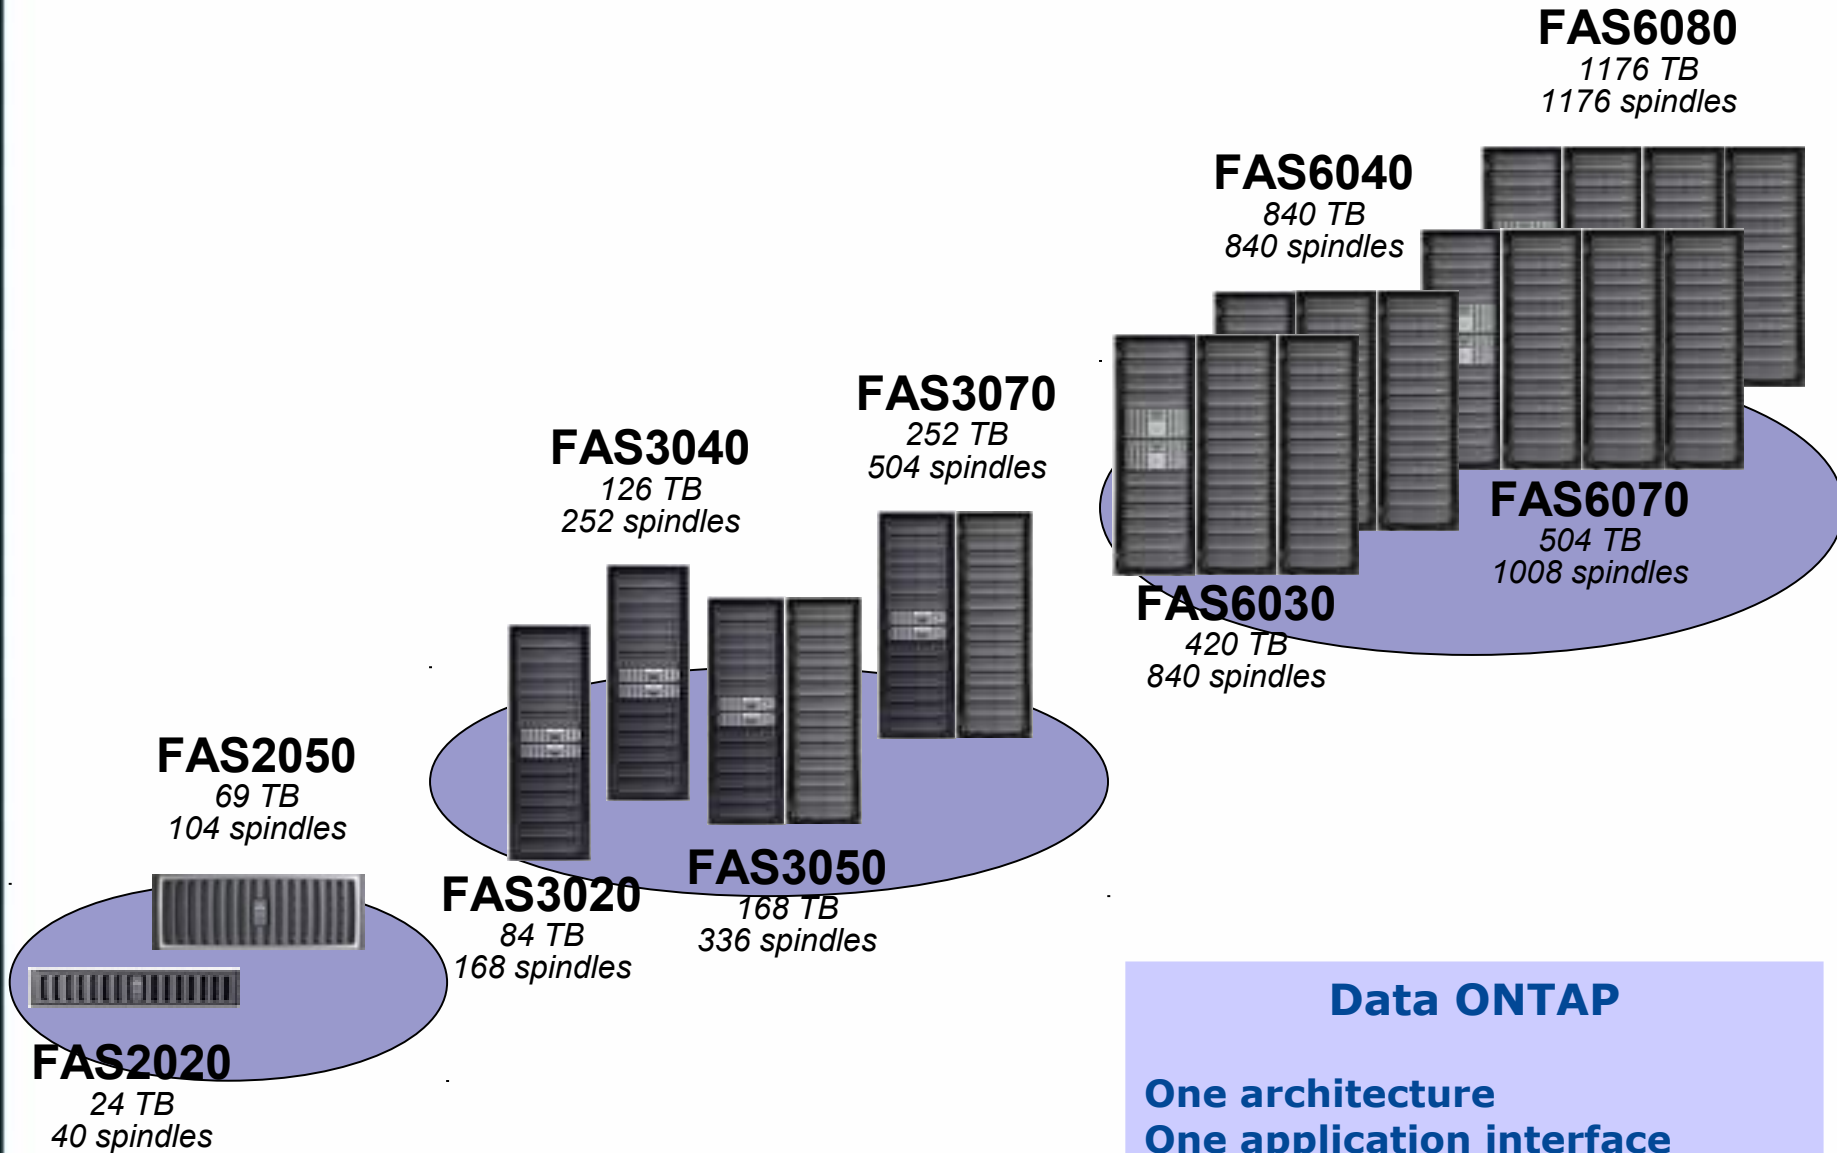


VTL (Virtual Tape Library)

StoreVault



NetApp Products (cont.)



Data ONTAP

- One architecture
- One application interface
- One management interface
- Total interoperability
- Learn one; know them all

NetApp Products (cont.)

IBM N-series

- N7000
 - N7700 (FAS6030) – 2 Gbps FC ports, 64-bit
 - N7900 (FAS6070) – 2 Gbps FC ports, 64-bit
 - N7600 (FAS6040) – 4 Gbps FC ports, 64-bit
 - N7800 (FAS6080) – 4 Gbps FC ports, 64-bit
- N5000
 - N5300 (FAS3040) – 2 Gbps FC ports, 32-bit
 - N5600 (FAS3070) – 2 Gbps FC ports, 32-bit
 - N5200 (FAS3020) – 4 Gbps FC ports, 64-bit
 - N5500 (FAS3050) – 4 Gbps FC ports, 64-bit
- N3000
 - N3300 (FAS2020) – 4 Gbps FC ports
 - N3600 (FAS2050) – 4 Gbps FC ports
- N3700 (FAS270)



NetApp Products (cont.)

Past

- FAS200 series
- FAS800 series
- FAS900 series

Present

- FAS2000 series (has recently (june 2007) replaced FAS200 series)
- FAS3000 series (replaces FAS800 & FAS900 series)
- FAS6000 series (very high-end)

Rebranding (March 2008):

Network Appliance → "NetApp"



Table of Contents

- NetApp Products
- **Storage Terminology**
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview

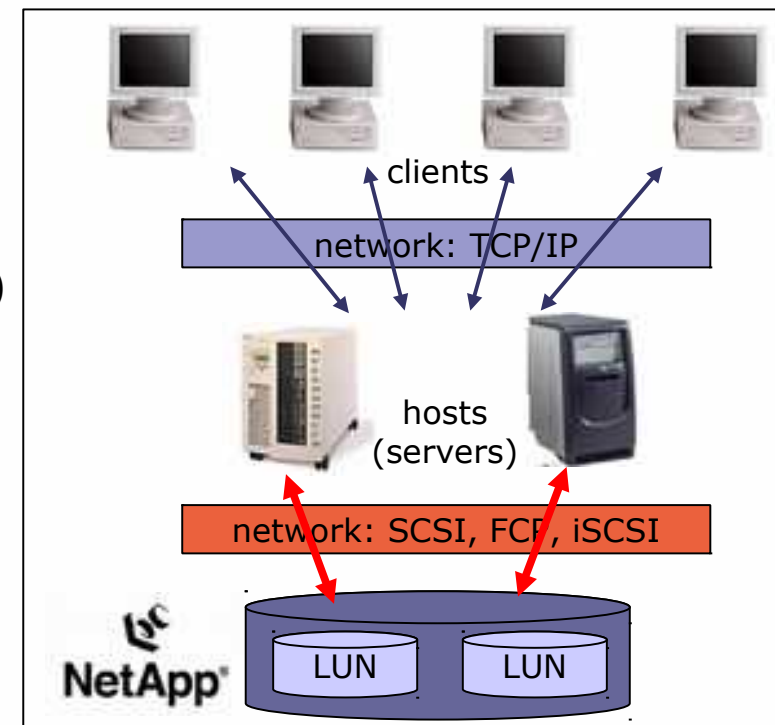
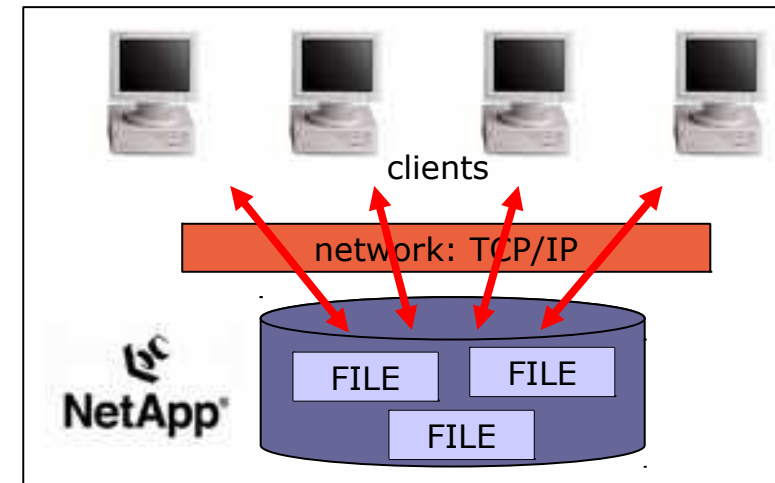


Storage Terminology

SAN vs. NAS

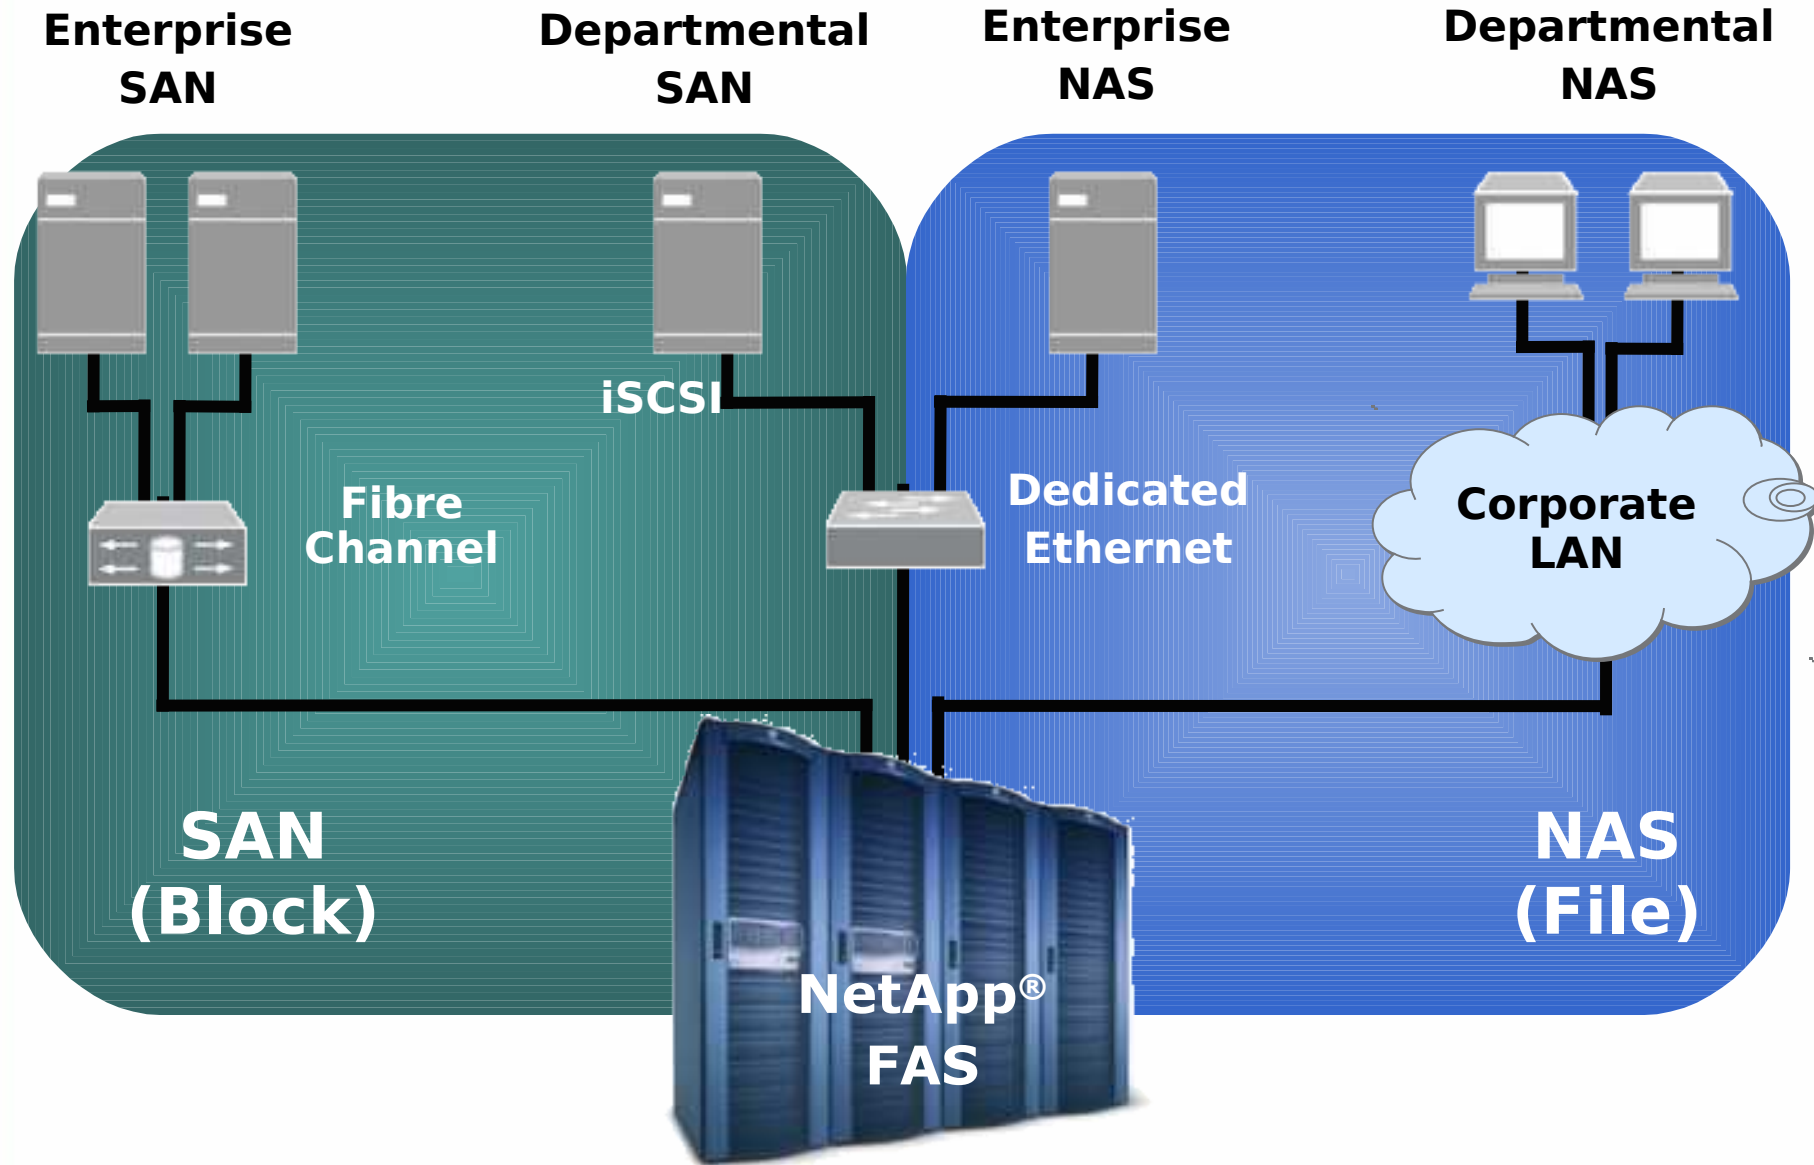
- NAS storage speaks to a FILE
 - **File** Access to data
 - NFS (Unix)
 - CIFS (Windows)
 - FTP, HTTP & WebDAV, DAFS
 - Design made to share data

- SAN storage speaks to a LUN
 - **Block** Access to data
 - SCSI
 - FCAL/FCP (encapsulated SCSI)
 - iSCSI (encapsulated SCSI)
 - Sharing of data difficult



Storage Terminology (cont.)

SAN vs. NAS (previous slide, presented differently)



Storage Terminology (cont.)

NAS Terminology

- **NAS**
 - "Network-Attached Storage"
 - File-level data storage connected to a computer network providing data access to heterogeneous network client
- **Client/server**
 - Computing architecture implemented over a computer network, allows devices to share files and resources
- **CIFS (or SMB) (Windows) and NFS (Unix)**
 - Two most commonly used NAS protocols
- **Share, export**
 - A CIFS server makes data available via *shares*, a Unix server makes data available via *exports*
- **Drive mapping, mounting**
 - CIFS clients typically map a network drive to access data stored on a server, Unix clients typically mount the remote resource



Storage Terminology (cont.)

SAN Terminology

- **SAN**
 - "Storage Area Network"
 - Device from storage manufacturer that provides centralized storage for server systems
- **LUN**
 - "Logical Unit Number"
 - A disk, presented by a SAN, to a host OS (Windows, Unix, ...) that looks like a locally attached disk to the host OS
- **Target**
 - The "machine" that offers a disk (LUN) to another machine, in other words, the SAN
- **Initiator**
 - The machine that expects to see a disk (LUN), in other words, the host OS
 - Typically, the host will only see LUNs after the appropriate initiator software has been installed, eg. FC/iSCSI drivers



Storage Terminology (cont.)

SAN Terminology (cont.)

– Fabric

- One or more fibre channel switches with target(s) and initiator(s) connected to them are referred to as a fabric
- Well-known vendors of fibre channel switches: Brocade, Cisco

(Example on next slide)

– HBA

- “Host Bus Adapter”
- Fibre channel card to connect a server or SAN to a fabric
- There are also iSCSI HBAs


– Multipathing (MPIO)

- The use of redundant storage network components responsible for transfer of data between the server and storage. These components include cabling, adapters and switches and the software that enables this



Storage Terminology (cont.)

Two Fabrics



Storage Terminology (cont.)

SAN Terminology (cont.)

– Zoning

- The partitioning of a fabric (or storage area network) into smaller subsets to restrict interference, add security, and to simplify management. If a SAN hosts several hundred disk drives, each system connected to the SAN doesn't need to see all of them
- Compare this to VLANs in networking

– Boot from SAN

- Put boot disk(s) of server on SAN
- Requires special HBA features
- FCP and iSCSI

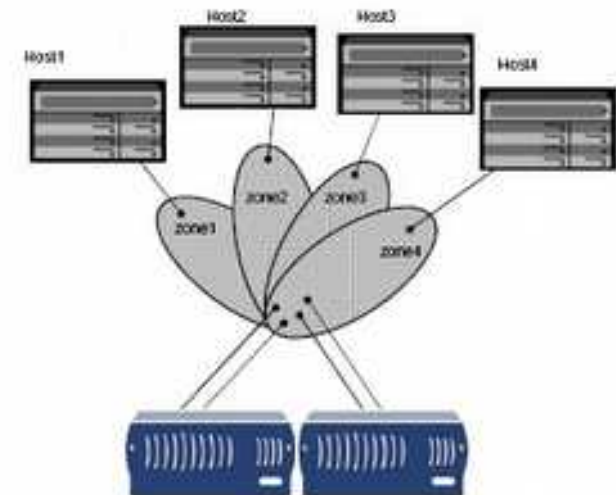


Table of Contents

- NetApp Products
- Storage Terminology
- **NetApp Terminology**
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



NetApp Terminology

Some NetApp-specific Terms ...

– Data ONTAP

= Operating system on Network Appliance filers and nearstores, borrows ideas from Unix (BSD)

eg. /etc/ directory on vol0

eg. inodes

- Same OS on every Filer/NearStore model, different architectures supported (Intel x86, AMD, 64-bit, MIPS, ... depending on hardware model)
- All features are activated through licenses, eg. clustering
- Recent version: 7.2.5(.1), 7.0.7
- “Release early, release often”-philosophy
 - **RC** = release candidate
 - **GA** = general availability, supported but not fully tested
 - **GD** = general deployment, factory-installed on machines
 - Older releases still supported: eg. 6.5.6
- ONTAP 7G
- ONTAP GX – spinfos, Spinserver, Spinnaker



NetApp Terminology (cont.)

Some NetApp-specific Terms ... (cont.)

– Head/filer




- Contains motherboard, network connections, fiber connections, console connections, connections to disks, ...

– (Disk) Shelf

- Contain disks
 - DS14(mk II): 14 disk per shelf
 - FAS2000: 12, 20, 25 disks per shelf

– **Note: FAS200 and FAS2000 series**

- Motherboard and first disk shelf are integrated (disk shelf can be turned into filer and vice versa)



NetApp Terminology (cont.)

What Can Be Upgraded ?

- Disk firmwares
 - Non-disruptively
- Shelf firmwares
 - Non-disruptively for FCAL shelves
 - Disruptively for (S)ATA shelves
- Motherboard firmware and diagnostics
 - Requires halt/reboot
- Data ONTAP
 - Requires reboot
- RLM/BMC (Remote LAN Module/Baseboard Management Controller)
 - No reboot required

When to Upgrade ?

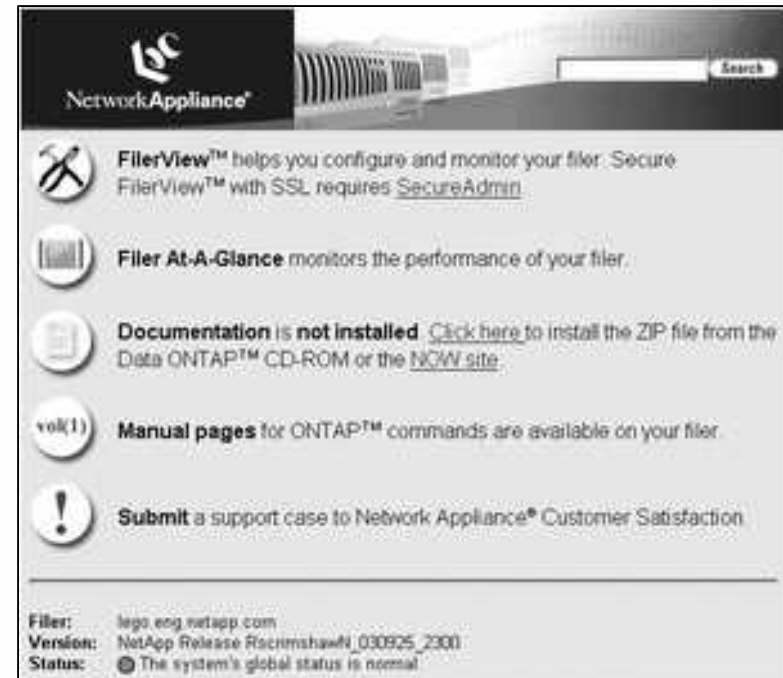
- NOW Site – Autosupport analysis
- Emails from NetApp



NetApp Terminology (cont.)

Basic Filer Management

- Filerview (http(s))
- Console cable
- Telnet
- ssh (Secure Shell)
- rsh (Remote Shell)
- Windows MMC
(Computer Management Snap-in)
- (snmp, ndmp)



Most day-to-day activities can be performed via the web interface

Command-line interface: not-so-commonly-used commands, eg "snap restore" + many more commands

2 most commonly used commands: "sysconfig" & "options"

NetApp Terminology (cont.)

Most Commonly Used Commands

man (man pages)

Browses through man(ual) pages of command documentation

sysconfig (-a, -r, -c, -t, -m)

Shows information about filer (hardware, disks, aggregates & RAID groups, ...)

options

Queries or changes values for various "registry" options

setup

Walks through initial setup questions: filer name, IP addresses, etc., but does not erase any data

cifs setup

Walks through CIFS setup questions: domain/workgroup membership etc.

sysstat -x -s 1

Prints out all-round performance statistics



NetApp Terminology (cont.)

Most Commonly Used Commands (cont.)

license

Adds/removes/prints licenses on filer

version (-b)

Prints out Data ONTAP & Diagnostics/Firmware version numbers

rdfile

Reads a text file and prints contents to console (Unix "cat")

wrfile

Reads from console and sends output to text file (Unix "Cat >")

snap (restore)

Performs snapshot operations, eg. restore from snapshot

snapvault

snapmirror

Manipulates/controls SnapMirror/SnapVault/OSSV operations from the command-line



NetApp Terminology (cont.)

Most Commonly Used Commands (cont.)

cf

Controls clustering, eg. enabling/disabling, forcing takeover & giveback

aggr

Creates/expands/destroys/manipulates aggregates, eg. change options

vol

Create/resizes/destroys/manipulates volumes, eg. change options

df

Shows free disk space (volumes, aggregates, also inodes)

qtree

Creates/manipulates qtrees (=special directories)



NetApp Terminology (cont.)

Most Commonly Used Commands (cont.)

`vif`

Creates/destroys/manipulates virtual network interfaces (eg. team interfaces for failover or load-balancing)

`ifconfig`

Sets network IP configuration (put in /etc/rc to survive reboots)

`ifstat`

Shows network interface statistics

`netdiag`

Performs basic network diagnostic testing

`ndmpd`

`ndmpcopy`

Manipulates NDMP settings, or use `ndmpcopy` to copy files via NDMP

`priv set/priv set advanced/priv set diag`

Goes into advanced/diagnostics mode



NetApp Terminology (cont.)

Autosupport

- What ?
 - Provides capability to configure filers to initiate automatic sending of email notifications to NetApp technical support and other designated addressees when specific events occur
- Why ?
 - Helps solve problems proactively
- How ?
 - Via email (smtp, preferred) or the web (http(s))
 - Configure via the web interface ...
 - ... or via the 'options' command

can only be changed
via command-line

```
filer> options autosupport.support.transport smtp
filer> options autosupport.doit testing123
```

if email subject contains the word "test", Netapp
will send a reply to all email recipients



NetApp Terminology (cont.)

Disks:

- NetApp currently uses 3 types of disks:
 - FCP (Fiber) – fast, expensive, on all models, originally in filers
 - SATA (Serial ATA) – slower, cheaper, on all models, originally on nearstores
 - SAS (Serial Attached SCSI) – fast, expensive, currently only on FAS20x0 series, poised to replace FCP in the long run
- Now:
 - Recent models can combine FC, SATA, & SAS disks
 - SATA is slower than FCP & SAS
 - FC and SATA not on same loop!

Note: “FCAL = Fiber Channel – Arbitrated Loop”

A fast, serial-based standard meant to replace the parallel SCSI standard

Primarily used to connect storage devices to servers

Software-compatible with SCSI



NetApp Terminology (cont.)

Disks, Aggregates, (Flexible) Volumes, LUNs

Aggregate = collection of disks, protected by RAID-4 or RAID-DP, can grow but cannot shrink.

Make aggregates as big as possible, with these limits:

- max 16 Tb
- don't mix disks of different types
- don't mix disks of different sizes


Create multiple aggregates if needed

Flexible Volume = logical space inside an aggregate, containing actual data, eg. files in the case of NAS

Can grow AND shrink

LUN = Logical Unit Number = logical space inside a volume, assigned to (a) server(s) to be used as a local disk.

Can grow and usually not shrink



NetApp Terminology (cont.)

WAFL

- = Write Anywhere Filesystem Layout, file system on NetApp filers and nearstores
- Unix-based, hence terms like "inodes", but allows NTFS-permissions (NTFS-security style)
- Formatting disks ? No: zeroing disks

Aggregate

- Logical group of disks, consists of parity disks and data disks
- Can be expanded on-the-fly, but **cannot shrink** in size!

Flexible Volume (Flexvol)

- Aggregate can contain multiple volumes
- Contain actual data (files and LUNs)
- Can grow and shrink



NetApp Terminology (cont.)

Filer Disks

- Data disks
- Spare disks
- Parity disks
- Double parity disks
- (Broken disks)
- (Partner disks)

When dealing with storage, ONTAP 7 will try to hide as much as possible from this from the storage administrator


On our filer, we will create an *aggregate* of many data disks + parity. This aggregate can be expanded. On this aggregate we will create *flexible volumes* that can grow and shrink in size, and that will contain actual data (files and LUNs)




NetApp Terminology (cont.)

Traditional RAID Levels

- RAID1 (mirroring)



- RAID5

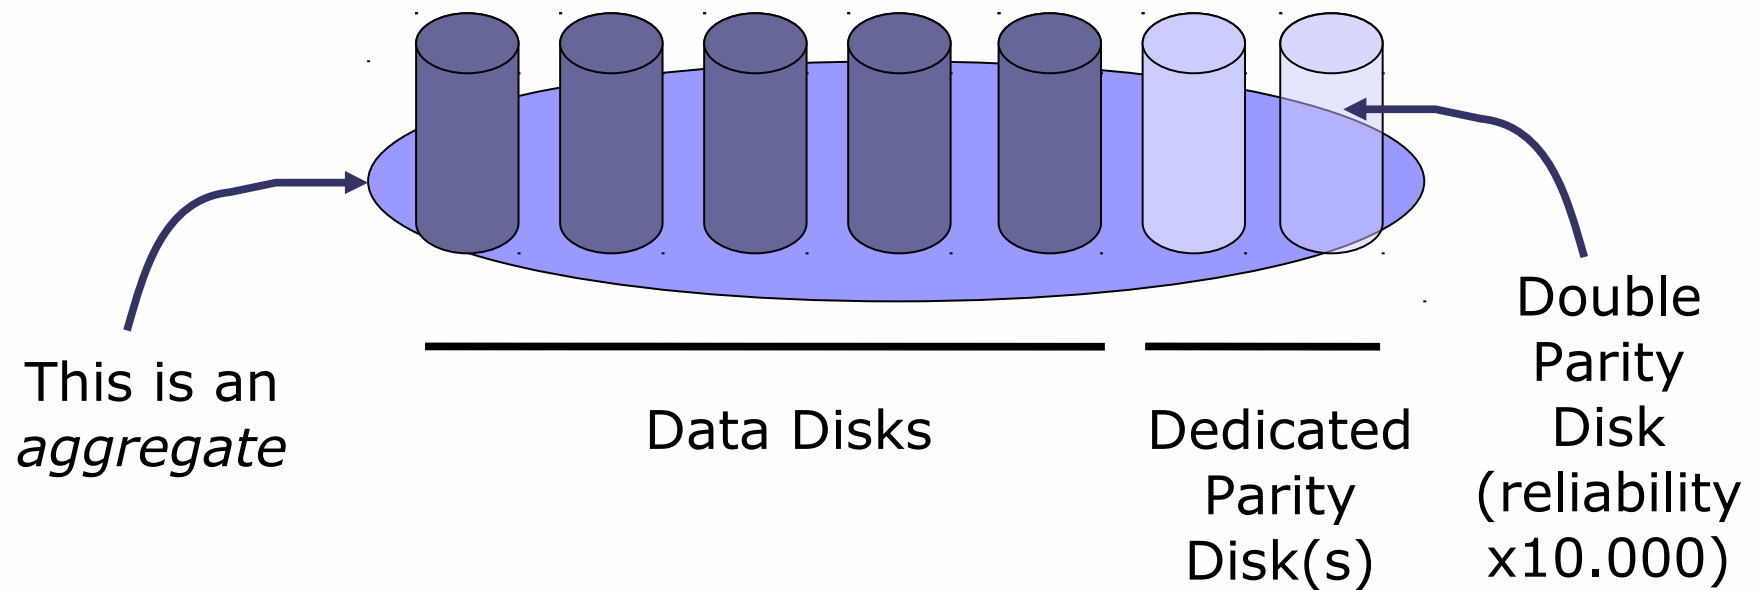


Very nice but unfortunately difficult, if not impossible to expand on-the-fly

NetApp Terminology (cont.)

RAID4 & RAID-DP

- NetApp uses RAID4 (or RAID-DP – double parity) as the only underlying RAID level (**no RAID1 or RAID5**)

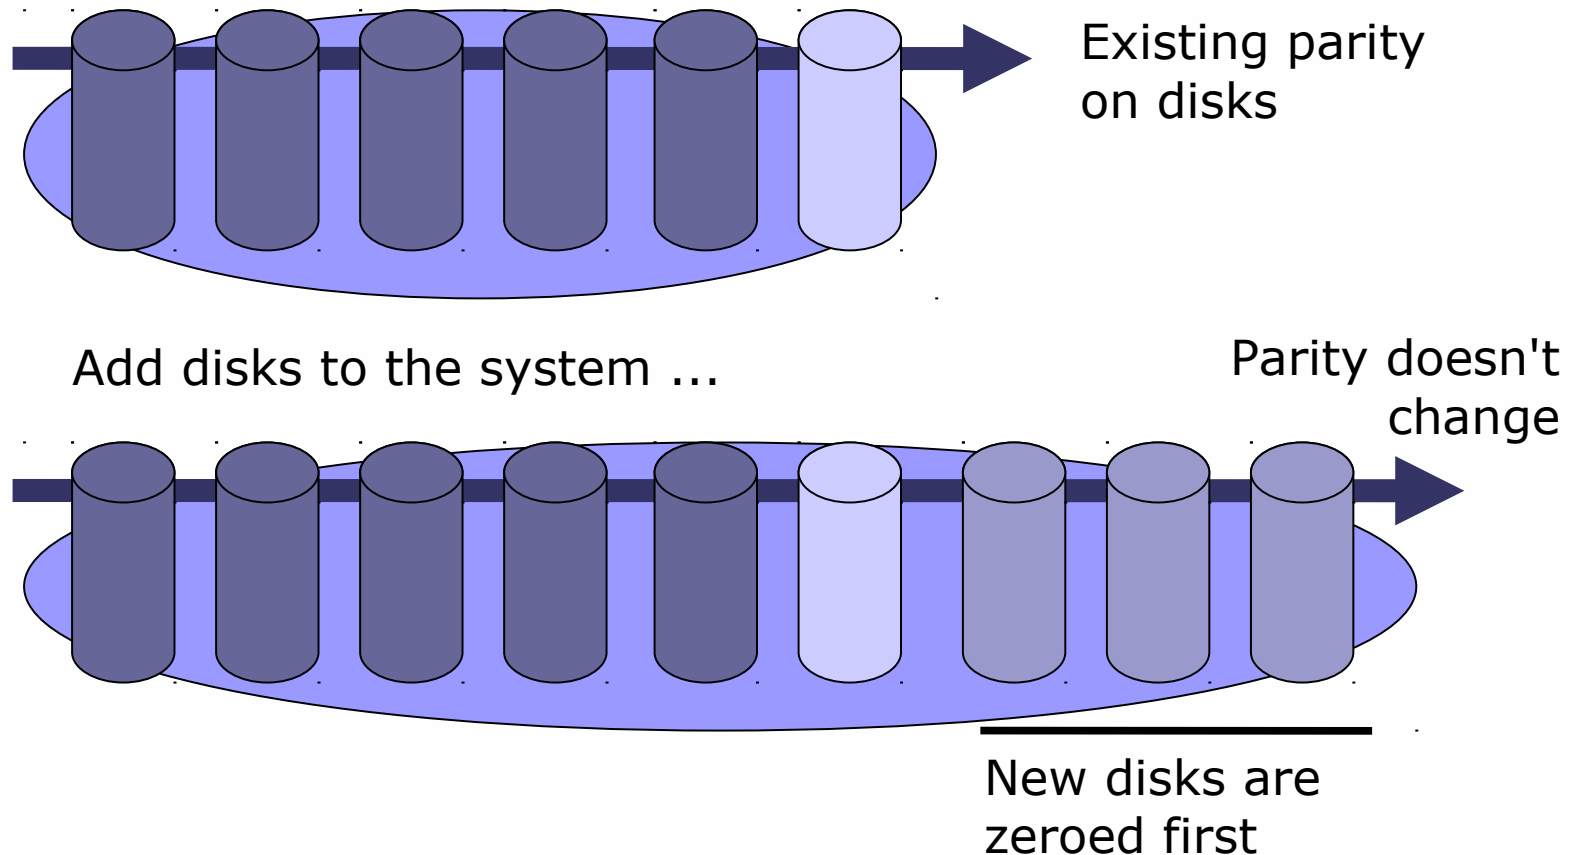


- You are **STRONGLY** encouraged to use RAID-DP instead of RAID-4 for better fault tolerance

NetApp Terminology (cont.)

RAID4 & RAID-DP (cont.)


- RAID4 advantages
 - Combined with WAFL filesystem: on-the-fly expansion of storage (no shrinking) without parity recalculation



NetApp Terminology (cont.)

RAID4 & RAID-DP (cont.)

- Traditional RAID4 has disadvantages
 - Parity disk becomes bottleneck on write operations
 - Parity disk is often the first disk to fail
- NetApp solves this by filer design
 - NVRAM (non-volatile RAM) in every filer
 - Battery-backed RAM (eg. 128 Mb on FAS270), split in two parts (four on cluster)
 - All write operations are written to NVRAM
 - When 50% of NVRAM is full, flush writes to disk:
 - “CP: Consistency Point”: spreads out writes over all disks
 - NVRAM will also be flushed every 10 seconds when not 50% full



NetApp Terminology (cont.)

RAID4 & RAID-DP (cont.)

```
filer> sysstat -x -s 1
CPU NFS CIFS HTTP Total Net kB/s Disk kB/s Tape kB/s Cache Cache CP CP Disk DAFS FCP iSCSI FCP kB/s
in out read write read write read write age hit time ty util in out in out
23% 0 0 0 17 105 45 18185 24 0 0 1 99% 0% - 94% 0 0 17 0 0
14% 0 0 0 2 75 3 10002 8 0 0 1 99% 0% - 63% 0 0 2 0 0
11% 0 0 0 8 110 8 8983 0 0 0 1 99% 0% - 43% 0 0 8 0 0
19% 0 0 0 1 6 2 13930 32 0 0 1 98% 8% Ss 72% 0 0 1 0 0
7% 0 0 0 8 130 803 3356 10020 0 0 1 96% 100% :v 14% 0 0 8 0 0
3% 0 0 0 1 6 2 1024 40 0 0 1 98% 100% Zf 10% 0 0 1 0 0
14% 0 0 0 8 132 186 4619 3612 0 0 1 99% 84% Z 40% 0 0 8 0 0
24% 0 0 0 19 130 56 22357 0 0 0 1 98% 0% - 100% 0 0 19 0 0
20% 0 0 0 31 161 158 20764 0 0 0 1 98% 0% - 99% 0 0 31 0 0
23% 0 0 0 16 72 30 22336 24 0 0 1 98% 0% - 100% 0 0 16 0 0
23% 0 0 0 48 253 221 20880 0 0 0 1 98% 0% - 100% 0 0 48 0 0
21% 0 0 0 16 121 49 20196 8 0 0 1 98% 0% - 100% 0 0 16 0 0
22% 0 0 0 8 130 3 20783 24 0 0 1 99% 0% - 100% 0 0 8 0 0
21% 0 0 0 5 4 18 20536 0 0 0 1 98% 0% - 100% 0 0 5 0 0
25% 0 0 0 23 140 73 21598 0 0 0 1 89% 0% - 96% 0 0 23 0 0
26% 0 0 0 14 144 13 20428 24 0 0 1 85% 0% - 100% 0 0 14 0 0
14% 0 0 0 38 31 958 14340 4080 0 0 1 87% 33% Zf 67% 0 0 38 0 0
5% 0 0 0 47 111 398 732 8960 0 0 1 76% 100% :v 10% 0 0 47 0 0
5% 0 0 0 81 91 417 1344 1376 0 0 1 83% 100% Zf 18% 0 0 81 0 0
CPU NFS CIFS HTTP Total Net kB/s Disk kB/s Tape kB/s Cache Cache CP CP Disk DAFS FCP iSCSI FCP kB/s
in out read write read write read write age hit time ty util in out in out
23% 0 0 0 22 55 108 9992 32 0 0 1 90% 6% : 54% 0 0 22 0 0
21% 0 0 0 15 115 39 10228 0 0 0 1 84% 0% - 49% 0 0 15 0 0
1% 0 0 0 5 20 3 260 32 0 0 1 97% 0% - 13% 0 0 5 0 0
22% 0 0 0 13 134 55 11029 0 0 0 1 83% 0% - 45% 0 0 13 0 0
21% 0 0 0 11 145 9 12442 0 0 0 1 83% 0% - 52% 0 0 11 0 0
3% 0 0 0 1 6 0 765 24 0 0 1 88% 0% - 10% 0 0 1 0 0
27% 0 0 0 21 158 60 14762 0 0 0 1 82% 0% - 60% 0 0 21 0 0
16% 0 0 0 14 57 70 8438 16 0 0 1 84% 0% -n 40% 0 0 14 0 0
18% 0 0 0 29 119 865 6472 8437 0 0 1 96% 100% Zf 22% 0 0 29 0 0
7% 0 0 0 39 86 95 4668 984 0 0 1 90% 85% Z 22% 0 0 39 0 0
27% 0 0 0 38 76 136 19660 20 0 0 1 86% 0% - 99% 0 0 38 0 0
25% 0 0 0 18 79 75 18064 4 0 0 1 87% 0% - 88% 0 0 18 0 0
22% 0 0 0 17 72 62 19337 8 0 0 1 96% 0% - 99% 0 0 17 0 0
22% 0 0 0 37 132 141 19076 8 0 0 1 96% 0% - 97% 0 0 37 0 0
20% 0 0 0 28 135 89 17974 16 0 0 1 97% 0% - 90% 0 0 28 0 0
22% 0 0 0 1 6 1 22696 0 0 0 1 99% 0% - 91% 0 0 1 0 0
22% 0 0 0 7 126 3 21224 12 0 0 1 99% 0% - 91% 0 0 7 0 0
20% 0 0 0 10 38 22 19776 20 0 0 1 98% 0% - 100% 0 0 10 0 0
22% 0 0 0 32 161 108 19592 0 0 0 1 95% 0% - 98% 0 0 32 0 0
27% 0 0 0 7 32 3 18347 16 0 0 1 96% 12% Ts 87% 0 0 7 0 0
```

NetApp Terminology (cont.)

Aggregates

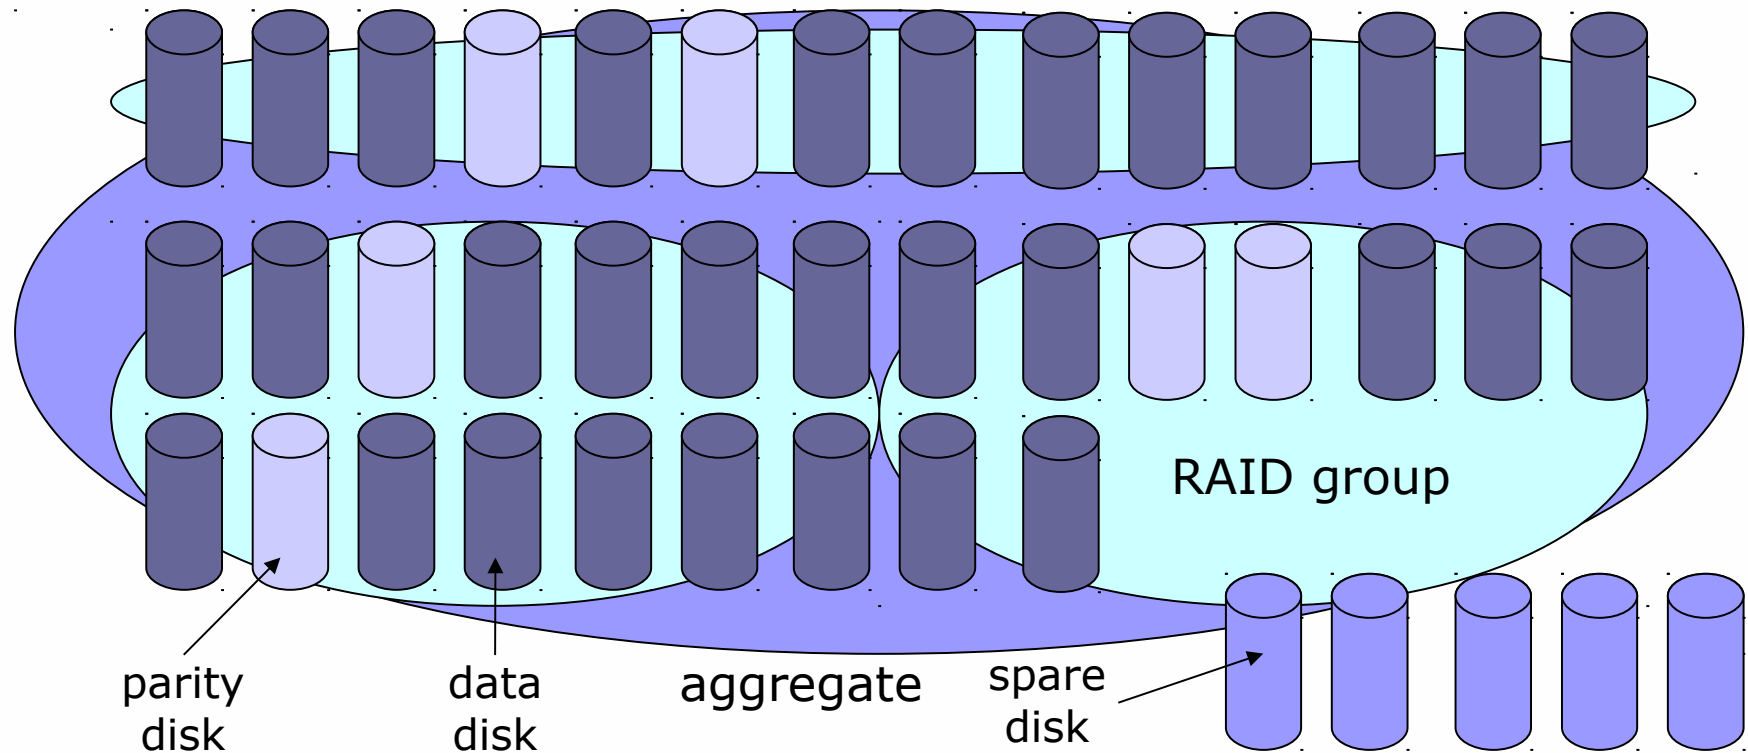
- Spare disks = not used
- Aggregate = collection of RAID-4/RAID-DP disks with parity or double parity, consists of one or more RAID groups
- Typically, all disks except spare disks in same aggregate
 - Exception: different disk sizes are in different aggregates (performance)
 - Exception: disks of different types (eg. FC vs. SATA)
 - Exception: Max 16 Tb raw capacity
- When disks are added to the system, aggregates can be expanded on-the-fly
- Aggregates **cannot shrink**, however!
- Aggregates do not contain data directly
- You must create flexible volumes (flexvols) inside aggregates. Volumes will contain data



NetApp Terminology (cont.)

RAID Group Size

- Every x data disks, provide 1 or 2 parity disks
- This "x" is the RAID group size, it's a property of an aggregate
- Default raid group sizes: 14, 16 disks ((S)ATA/FCP)



NetApp Terminology (cont.)

NetApp Cluster

- Not like eg. Windows cluster: No "shared" resources
- 2 nodes only, model name contains "c" or "ha", eg.
 - FAS270 vs. FAS270c
 - FAS3020 vs. FAS3020ha
- Hardware: a heartbeat cable between both systems
- Hardware: systems have connection to disks of other system
- In case of failure of one system: other system will notice, take over the disks and boot the first system from its disks
- CIFS users will get disconnected, but generally client will recover
- Same story for LUNs: retransmissions
- Note: NVRAM synchronization on clusters



NetApp Terminology (cont.)

What about all those Blinking Lights ?

- Every head and every shelf have plastic quick reference cards that can be pulled out from under the front panel
- eg. disk numbering



NOTE Set data rate switch to appropriate speed.

1 Gb/2 Gb data rate switch

STEP 1 Match OPS panel LEDs with the following possible conditions and perform action from key section

	FAULT CONDITIONS									
	1	2	3	4	5	6	7	8	9	10
POWER	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
FAULT	OFF	OFF	OFF	OFF	OFF	OFF	OFF	OFF	OFF	OFF
LOOP A	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
LOOP B	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
SYSTEM	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
REAR PANELS	MODULE A	MODULE B	MODULE A	MODULE B	PSU	FAN	FAN	FAN	FAN	FAN

KEY

- ON - No fault indicated.
- OFF - Module A fault - Check module A.
- OFF - Module B fault - Check module B.
- OFF - Module A loop A fault - Check module A or loop A.
- OFF - Module B loop B fault - Check module B or loop B.
- OFF - Enclosure fault - Contact Network Appliance customer support.
- OFF - Any PSU or fan fault - Check power supply unit LED panel status.
- OFF - Temperature fault - Check environmental conditions.
- OFF - OPS panel hardware fault - Contact Network Appliance customer support.
- OFF - Loop speed switch setting changed while powered up - Power cycle the disk shelf.
- OFF - Incorrect loop speed set for LRC - Set the loop speed for all DS14Mk2 FC disk shelves in the loop to 1Gb and power cycle the disk shelves.

Legend

- ON
- OFF
- ON/OFF
- FLASH
- 1 SEC. BEEP
- 20SEC. INTER.
- CONTINUOUS BEEP

DS14 - Mk2 - FC
OPS Panel 1 of 4

Table of Contents

- NetApp Products
- Storage Terminology
- NetApp Terminology
- **NetApp Hardware Essentials**
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Table of Contents

- Cables and Connectors
- Shelves & Shelf Modules



Shelves & Shelf Modules

- Shelf Modules
 - LRC (Loop Redundancy Circuit): older hardware
 - ESH, ESH2, ESH4 (Embedded Switched Hub): newer technology, ESH2+ has autotermination
 - AT-FC & AT-FC2: Used in R150/R200, single-path only
 - AT-FCX: newer technology
- Disk Shelves
 - DS14 (older hardware)
 - DS14Mk2-FC
 - DS14Mk2-AT
 - DS14Mk4-FC
 - DS12-ESAS
 - DS20-ESAS



Shelves & Shelf Modules

Shelf Modules



Shelf module is inserted into disk shelf cabinet

Can be disk shelf module or a "shrunk head" controller module with RAM, NICs, FCP HBAs, ...

This is a FAS270 "shrunk head" module. It transforms a regular shelf into a FAS270 or FAS270c (if 2 FAS270 modules are used)



Shelves & Shelf Modules

Shelf Modules



LRC shelf module



ESH shelf module –
notice the termination
switch





ESH2 shelf module – runs at 2
Gbps, is autoterminated

ESH4 shelf module (4 Gbps)
looks very similar

Shelves & Shelf Modules

NetApp Disk Shelves: DS14 Mk2 - FC



shelf ID

in out

1x module

2x modules

ESH

ESH2




GBIC

*Why 2x modules ?
→ redundancy
or clustered
systems*



Shelves & Shelf Modules

NetApp Disk Shelves: DS14 AT-FCX



shelf
ID in out



Data Cables



DB9 ("serial" or "console") cable is needed for connection to a controller

It is required during initial setup when there is no network connection to the filer



Ethernet network cables are needed for network connectivity (filer management, iSCSI, NFS, CIFS access, ...)



Data Cables (cont.)



FCP cables used can be optical or copper

Optical cables require the use of SFPs (Small Form-factor Pluggable Transceiver) on filer or shelf

Copper is via SFP connections or HSSDC2 for connection of shelves to FAS270



Data Cables (cont.)



SCSI cables for connection to tape devices





InfiniBand for cluster interconnect, now MTP cable is used with converter



Basic Filer Models

FAS250, FAS270, and FAS270c




Basic Filer Models (cont.)

FAS2020 and FAS2020ha




Basic Filer Models (cont.)

FAS2050 and FAS2050ha




Basic Filer Models (cont.)

FAS3020, FAS3040, FAS3050, FAS3070




Basic Filer Models (cont.)

FAS6030, FAS6040, FAS6070, FAS6080



Basic Filer Models (cont.)

FAS920, FAS940, FAS960, FAS980



Basic Filer Models (cont.)

R200





NetApp Hardware Essentials


(cont.)

Some NetApp-specific Terms ... (cont.)

– (Disk) Loop



0a



Shelves can be daisy-chained into a loop of **up to 6** shelves

...


by connecting the "out" of one shelf to the "in" of the next shelf

NetApp Hardware Essentials

(cont.)

Some NetApp-specific Terms ... (cont.)


- When adding shelves, one can either add shelves to existing loops, or create additional loops



NetApp Hardware Essentials

(cont.)

DS14 MkII Disk Shelf can be turned into FAS250/FAS270/FAS270c and vice versa



FAS250 module
"shrunk head"



ESH2 module
(autotermination)




ESH module



Backplane Speed Switch
(1/2/4 Gbps)

Redundant Power
Supply Units

Top and Bottom
Modules
(can be FAS250,
FAS270 or ESH(2))




Shelf ID switch


NetApp Hardware Essentials

(cont.)


Various Disk Shelf Modules (FC only)




LRC with Copper Interfaces



LRC with Optical Input and Copper Output



ESH



ESH2 (modern)



NetApp Hardware Essentials

(cont.)

NetApp Filer Models: FAS250



shelf ID

Fibre (FC) connection for tape backup

2x Gigabit NICs, can be teamed (VIF)

Connection for serial console cable



NetApp Hardware Essentials

(cont.)

NetApp Filer Models: FAS270(c)



Second module installed = FAS270c (cluster)

shelf ID

Fibre (FC) connection for SAN & tape backup

Fibre (FC) connection to additional disk shelves

2x Gigabit NICs, can be teamed (VIF)


Connection for serial console cable



NetApp Hardware Essentials

(cont.)

NetApp Filer Models: FAS3020, FAS3050, FAS3070



4x Modular I/O Expansion Slots

Connection for serial console cable

4x Fibre (FC) connection for disk shelves or FC SAN:
0a, 0b, 0c, 0d

1x RLM NIC (Remote LAN Module)

4x Gigabit NICs, can be teamed (VIF):
e0a, e0b, e0c, e0d

SCSI connection for tape backup:
0e (not on all models)

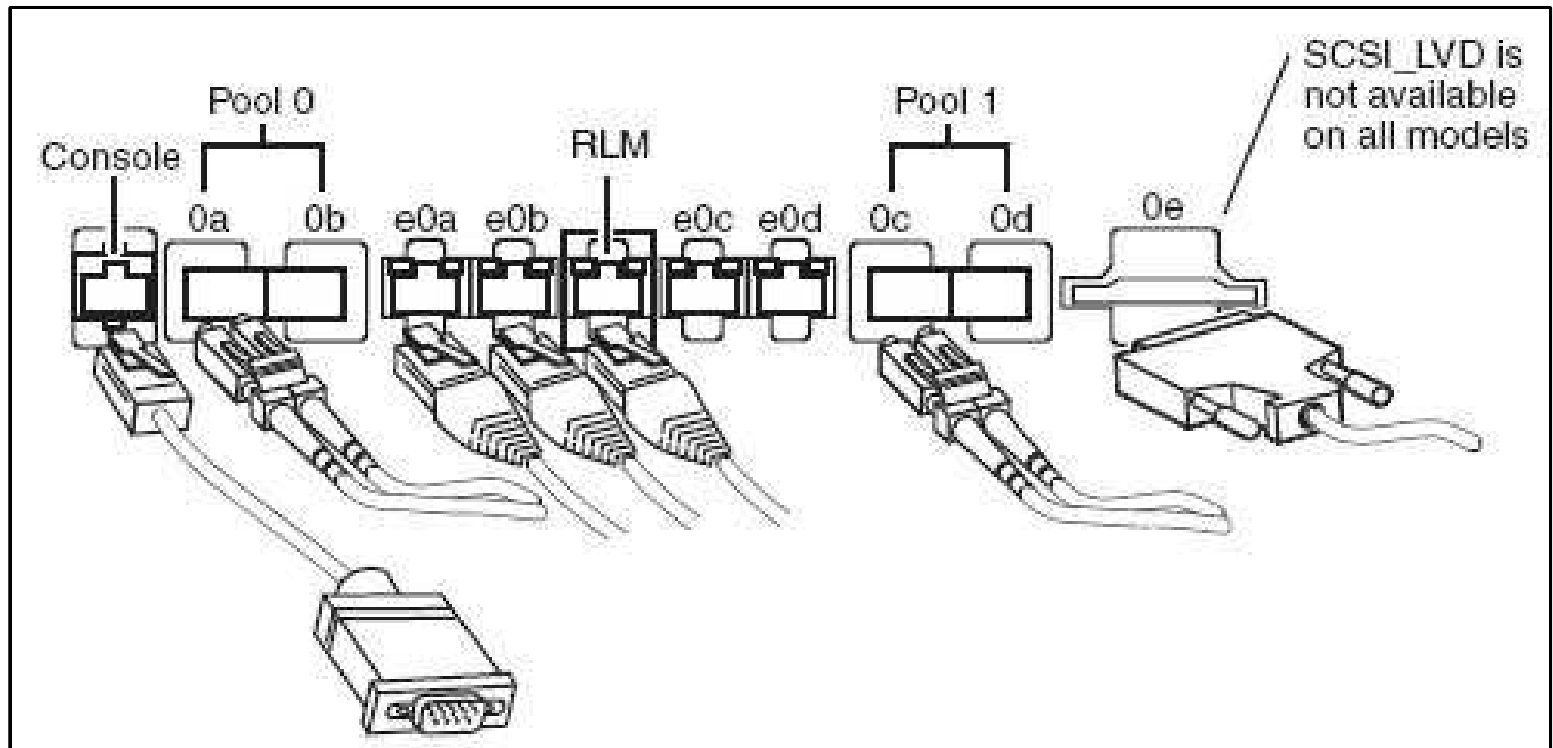


NetApp Hardware Essentials

(cont.)

NetApp Filer Models: FAS3020, FAS3050, FAS3070

Connections




NetApp Hardware Essentials

(cont.)

NetApp Filer Models: FAS60x0



10x Modular Expansion Slots



6x Gigabit
NICs, can
be teamed
(VIF):
e0a, e0b,
e0c, e0d,
e0e, e0f

1x RLM
NIC
(Remote
LAN
Module)

8x Fibre (FC) connection for disk shelves
or FC SAN: 0a, 0b, 0c, 0d, 0e, 0f, 0g, 0h


Connection
for serial
console
cable



NetApp Hardware Essentials


(cont.)

NetApp Filer Models: FAS960 (older model)




11x Modular Expansion Slots

1x 10/100 Mbps NIC



Connection for serial console cable



NetApp Hardware Essentials


(cont.)

Common Cabling Examples

Standard Filer Cabling: FAS270



Total capacity = 3 shelves
FAS270, shelf ID = 1,
disk 0b.16 → 0b.29



Disk shelf, shelf ID = 2
disk 0b.32 → 0b.45



Disk shelf, shelf ID = 3
disk 0b.48 → 0b.61

One "FCAL loop"




NetApp Hardware Essentials

(cont.)

Common Cabling Examples (cont.)

Standard Filer Cabling: FAS3020/3050



Total capacity = 2 shelves
FAS3020/3050, FC port 0a
set to "initiator"

Disk shelf, shelf ID = 1
disk 0a.16 → 0a.29

Disk shelf, shelf ID = 2
disk 0a.32 → 0a.45

One "FCAL loop"




NetApp Hardware Essentials


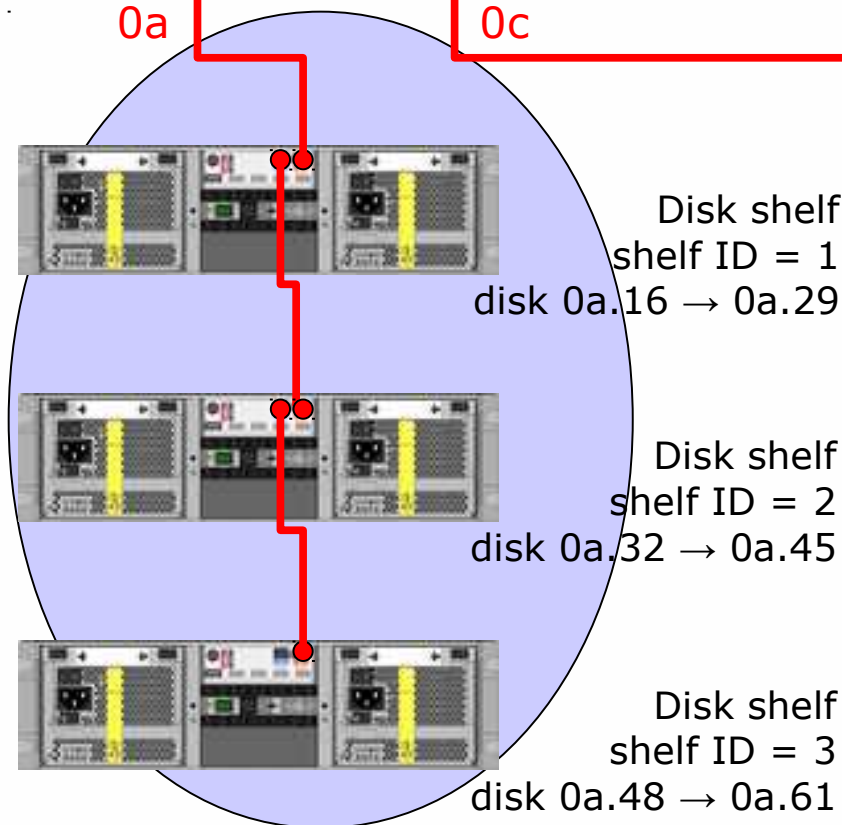
(cont.)

Common Cabling Examples (cont.)

Standard Filer Cabling: FAS3020/3050, two disk loops



Total capacity = 5 shelves
FAS3020/3050, FC port 0a
and 0c set to "initiator"



- Two "FCAL loops", why ?
- when backplane speed is maxed out (2 Gbps)
 - don't mix FC and SATA disks in same loop

NetApp Hardware Essentials


(cont.)

Common Cabling Examples (cont.)

Standard Filer Cabling: FAS3020/3050, two disk loops Example

```
filer> sysconfig -a
...
slot 0: FC Host Adapter 0a (Dual-channel, QLogic 2322 rev. 3, 64-bit,
L-port, <UP>)
Firmware rev: 3.3.10
Host Loop Id: 7 FC Node Name: 5:00a:098200:006b3b
Cacheline size: 16 FC Packet size: 2048
SRAM parity: Yes External GBIC: No
Link Data Rate: 2 Gbit
41: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9ZV1A)
42: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9VNTA)
39: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W981KA)
38: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA33HA)
37: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA171A)
36: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA2W6A)
32: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA3B1A)
35: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9VDP A)
34: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W84HX A)
33: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9JMS A)
40: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9BZ9 A)
29: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WBKGVA)
28: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WAX0TA)
27: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WASYRA)
25: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W74NGA)
26: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W768HA)
23: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W71TEA)
22: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WAAMHA)
21: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5VAJH7A)
20: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA2W3A)
16: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WB M9VA)
19: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WB DXSA)
18: NETAPP X274_SCHT6146F10 NA08 136.0GB 520B/sect (3HY4FWVP)
17: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WBJ20A)
24: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA5MAA)
45: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W82V9A)
44: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9AR9A)
43: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W7069A)
Shelf 1: ESH2 Firmware rev. ESH A: 14 ESH B: 14
Shelf 2: ESH2 Firmware rev. ESH A: 14 ESH B: 14
I/O base 0xee00, size 0x100
memory mapped I/O base 0xe1940000, size 0x1000
slot 0: FC Host Adapter 0b (Dual-channel, QLogic 2322 rev. 3, 64-bit,
L-port, <OFFLINE (hard)>)
...
```

```
...
slot 0: FC Host Adapter 0c (Dual-channel, QLogic 2322 rev. 3, 64-bit,
L-port, <UP>)
Firmware rev: 3.3.10
Host Loop Id: 7 FC Node Name: 5:00a:098000:006b3b
Cacheline size: 16 FC Packet size: 2048
SRAM parity: Yes External GBIC: No
Link Data Rate: 2 Gbit
21: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA2USA)
20: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA5NJA)
17: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA3J7A)
22: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9U72A)
23: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9TRMA)
25: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W8475A)
26: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA006A)
29: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA394A)
27: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA5SLA)
28: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9GDEA)
24: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA30AA)
16: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA3TMA)
18: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WA5UKA)
19: NETAPP X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5W9UBDA)
Shelf 1: ESH2 Firmware rev. ESH A: 14 ESH B: 14
I/O base 0xee00, size 0x100
memory mapped I/O base 0xe1940000, size 0x1000
slot 0: FC Host Adapter 0d (Dual-channel, QLogic 2322 rev. 3, 64-bit,
L-port, <OFFLINE (hard)>)
...
```



Note: this is from looking at one machine, it may very well be a clustered system !




NetApp Hardware Essentials

(cont.)

Common Cabling Examples (cont.)

File Cabling: FAS3020/3050, double connected disk loops for redundancy & load balancing



Total capacity = 3 shelves
 FAS3020/3050, FC port 0a
 and 0c set to "initiator"


Two "FCAL loops" for redundancy & load balancing – the same disk can be addresses as 0a.16/0c.16 etc.

NetApp Hardware Essentials

(cont.)

Common Cabling Examples

Cluster Filer Cabling: FAS270c



*Total capacity = 3 shelves,
divided over 2 filers*

FAS270c, shelf ID = 1,
disk 0b.16 → 0b.29

Disk shelf, shelf ID = 2
disk 0b.32 → 0b.45

Disk shelf, shelf ID = 3
disk 0b.48 → 0b.61

Two "FCAL loops", two filers
Who "owns" disks ?

Software-based disk ownership: "disk assign" command

NetApp Hardware Essentials

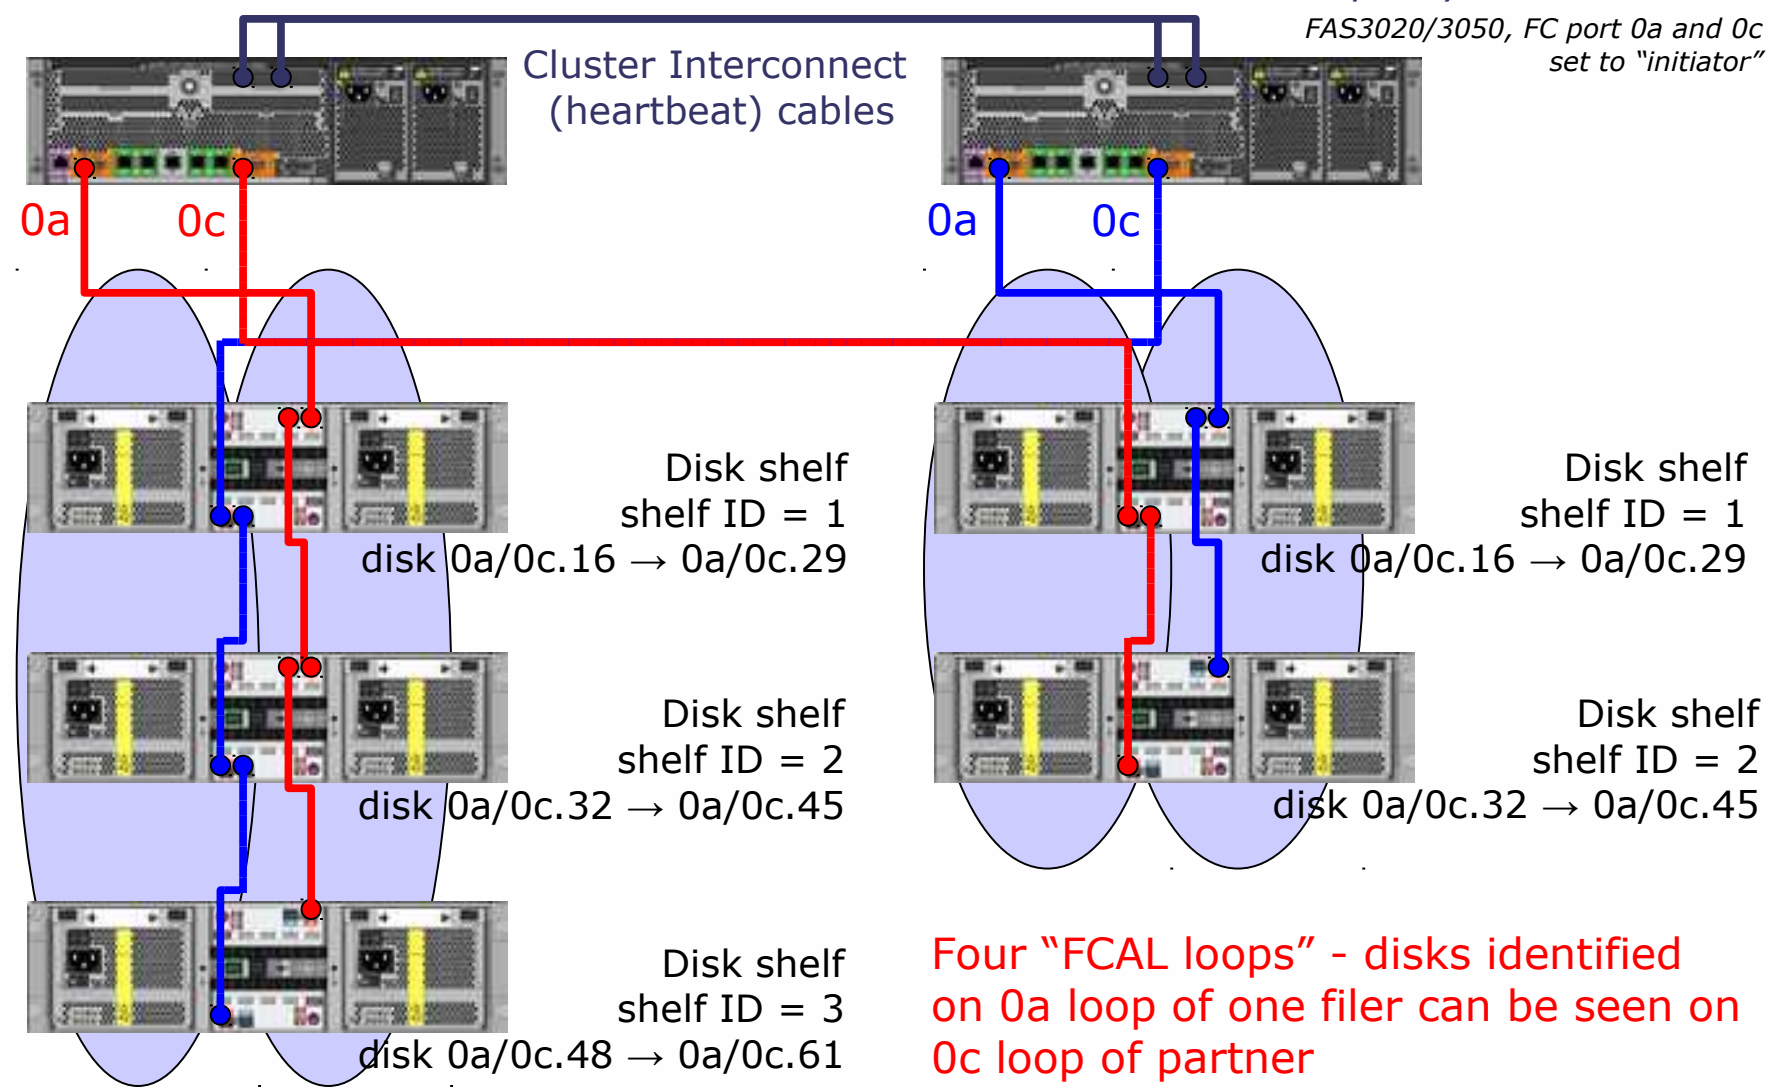
(cont.)

Common Cabling Examples (cont.)

Cluster Filer Cabling: FAS30x0 Standard Cluster

Total capacity = 3 + 2 shelves

FAS3020/3050, FC port 0a and 0c set to "initiator"



NetApp Hardware Essentials

(cont.)

Common Cabling Examples (cont.)

Cluster Filer Cabling: FAS3020/3050 Standard Cluster

Important Notes about Previous Slide

- Disk ownership ? 2 "independent" nodes, who owns what disks ?
 - **Software-based** ownership is still possible (eg. upgrade from FAS270c): in this case, disks can be literally anywhere in shelves
 - Typically: **Hardware-based**:
 - The filer that connects to the TOP module of a shelf controls the disks in that shelf under normal (ie. non-failover) circumstances



Whoever connects to this module, owns the disks in this shelf under normal circumstances

- So-called "mailbox disks" (and backup mailbox disks) act as quorum
- If upgrading from FAS270c, remove ownership ("disk" command from maintenance mode) and put disks in correct shelves
- Why always 0a/0c ?
 - 0b/0d for additional shelves
 - It's **not always 0a/0c** - Check "**FC Config Guide**" on NOW site
 - Different scenarios are possible, eg. 0a/0b & 0c/0d in FC SAN configs!
 - Can get very complicated in dual fabric SAN environments
 - Additional FC cards in PCI(-Express) slots are possible !

NetApp Hardware Essentials

(cont.)

Common Cabling Examples (cont.)

Important Note !

For a more thorough overview of all supported cabling configurations, you must read (and follow) the FC Config Guide, available from the NOW site !

- Details setups with single/dual fabrics
- Details FC adapter cards & cabling changes
- Discusses fcp mode (single, partner, standby, dual_fabric, ...) (SAN setups are not the topic of this presentation)

So far, we are not protected against complete shelf failures.
We need a **Metrocluster** design to provide this feature



NetApp Hardware Essentials

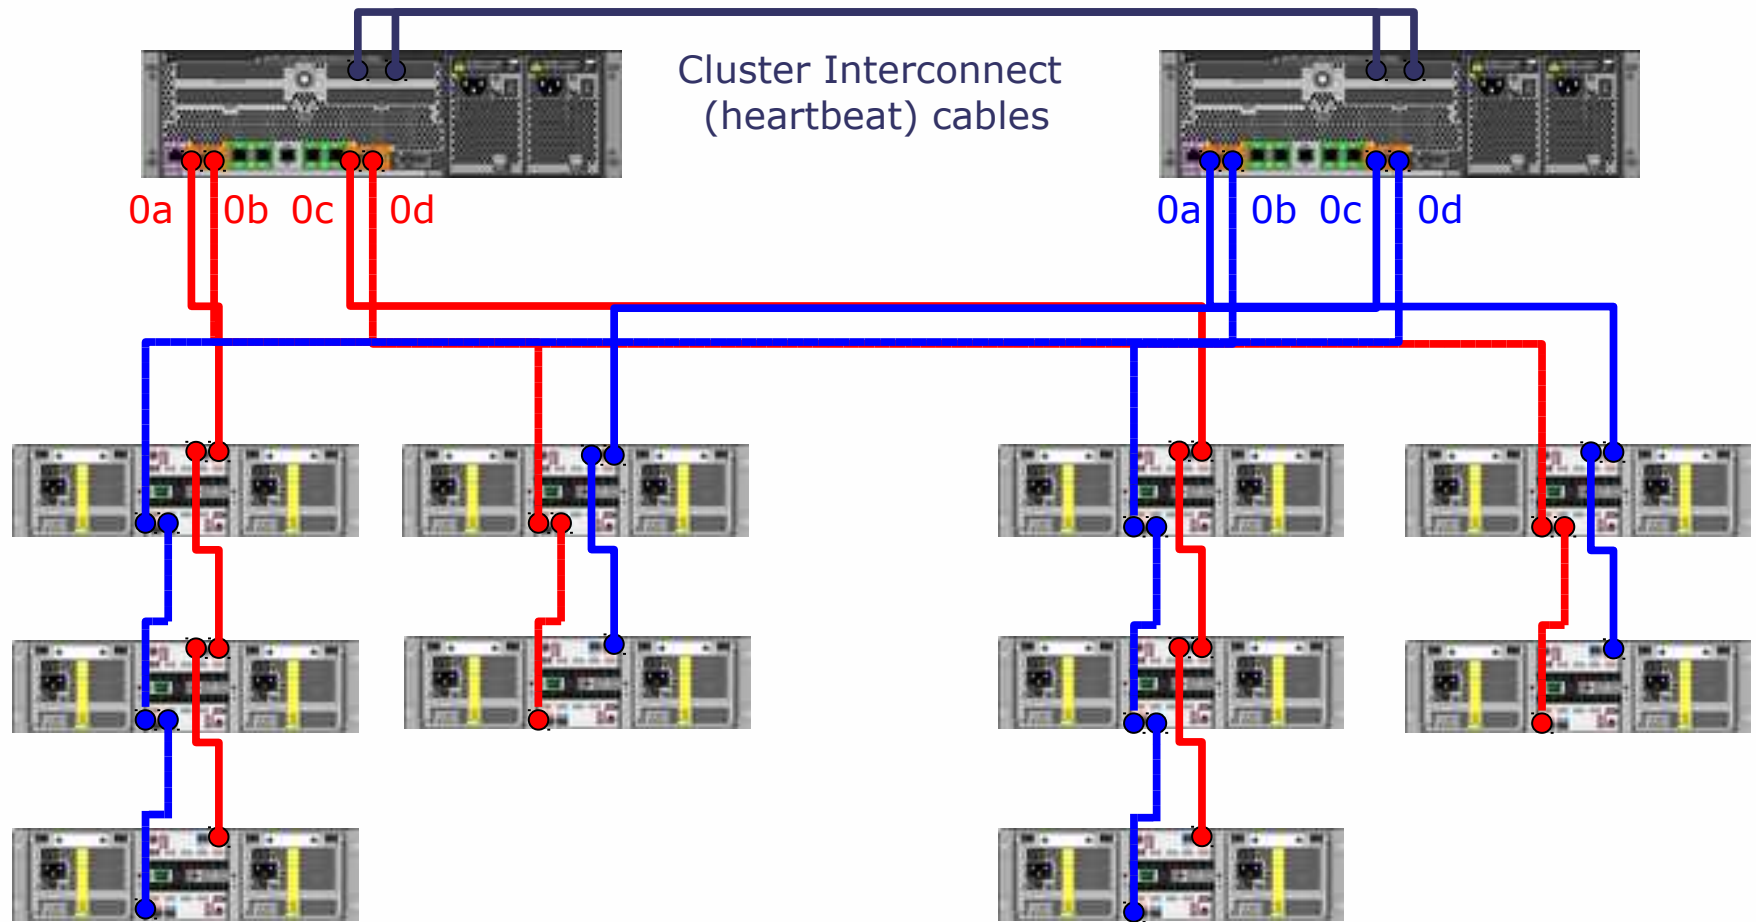
(cont.)

Common Cabling Examples (cont.)

Cluster Filer Cabling: FAS3020/3050 Metrocluster (stretched)

Total capacity = 3 + 2 shelves

FAS3020/3050, FC ports 0a, 0b, 0c, and 0d set to "initiator"



First node, disks @ local site, pool0

Second node, disks @ remote site, pool1 (mirror)

First node, disks @ remote site, pool1 (mirror)

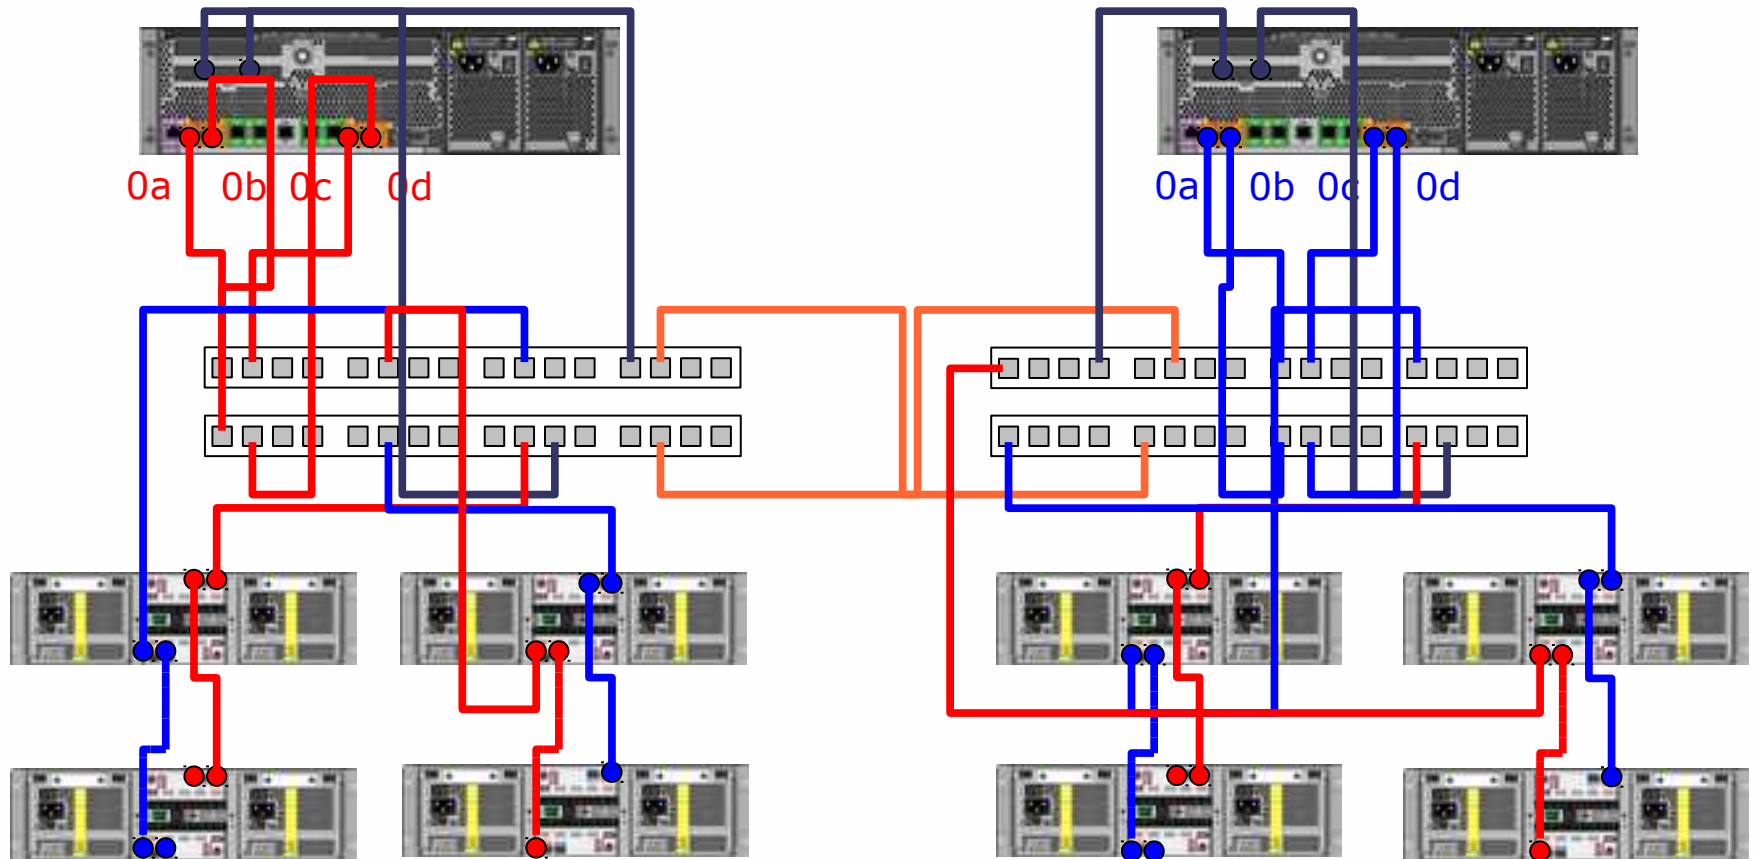
Second node, disks @ local site, pool0

NetApp Hardware Essentials

(cont.)

FAS3020/3050 Switched Metrocluster

Total capacity =
3 + 2 shelves



First node, disks @
local site, pool0

Second node, disks @
remote site, pool1 (mirror)

First node, disks @ remote
site, pool1 (mirror)

Second node, disks @
local site, pool0

Table of Contents

- NetApp Products
- Storage Terminology
- NetApp Terminology
- NetApp Hardware Essentials
- **Where to Get Help ?**
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Where to Get Help ? (cont.)

Confused about a syslog message ?

- Syslog translator:
- <https://now.netapp.com/eservice/ems>



Software* Release
Search String*¹

¹The search string must include the EMS Identifier; please enter the entire Syslog message when possible.

Tue Jun 3 09:11:32 BST [syntapfiler: cf.disk.shelfCountMismatchOK :CRITICAL]: cfdisk: previous shelf count mismatch resolved



EMS Identifier

Entire Syslog Message

Where to Get Help ? (cont.)

Filer Model Restrictions

Where to find out maximum number of shelves, maximum aggregate size, etc. of particular NetApp model and ONTAP version ?

→ NOW Site, System Configuration Guide:

<http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/>

The screenshot shows a web browser window displaying the "NetApp FAS270 Disk Shelf Support" page. The page content includes a table with the following data:

Disk shelf	Disk drives ^{1, 2, 3}	Standard power supply ⁴ optional AT ⁵	Notes
DiskShelf10k2 FC	36-GB Fibre Channel (PN: X270A or X271) 72-GB Fibre Channel (PN: X270A or X271) 144-GB Fibre Channel (PN: X270A/X271) ⁶ 300-GB Fibre Channel (PN: X270A)	AC power supply: 110002V PN: X511A Two power supplies per shelf required	Up to three Fibre Channel (DiskShelf10k2 FC) shelves are supported in any installable configuration. Empty shelves: With AC power supplies: PN: X550 Without power supplies: PN: X551 Rack mount kit: PN: 3008A Copper LRC module: PN: X5501A Optical LRC module: PN: X5502A EIOF module: PN: X511A FC SFP-to-SFP Patch Cable: PN: X5530 FC SFP-to-10GbE Patch Cable: PN: X5531
DiskShelf10k2 AT	350-GB ATA (PN: X2820-R5) 320-GB ATA (PN: X2820-R5) 500-GB ATA (PN: X282A-R5)	Maximum of 2 power supplies per shelf (PN: 9311N) Maximum of 14 per shelf	Up to two DiskShelf10k2 AT shelves are supported in any installable configuration. AT rack module: PN: X5512A-R5 Rack mount kit: PN: 3008A Optical SFP: PN: X5529 FC SFP-to-SFP Patch Cable: PN: X5530 Note: The first shelf in the loop must have optical modules.

Footnotes:

- 1: The maximum number of drives for an expansion shelf connecting to the FAS270/FAS270c is ten FC/AL drives or five ATA drives.
- 2: Each number is a multiple of the number of shelves supported and the size of the shelf. It does not indicate storage size.
- 3: X271, X270, and X270c are 15,000 RPM drives. See technical support TR8086 for information about using X271, X270, or X270c with other drives in the same appliance.

Last updated: April 11, 2006

Where to Get Help ? (cont.)

Exercises

- On the NOW site, look up the documentation for:
 - SnapManager for SQL (latest version)
 - SnapManager for Exchange (latest version)
- On the NOW site, look up the toolchest
- On the NOW site, locate the System Configuration Guide
- On the NOW site, download the latest firmware for your filer
- On the NOW site, locate your NetApp products and serial numbers
- On the NOW site, locate the Autosupport analysis for your filer
- On the NOW site, look up replacements parts for your filer
- On the NOW site, search the Bugs Online database



Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- **Disks, RAID4, Aggregates and Space Calculation**
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Disks, RAID4, Aggregates, and Space Calculation (cont.)

Space Calculation

... a.k.a. what you will lose

no
control

- Disk vendors lie about actual disk size
 - 144 Gb disk = 136 Gb capacity
- WAFL reserves 10% of a disk's space (unreclaimable)
- Parity disks and double parity disks don't contain data
- The system needs one, possibly two spare disks
- WAFL will reserve 5% snapshot reserve for aggregates
- WAFL will reserve 20% snapshot reserve for volumes
- NAS snapshots consume about 1% of space every day
- SAN snapshots consume 10%-20% of space every day
- LUNs need to have snapshot reservation enabled (x2)
- LUNs & SnapDrive mounts (.rws files) need reservation too

control



Disks, RAID4, Aggregates, and Space Calculation (cont.)

10% WAFL Space

- A disk that is 100% full has terrible write performance (fragmentation); by limiting the disk utilization to 90% of its full size, a "full" WAFL data disk still has somewhat "decent" performance
- WAFL metadata (eg. related to snapshot overhead) is stored in the 10% reserved space
- Bad blocks do not necessarily mean that a disk should be failed. However, the bad blocks should be relocated to a different disk location – WAFL 10% is used



Disks, RAID4, Aggregates, and Space Calculation (cont.)

See Also

<http://www.secalc.com/>



Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- **When Disks Go Bad ...**
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



When Disks Go Bad ...

Possible Failures

- RAID4 protects against single-disk failures in a RAID group
- RAID_DP protects against double-disk failures in a RAID group
- Plex/shelf mirroring (SyncMirror license, as in Mirrored Cluster or MetroCluster) protects against multiple-disk failures in same RAID group
- System will use one or more spare disks to recover in the background, while system keeps on running
- In all other failure scenarios, your only resort for data recovery is to quickly pick up a religion of choice and start praying really hard (I recommend Voodoo - has been known to work in some cases)



When Disks Go Bad ... (cont.)

Possible Failures (cont.)

- If a system is in a state where one more loss of a disk in a RAID group will result in DATA loss for that RAID group (i.e. there are not enough spare disks), the system is in **DEGRADED mode** and will **shut down in 24** hours as a data protection measure
- Idem when low NVRAM battery (will not boot until battery is sufficiently reloaded)
- Notes:
 - Timeout is configurable via "options raid.timeout" (default 24 hours)
 - System with no spares but at least one parity/dparity disk in RAID groups is NOT degraded
 - Again: no spare disks available does not necessarily mean "degraded"



When Disks Go Bad ... (cont.)

How to Replace Disks ?

- Important commands (priv set advanced):
 - `blink_on <disk-id>` `led_on <disk-id>`
 - `blink_off <disk-id>` `led_off <disk-id>`
- Use them creatively to identify failed disks
 - Normally, failure light should be on
 - If not, make disk LED blink
 - If LED is broken, make disk to the left & right blink
- To replace failed disk
 - `disk remove ...`



Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- **Volumes**
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Volumes

- Contain actual data
- WAFL (Write Anywhere Filesystem Layout) filesystem (NOT NTFS, ...)
 - WAFL allows NTFS permissions on files & dirs, though
- 2 types:
 - Traditional, “**TradVols**” (ONTAP 6 and earlier)
 - Flexible, “**FlexVols**” (use this!)
 - Can grow and shrink on-the-fly !
- Will contain either:
 - Files (NAS)
 - LUNs (SAN)




Volumes (cont.)


Why Flexvols ?

→ **Maximize Storage Utilization and Performance with Virtualization**

- Less capacity utilization
- Simplify provisioning & data management
- Thin provisioning possible



Before Data ONTAP 7G



With Data ONTAP 7G




Volumes (cont.)

Why FlexVols ? (cont.)

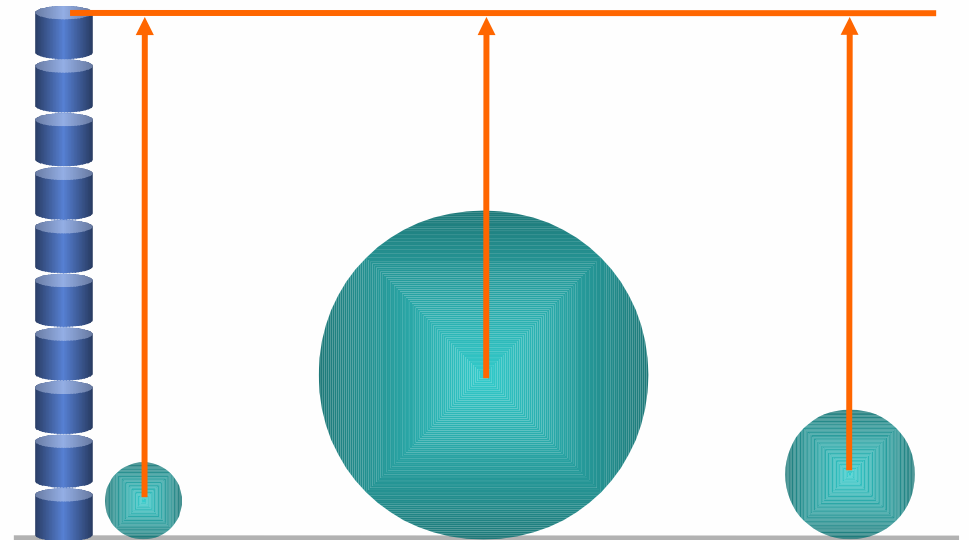
→ Regular Volumes

- Volume performance limited by number of disks it has
- "Hot" volumes can't be helped by disks on other volumes



→ FlexVol Volumes

- Spindle sharing makes total aggregate performance available to all volumes



(Note: FlexShare in DATA
ONTAP 7.2 (7.1))



Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- **Snapshots**
- Qtrees
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Snapshots

= A frozen, read-only image of a traditional volume, a flexible volume, or an aggregate that reflects the state of the new file system at the time the snapshot was created

Notes:

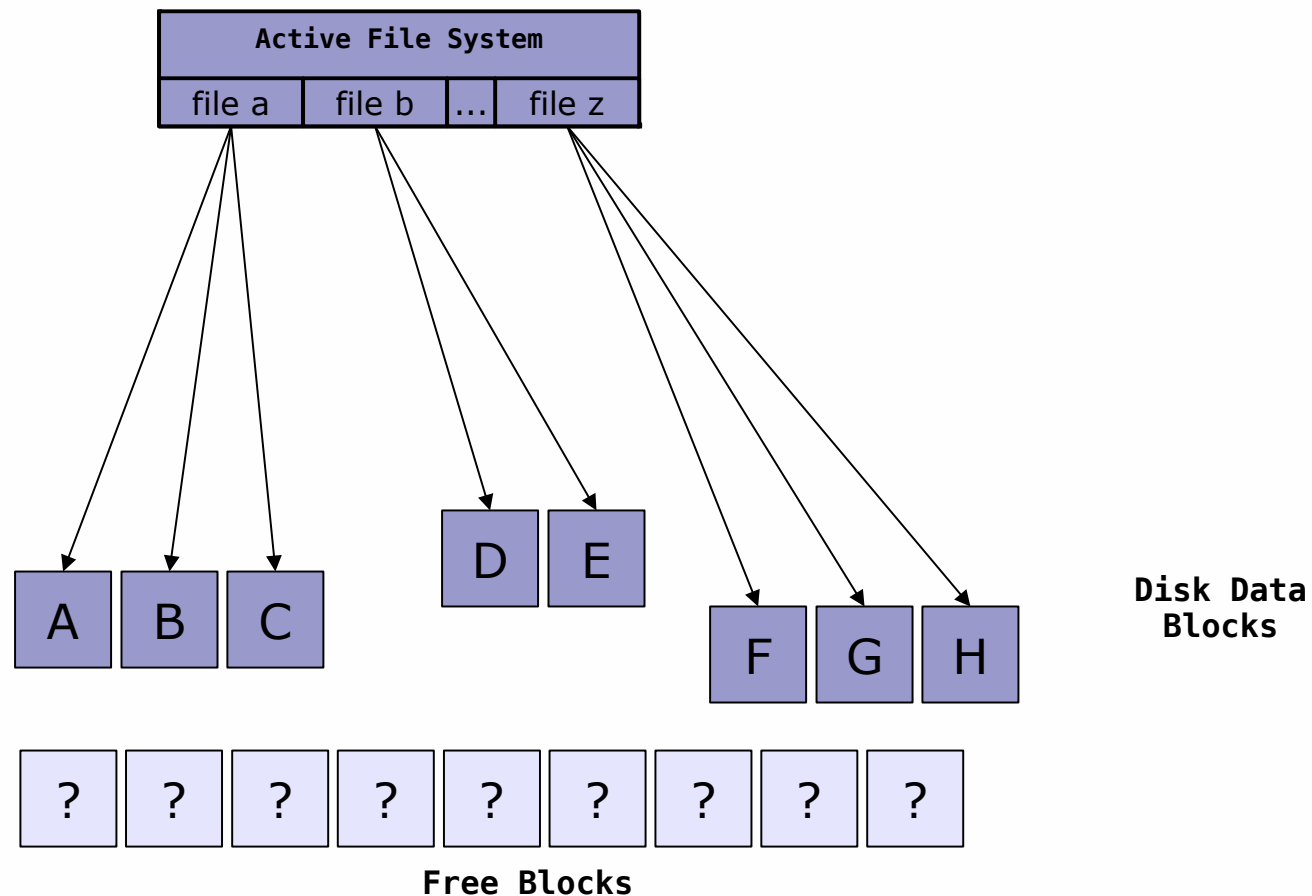
- Up to 255 snapshots per volume
- Can be scheduled
- Maximum space occupied can be specified (default 20%)
- File permissions are handled



Snapshots (cont.)

A Bird's Eye View at Snapshots & SnapRestore

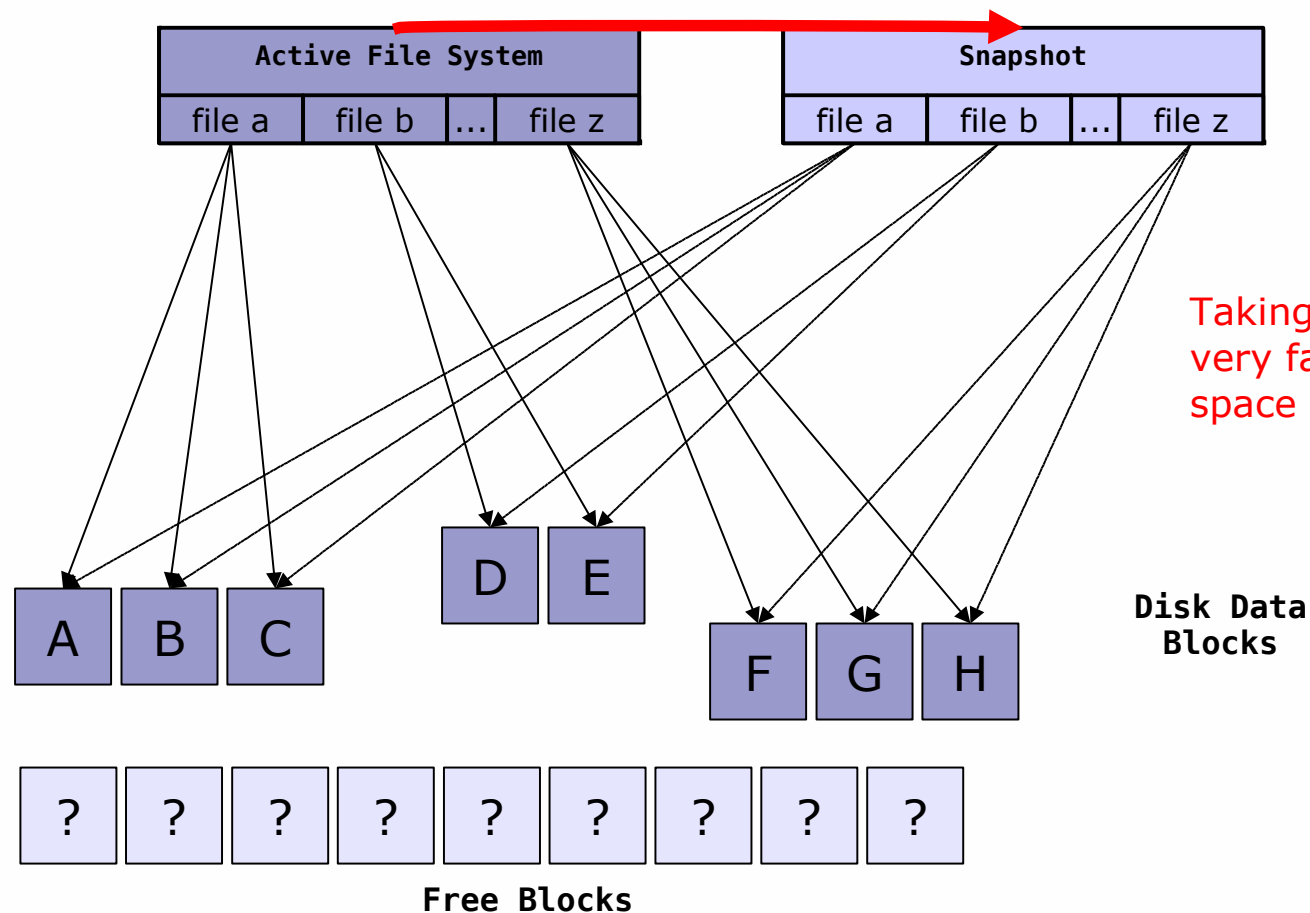
"Snapshots"



Snapshots (cont.)

A Bird's Eye View at Snapshots & SnapRestore

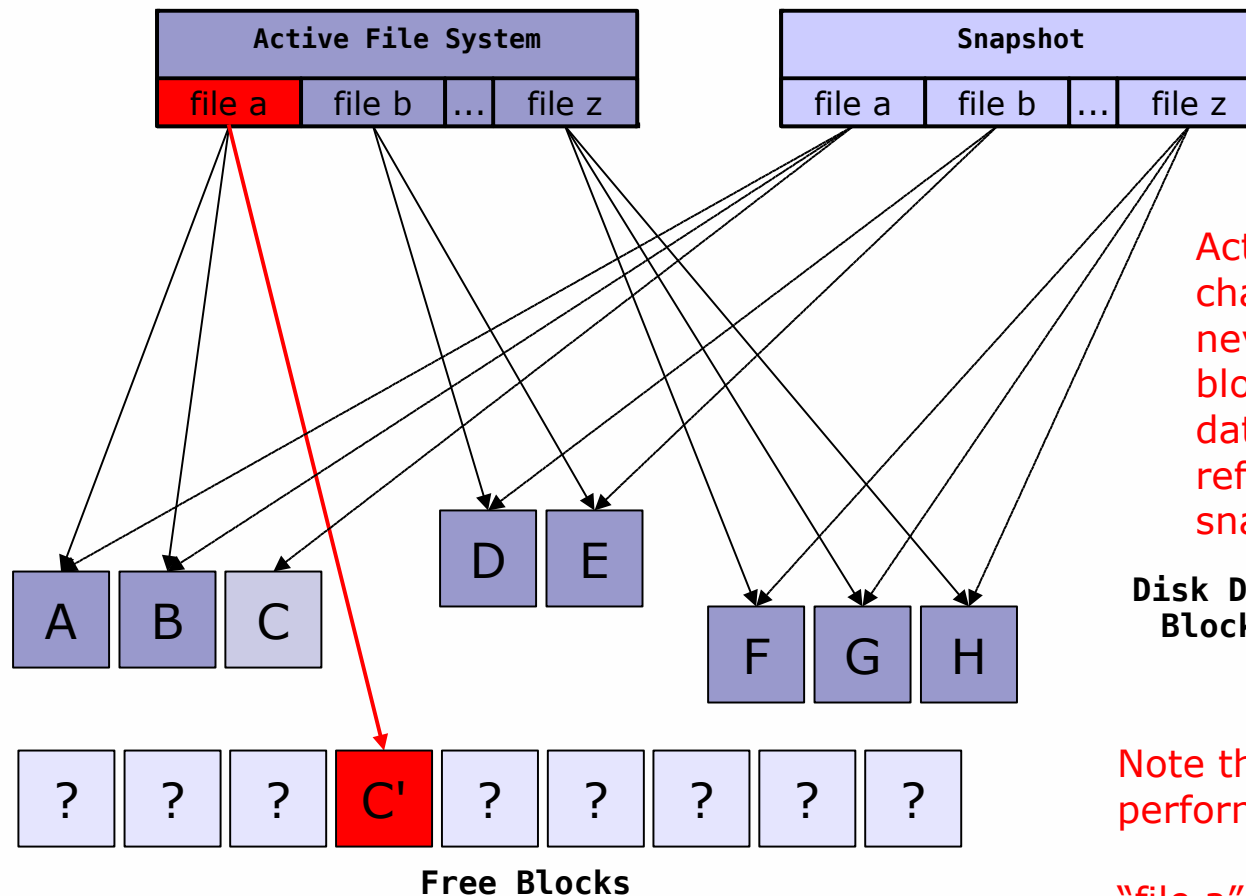
"Snapshots"



Snapshots (cont.)

A Bird's Eye View at Snapshots & SnapRestore

"Snapshots"



Active file system changes = WAFL never overwrites data blocks, old ("freed") data blocks are referenced to from the snapshot

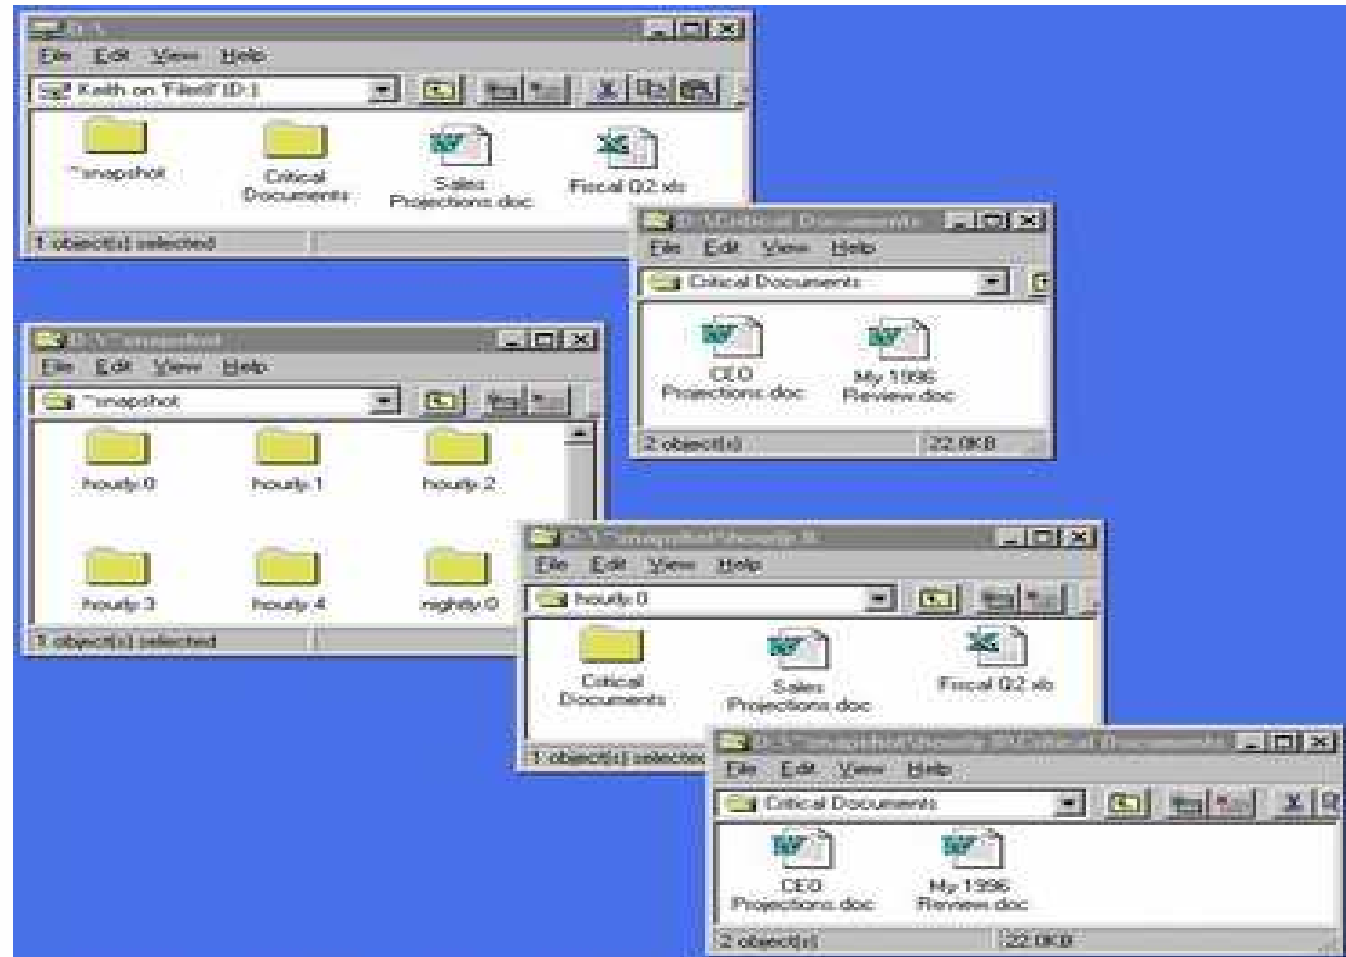
Disk Data Blocks

Note that there is no performance impact

"file a" now consists of blocks A, B, C'

Snapshots (cont.)

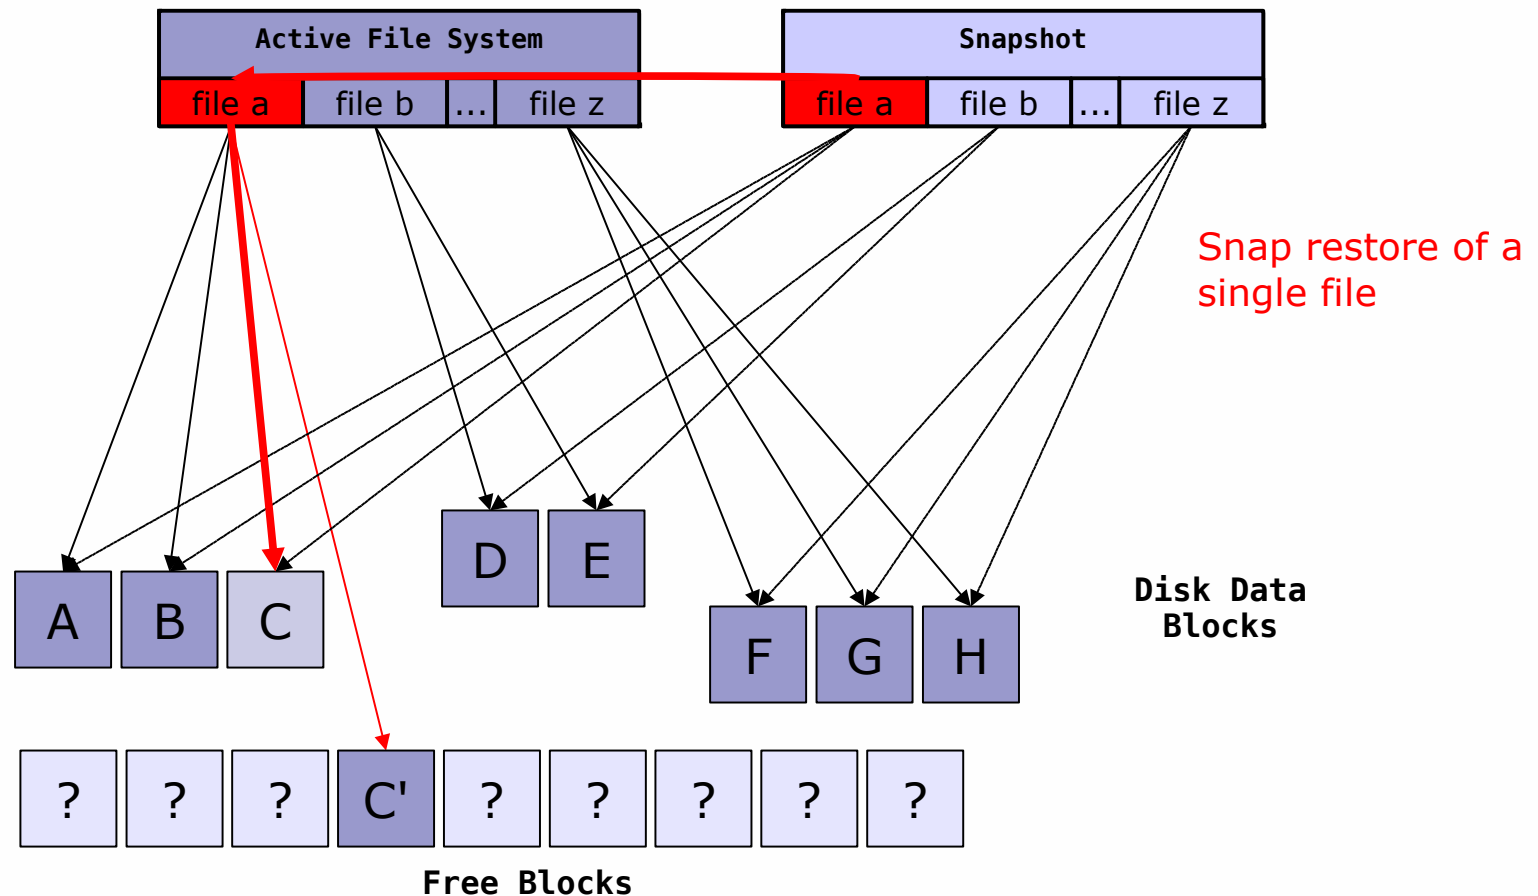
A Bird's Eye View at Snapshots & SnapRestore (cont.)



Snapshots (cont.)

A Bird's Eye View at Snapshots & SnapRestore

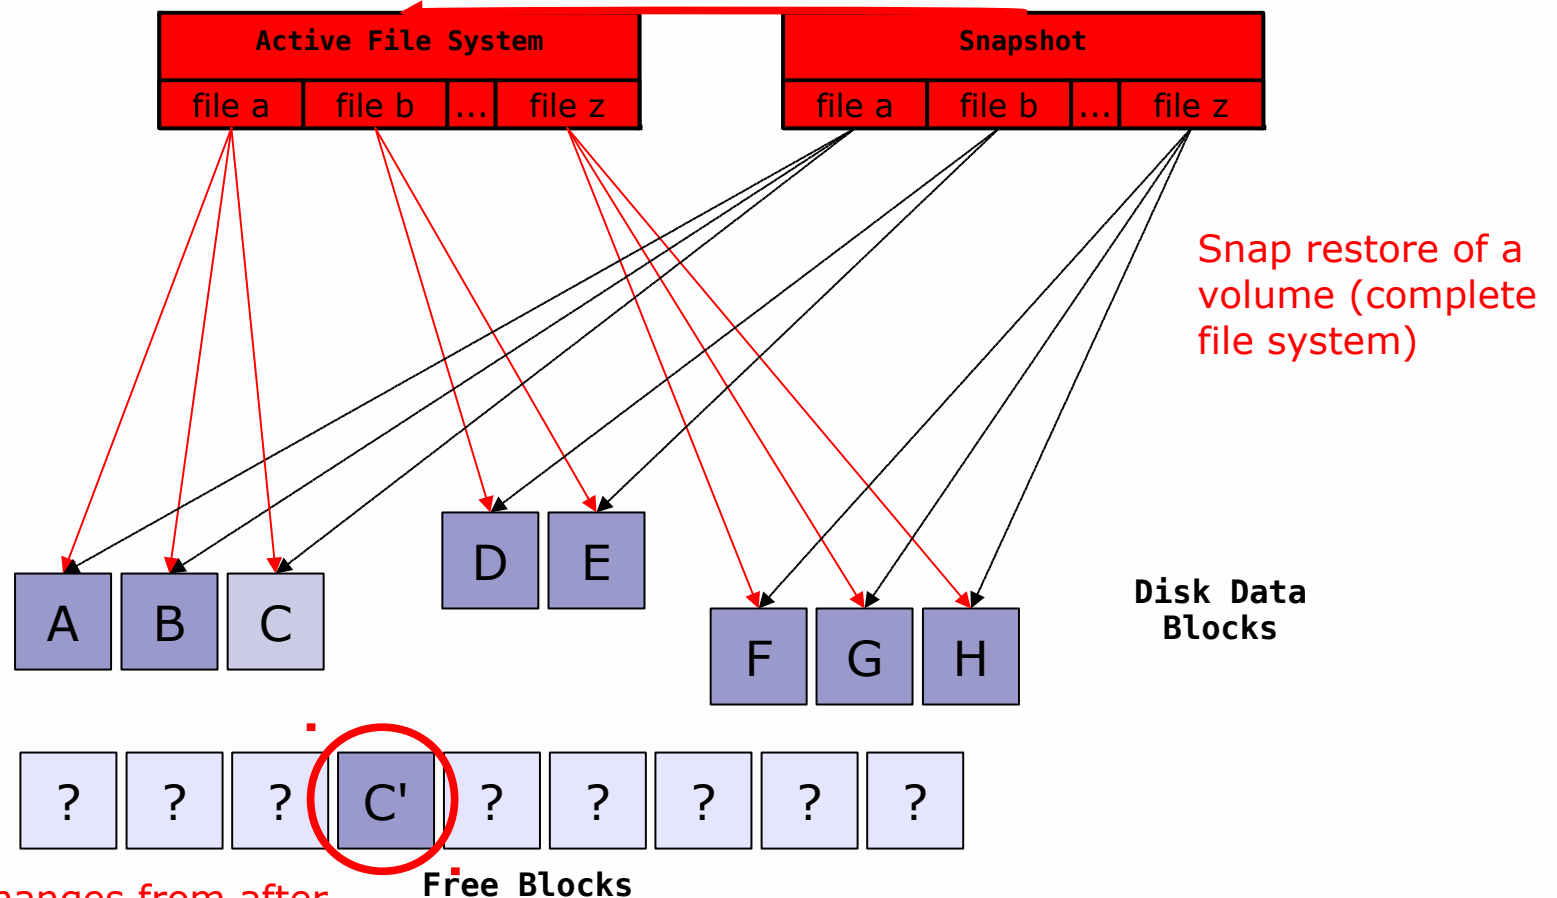
"Snapshots"



Snapshots (cont.)

A Bird's Eye View at Snapshots & SnapRestore

"Snapshots"



Snap restore of a volume (complete file system)

Disk Data Blocks

Free Blocks

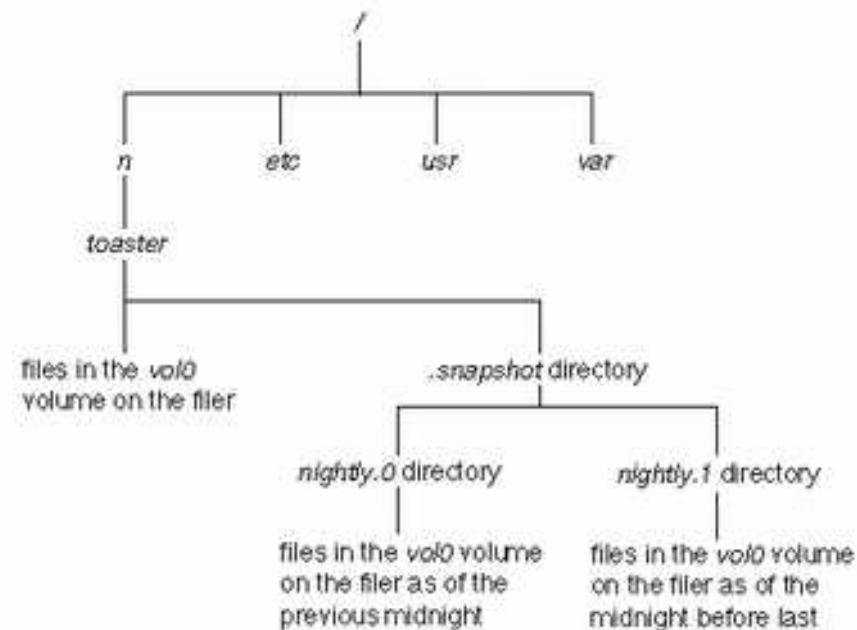
Any changes from after the restored file system (C') are irrevocably lost !

Snapshots (cont.)

Accessing Snapshots from Clients

NFS clients

- `.snapshot` directory



CIFS clients

- `~snapshot`, `~snapsht`, `.snapshot`

Snapshots (cont.)

The Problem of Consistent Snapshots

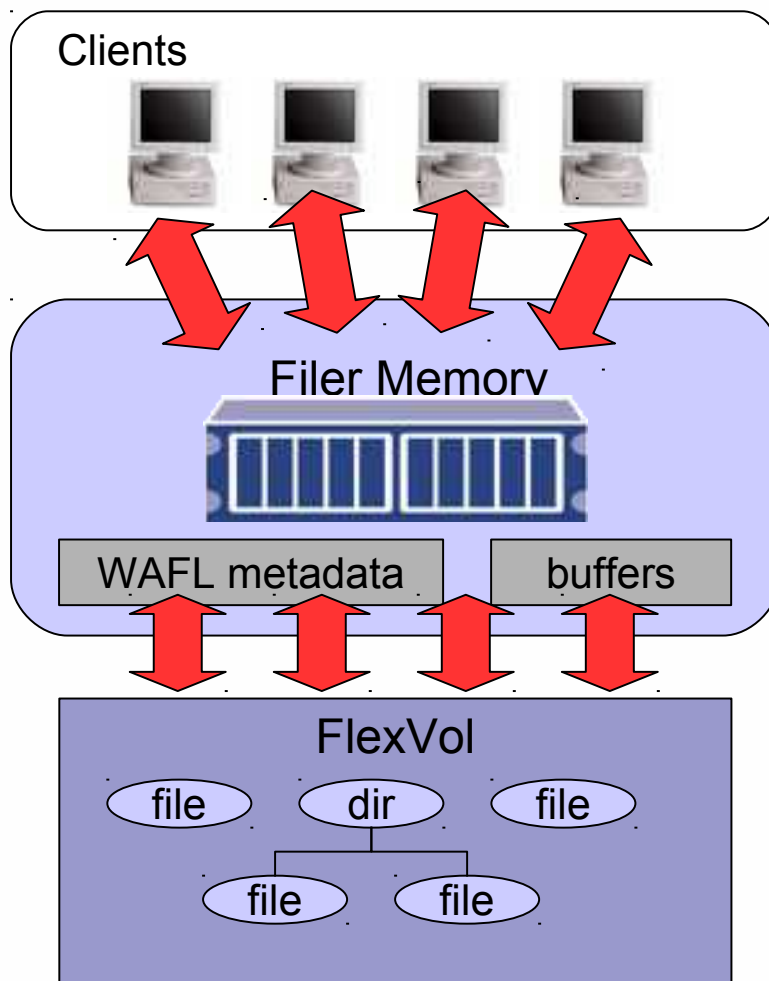
NAS clients modify files

The NetApp filer manages WAFL metadata and buffers in-memory

Eventually the modifications are written out to disk

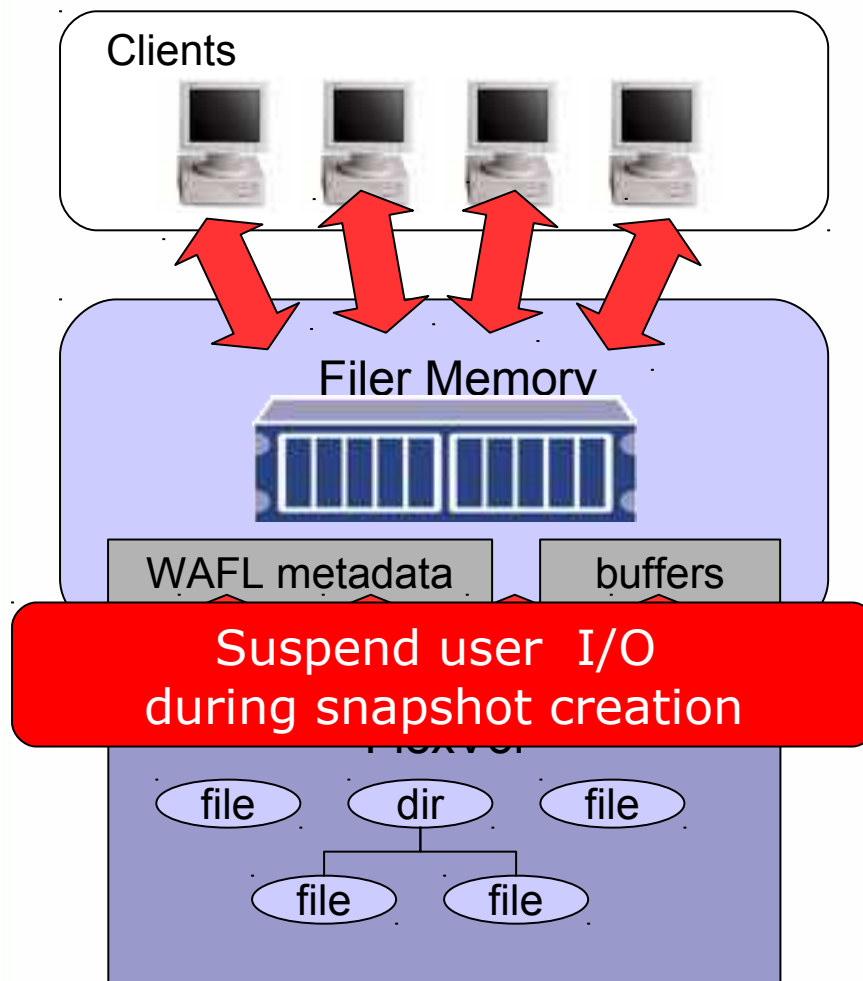
What happens when we take a snapshot of a flexible volume while clients are actively modifying files ?

Compare this problem with backup software: "Backing Up Open Files"



Snapshots (cont.)

The Problem of Consistent Snapshots (cont.)



During snapshot creation, the necessary buffers are flushed to disk, then user I/O is suspended to a volume

After snapshot creation, user I/O operations to the volume are resumed

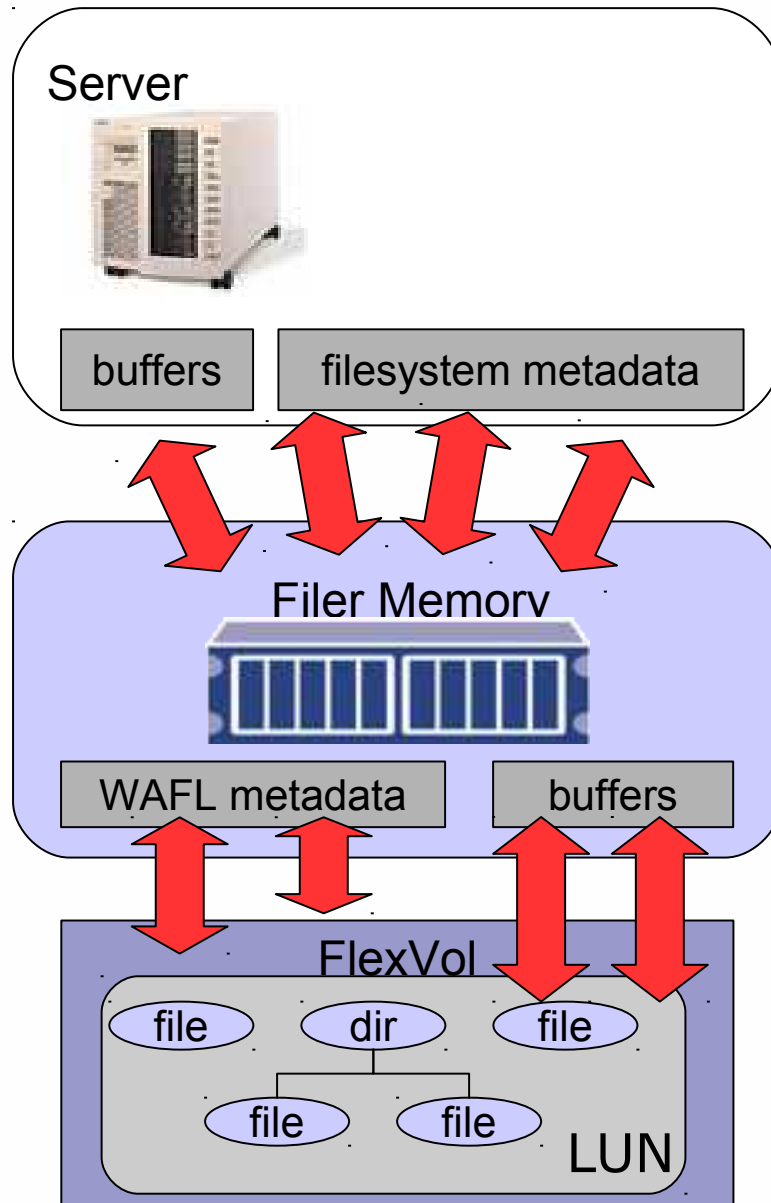
GOOD: WAFL will make sure volume meta-data is consistent on-disk

BAD: WAFL ignores any consistency issues in files (eg. "open Access databases", .pst files)

**Do we really care about this ?
No, this is the best we can do anyway**

Snapshots (cont.)

The Problem of Consistent Snapshots (cont.)



SAN complicates things!

A server will have its own buffers and filesystem metadata (eg. NTFS on Windows systems)

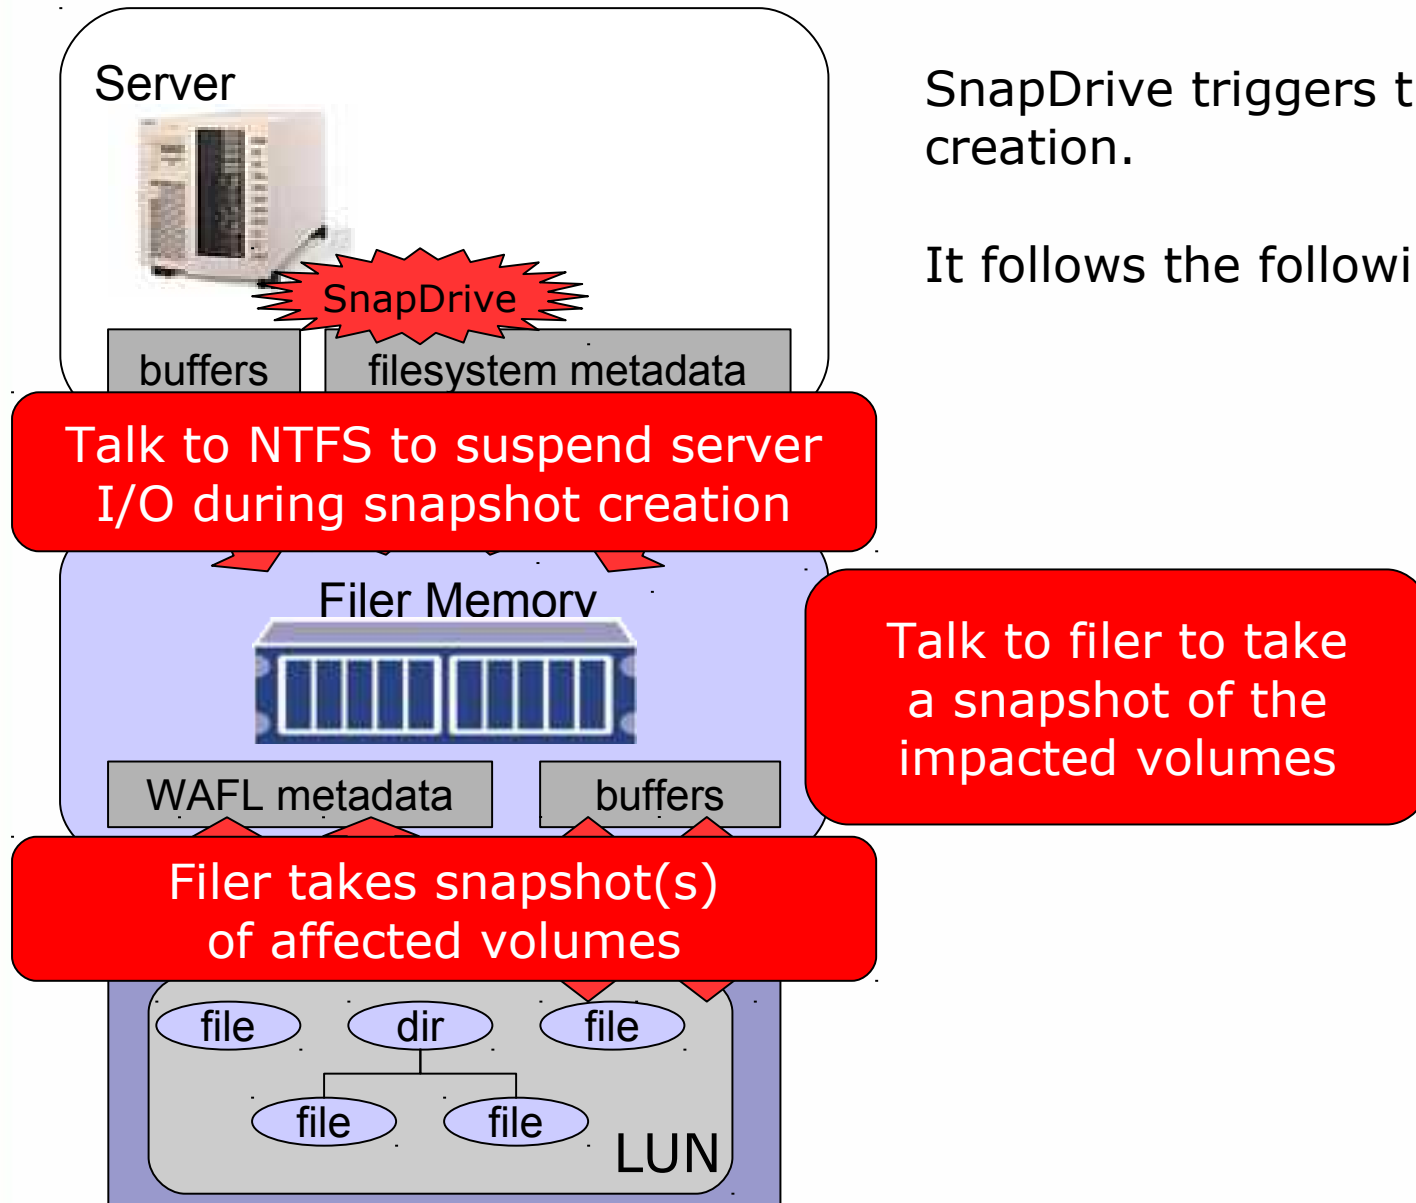
A WAFL-consistent snapshot will not necessarily be consistent from the server's filesystem's point of view, as filesystem metadata and buffers have not been flushed & synced correctly

Solution: we need some software (driver) on the server to talk to the local filesystem (eg. NTFS) and freeze/thaw it before/after a NetApp snapshot creation

SnapDrive !

Snapshots (cont.)

The Problem of Consistent Snapshots (cont.)

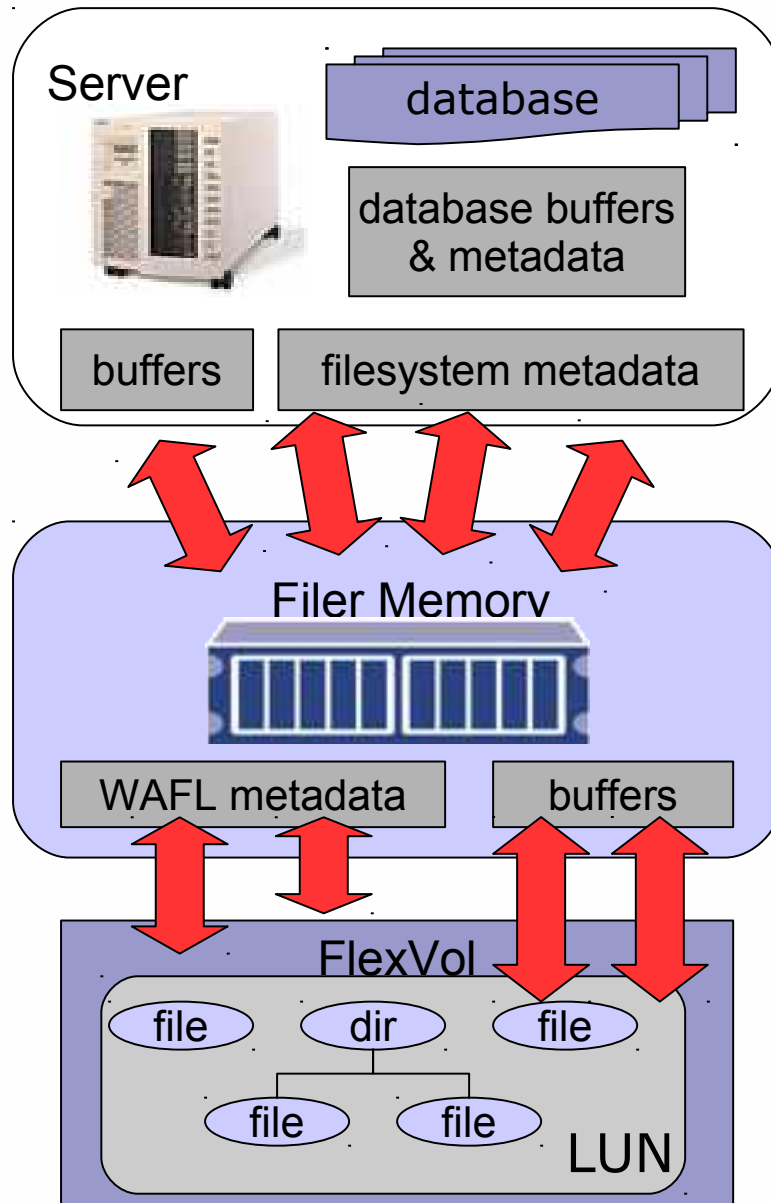


SnapDrive triggers the snapshot creation.

It follows the following steps:

Snapshots (cont.)

The Problem of Consistent Snapshots (cont.)



Running a database on your SAN complicates things even more if you want to take snapshot backups of your data

A special application that talks to the database's backup API is necessary

SnapManager !

Snapshots (cont.)

The Problem of Consistent Snapshots (cont.)

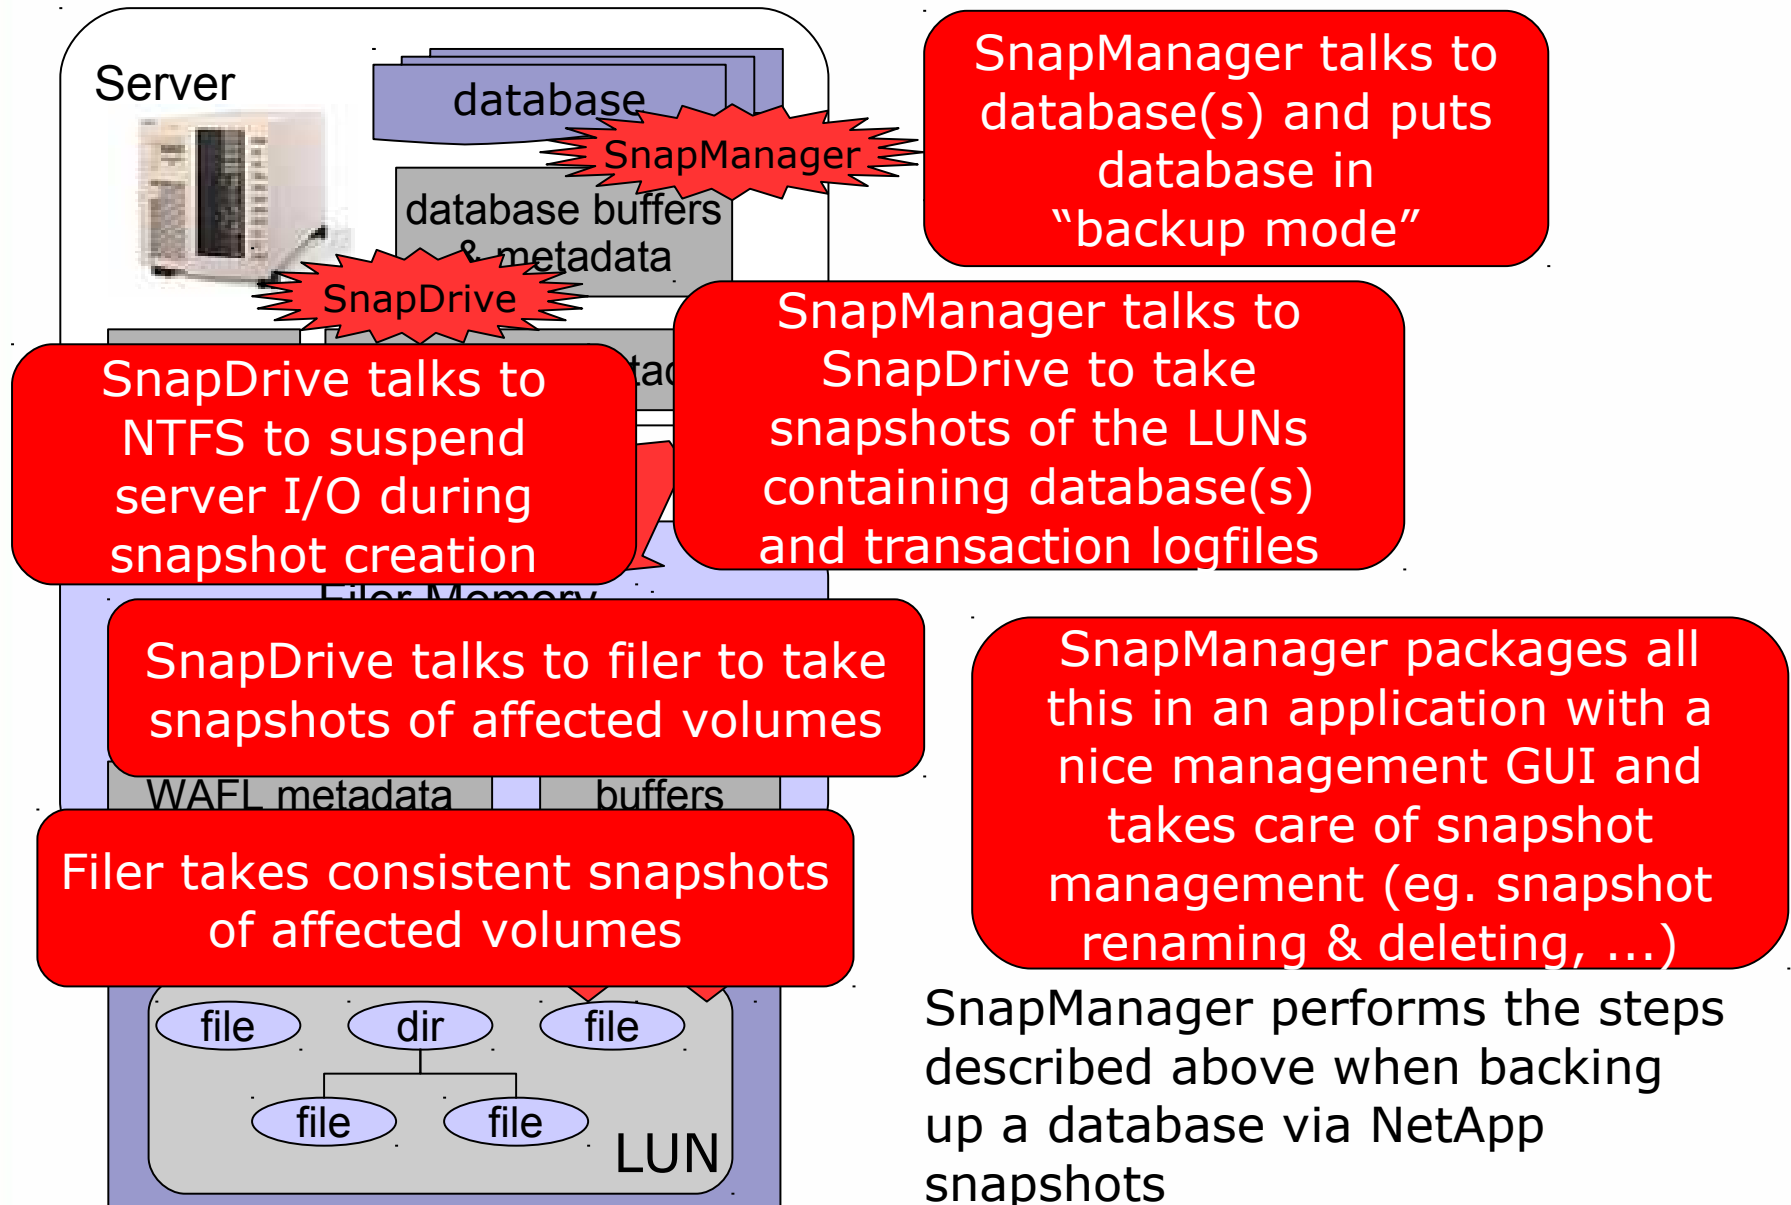


Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- **Qtrees**
- LUNs
- Network Configuration: VIFs
- Replication Technologies Overview



Qtrees

- = A directory with special properties
- Originally, Q = quota, “quota-tree”, can be used to set a quota on a particular directory
 - Nowadays, we have FlexVols, in a way already quota-limited
- Security style & oplocks settings can be different than rest of volume
 - Nowadays, we have FlexVols, can have different security styles & oplocks settings
- Less important now
- **BUT**: Still important when dealing with replication technologies:
 - SnapMirror = we can replicate whole volumes OR qtrees
 - SnapVault = we can only replicate qtrees
 - OSSV (Open Systems SnapVault) = we can only replicate directories to qtrees



Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- **LUNs**
- Network Configuration: VIFs
- Replication Technologies Overview



LUNs

- Look like big files on WAFL filesystem
- Are formatted and handled by host OS
- Mapped via FCP or iSCSI
- See SnapDrive & SAN course for more info



Table of Contents

- NetApp Products
- Storage Terminology
- Current NetApp Hardware
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- **Network Configuration: VIFs**
- Replication Technologies Overview



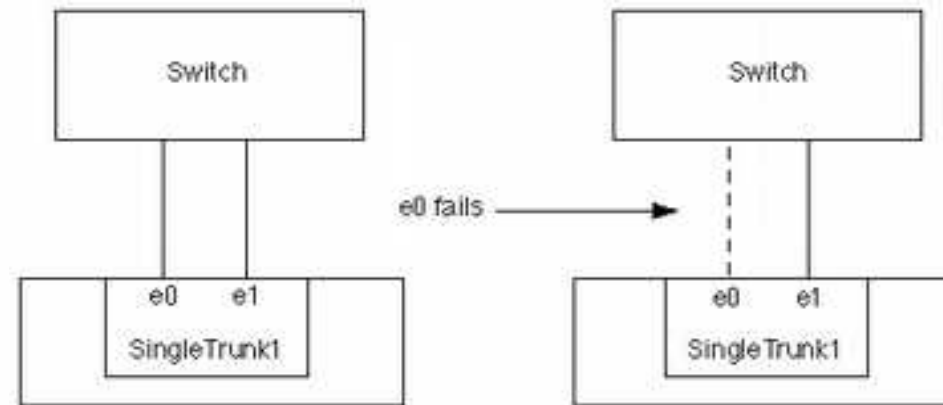
Network Configuration

- Give IP address to physical interfaces, or ...
- ... create VIFs and give IP address to VIF
- VIF = virtual interface
- 2 types (but can be stacked)
 - Single-mode VIF
 - 1 Active link, others are passive, standby links
 - Failover when link is down
 - No configuration needed on switches
 - Multi-mode VIF
 - Multiple links are active at the same time
 - Loadbalancing and failover
 - Loadbalancing based on IP address, MAC address, or round robin
 - Requires support & configuration on switches:
IEEE 802.3ad (static) (dynamic is supported on 7.2)



Network Configuration (cont.)

Single-mode VIF:



Multi-mode VIF:

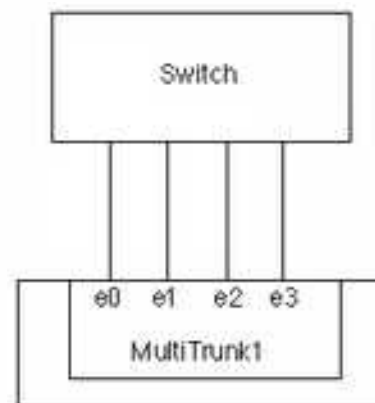


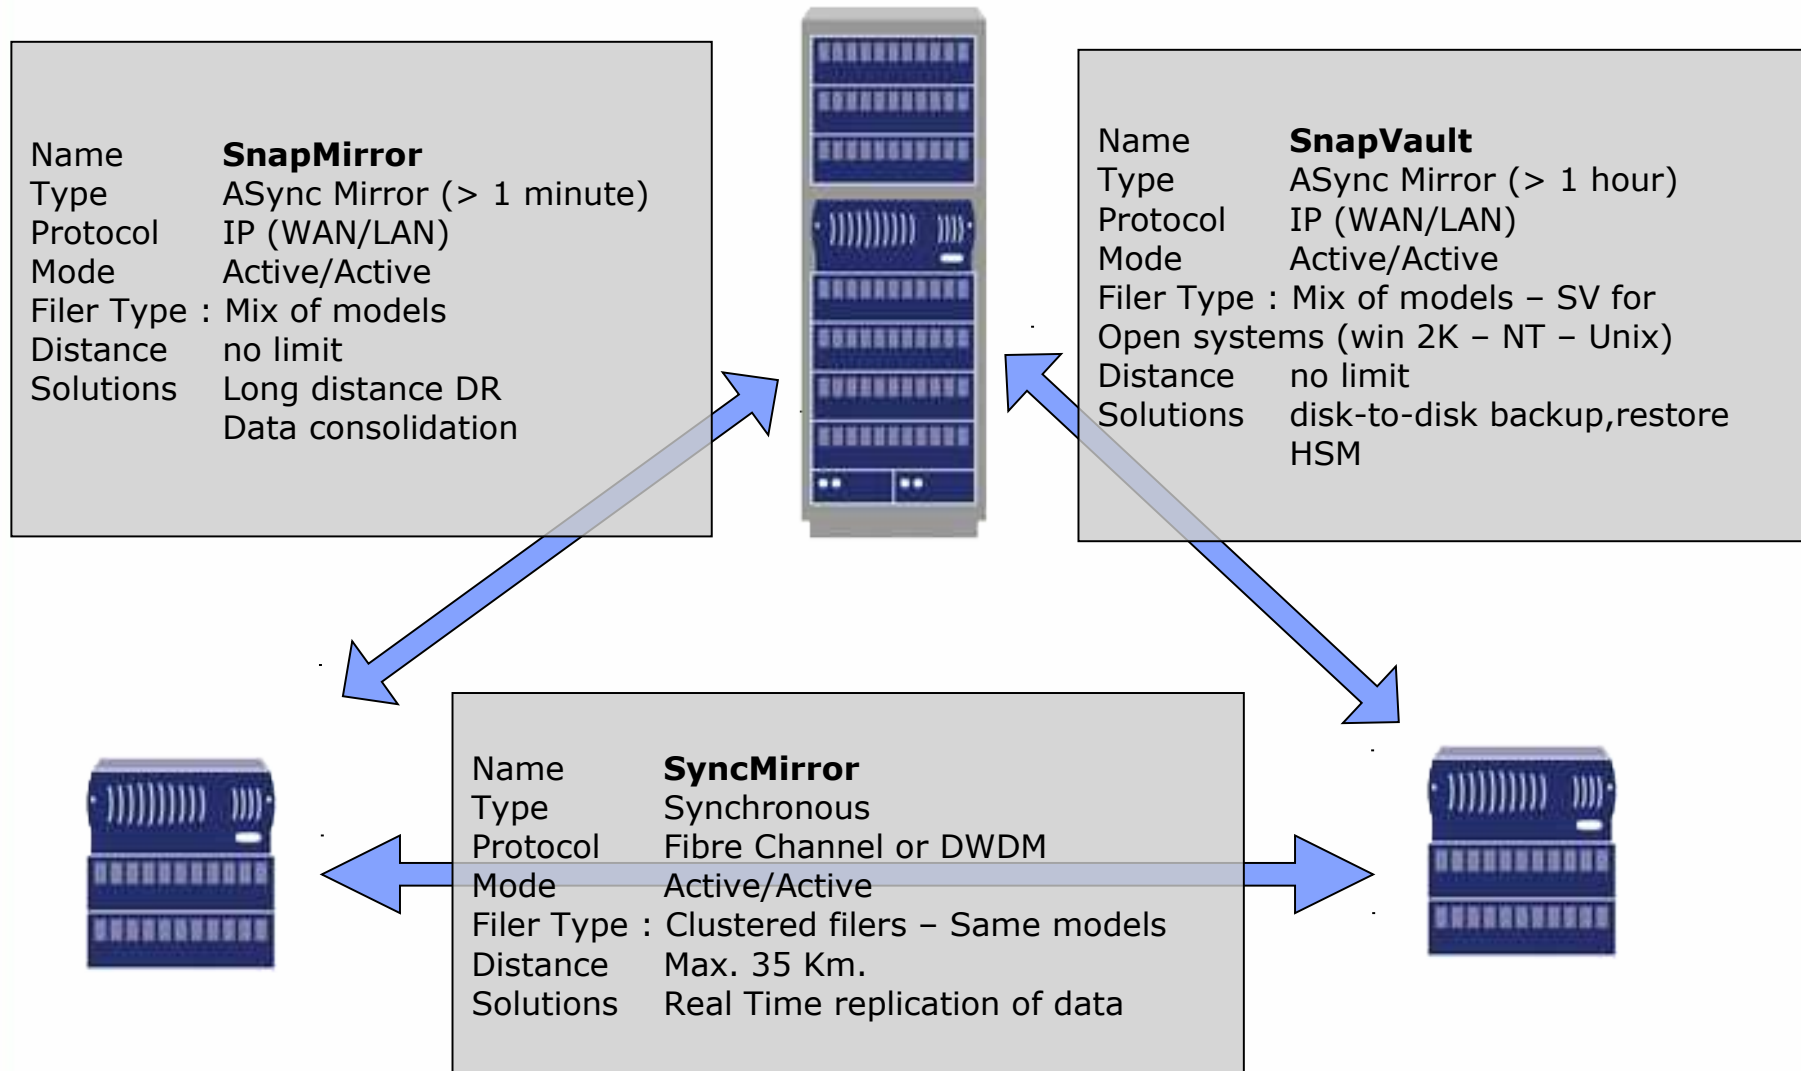
Table of Contents

- NetApp Products
- Storage Terminology
- NetApp Terminology
- NetApp Hardware Essentials
- Where to Get Help ?
- Disks, RAID4, Aggregates and Space Calculation
- When Disks Go Bad ...
- Volumes
- Snapshots
- Qtrees
- LUNs
- Network Configuration: VIFs
- **Replication Technologies**



Replication Technologies

SnapMirror, SnapVault (and OSSV), SyncMirror



Replication Technologies (cont.)

This is about ...

- SnapMirror (Disaster Recovery)
- SnapVault (Backup/Restore)
- SyncMirror is not a replication technology (=data redundancy)

Overview

- SnapMirror
 - Disaster Recovery Technology (get backup site online ASAP)
 - Source & destination must be NetApps (both need a "SnapMirror" license)
 - Asynchronous replication (= periodic updates)
 - Pull based (destination contacts source and asks for changes)
 - Replicate:
 - Volumes
 - Qtrees
 - There is a variant called synchronous SnapMirror (push



Replication Technologies (cont.)

Overview (cont.)

- SnapVault
 - Backup technology (restore data from a remote location)
 - Source & destination must be NetApps (Source needs "SnapVault Primary" license, destination needs "SnapVault Secondary" license – you need two different NetApps !)
 - Asynchronous replication (=periodic updates)
 - Pull based (destination contacts source and asks for changes)
 - Replicate:
 - Only qtrees can be snapvaulted (data must be in qtrees!)
- OSSV (Open Systems SnapVault)
 - Backup technology, see SnapVault
 - Source must be a server (Windows, Unix, Linux)
 - Destination must be a NetApp
 - Backs up directories to qtrees
 - Licenses must be installed on NetApp: "SnapVault Secondary" + "SnapVault Primary for Windows/Unix/Linux"



Replication Technologies (cont.)

Overview (cont.)

- SyncMirror
 - Not a replication technology
 - Allows to synchronously mirror an aggregate
 - Used eg. in Metrocluster for data redundancy



Replication Technologies (cont.)

Images Used



NetApp
Filer




Servers (Windows,
Unix, Linux)



Volume
(with
snapshots)



snapshots



Volume
(with
qtrees &
snapshots)



Replication Technologies (cont.)

Images Used (cont.)



SnapMirror



Synchronous SnapMirror



SnapVault



OSSV (Open Systems SnapVault)



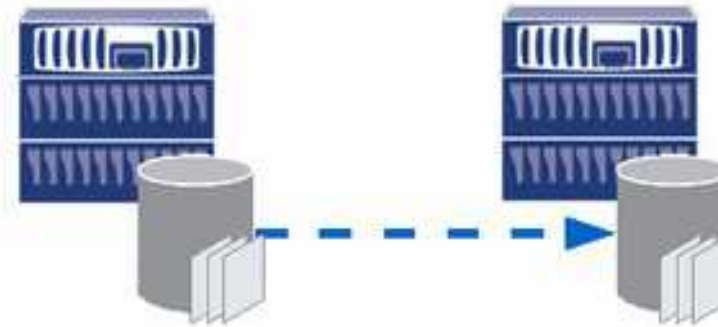


SyncMirror




Replication Technologies (cont.)

Volume SnapMirror (VSM)




Replication Technologies (cont.)

Volume SnapMirror (VSM) (cont.)





Replication Technologies (cont.)

Qtree SnapMirror (QSM)




Replication Technologies (cont.)

SnapVault




Replication Technologies (cont.)

OSSV (Open Systems SnapVault)




Replication Technologies (cont.)

SnapVault & OSSV Combined



Replication Technologies (cont.)

Only Valid 3-tier Backup & Disaster Recovery Design



Replication Technologies (cont.)

SyncMirror



Drawing is not 100% accurate, there is no replication, data is written simultaneously to two locations

Metrocluster:
cluster controllers
& make data fully
redundant +
geographically
spread out

