Reliable Storage Servers

Peter Ashford

Ashford Computer Consulting Service

8/24/2017

Intended Audience

The intended audience for this presentation is people who perform one or more of the following tasks on storage servers:

- Specify
- Assemble
- Install
- Configure
- Administer
- Maintain

What We'll Cover

This presentation covers selection and configuration of Hardware, Software and File-Systems for a large high-reliability storage server.

Several different methods are described for improving reliability. Select the appropriate methods for the data.

These methods are being described for use with stand-alone file servers, including those built in-house, but they apply to all types of storage server.

Why Stand-Alone File Server?

Many organizations use clustered storage servers (scale-out) and multi-head SANs. Why should they care about large stand-alone file servers?

- The actual storage in both of these is similar to the large file-servers discussed here.
- Large file servers can function as storage nodes for a clustered storage server.
- Large file-servers can function as a SAN by creating Files/Volumes and sharing them out through iSCSI or Fiber Channel.

Background

Direct Attach Storage - DAS

The drives are directly attached to the application server.

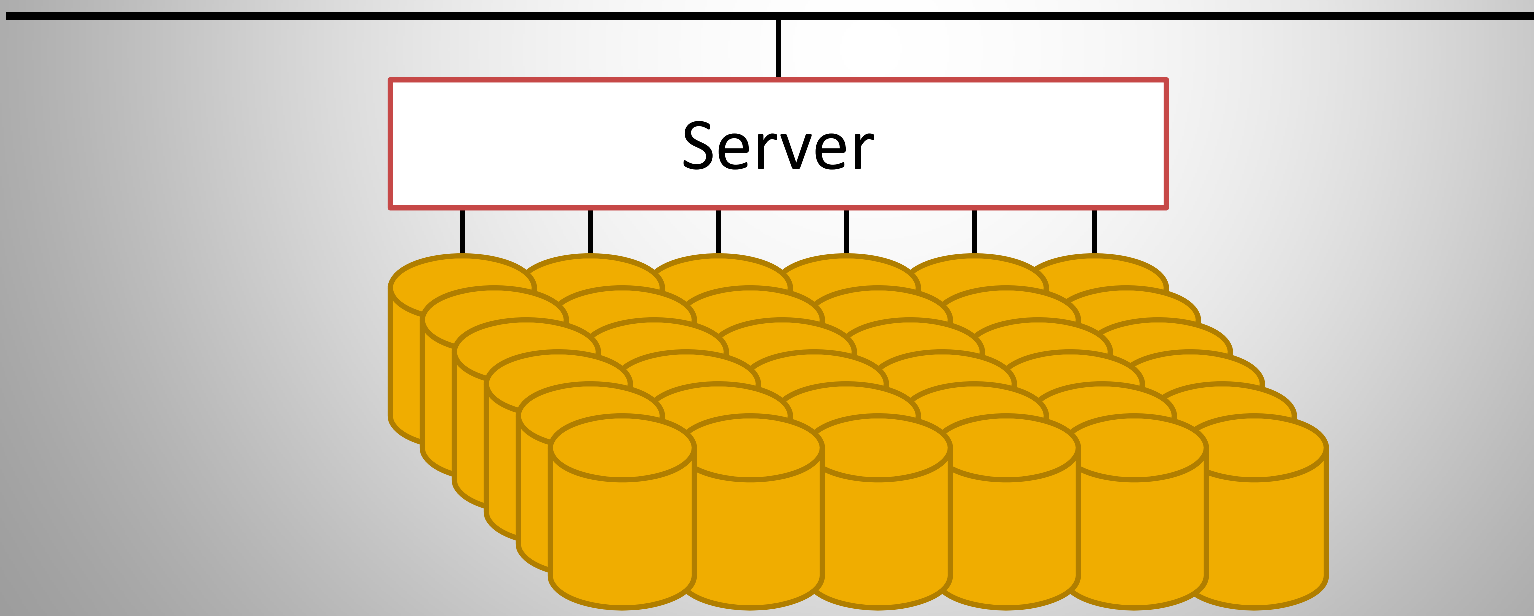

Background

Network Attach Storage - NAS

The drives are directly attached to a computer that delivers files to the application servers and clients.

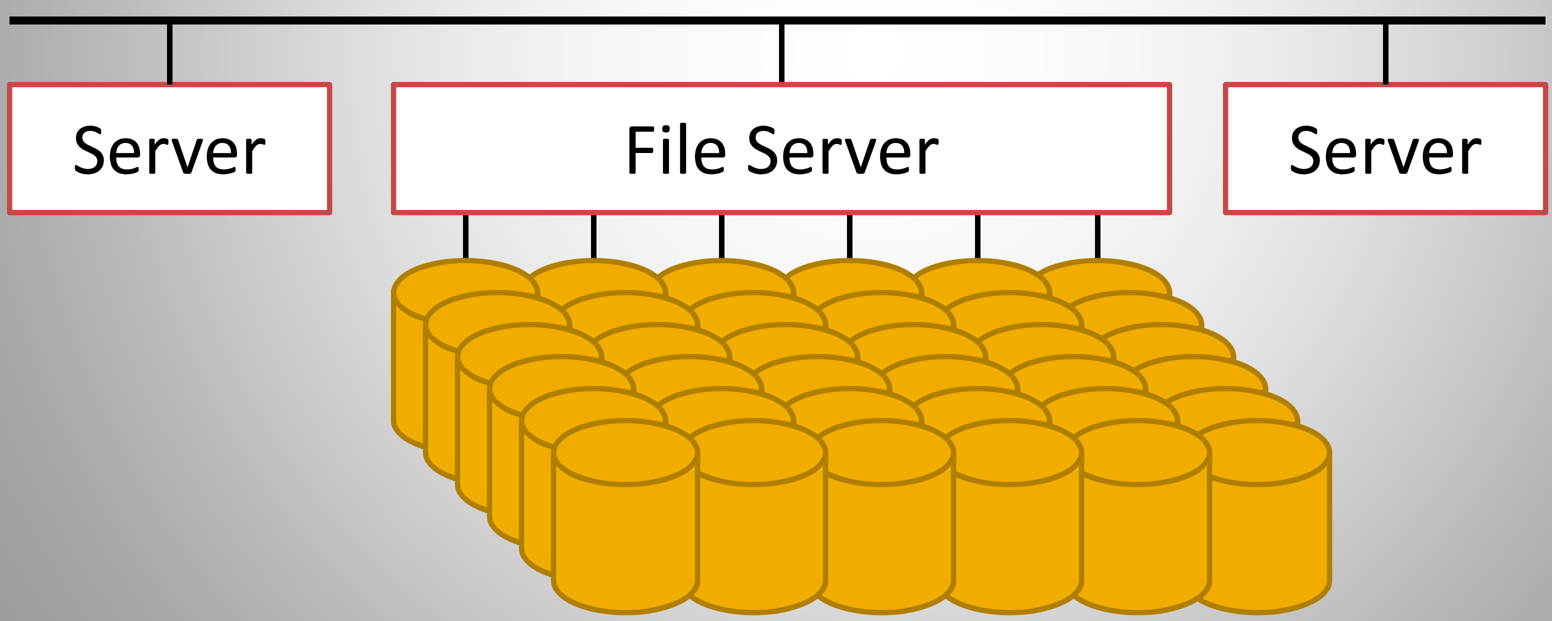

Background

Storage Area Network - SAN

The drives are directly attached to a computer that delivers blocks to the application servers.

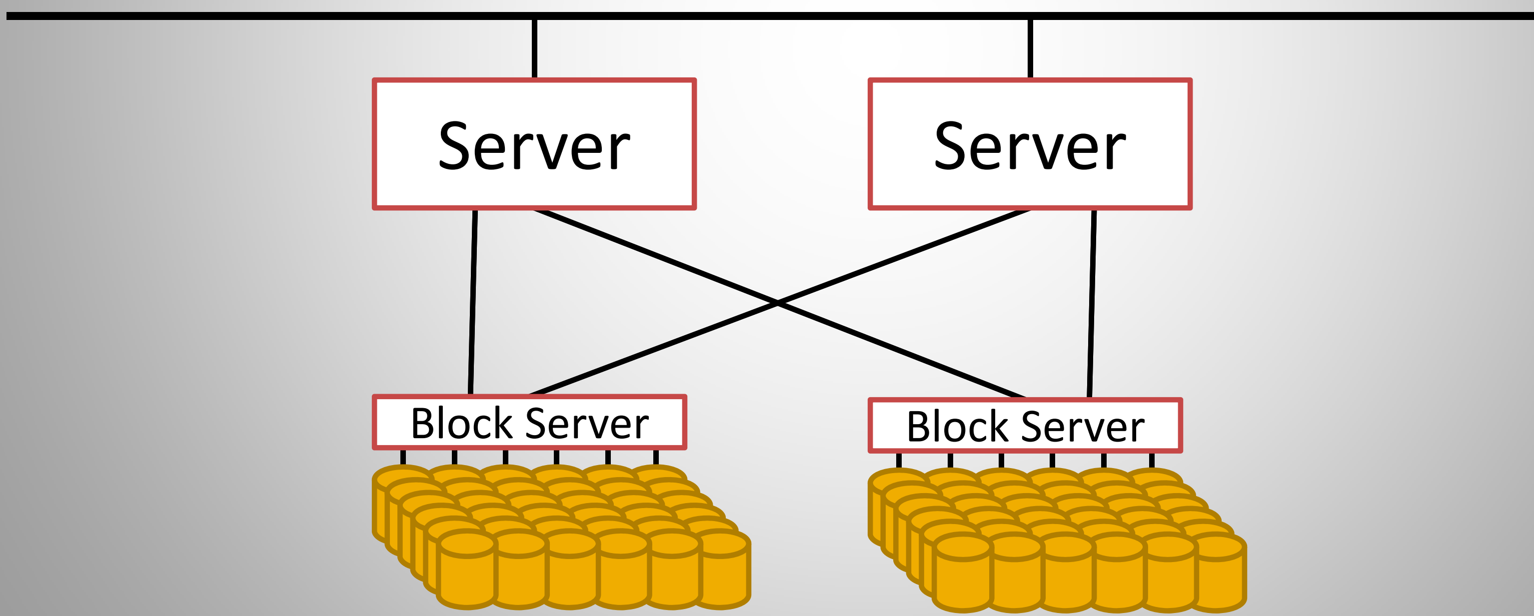

Background

Clustered Storage Server

The drives are directly attached to multiple computers that deliver data to the application servers.

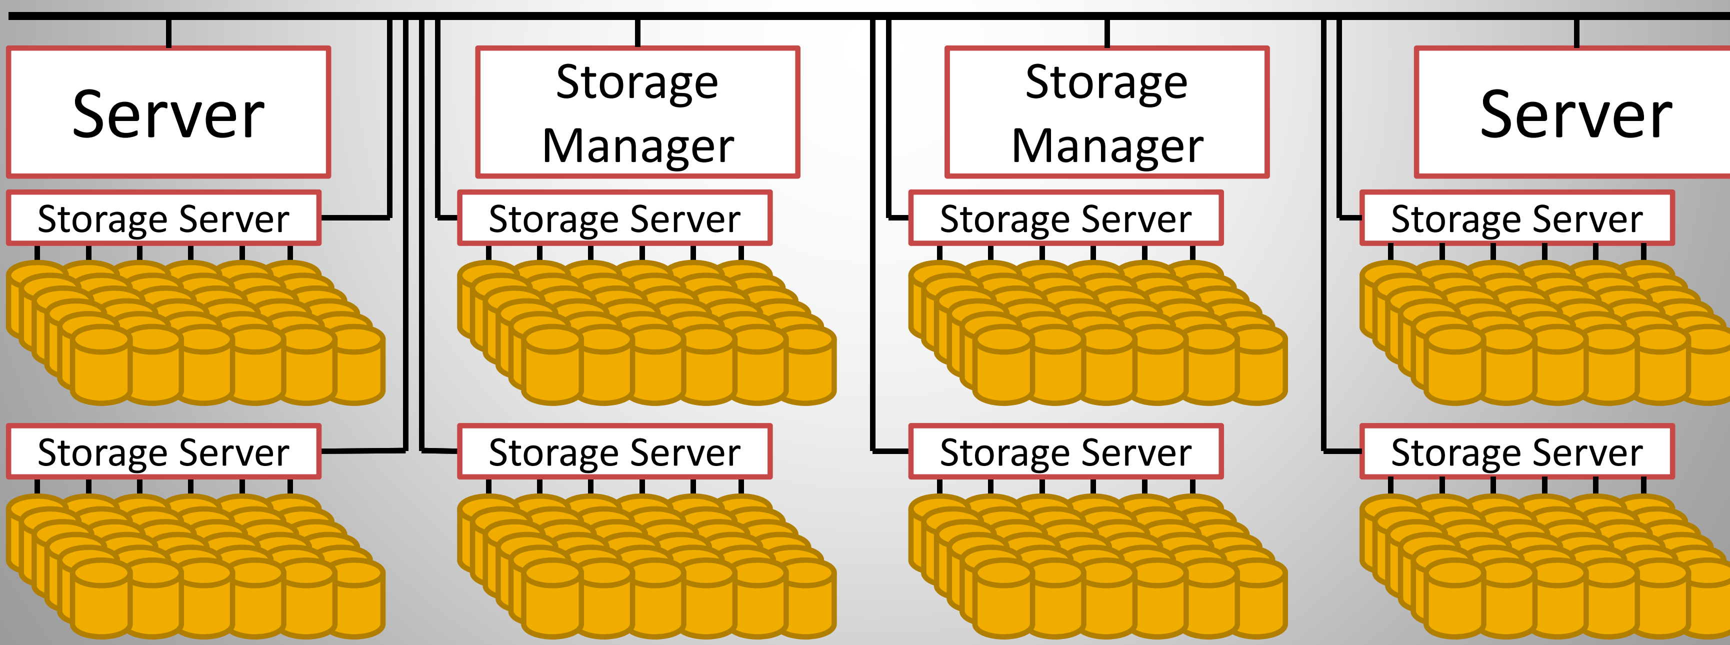

Background

Scale-up vs. Scale-out

- “Scale-up” is used to describe “Larger”. This is one of the focuses of this presentation.
- “Scale-out” is used to describe “More”. While this presentation doesn’t directly address Scale-out, it does address the Storage Nodes that are added to a Scale-out system.

Background

	% Availability	Downtime per Year
Two Nines	99%	3.65 days
Three Nines	99.9%	8.76 hours
Four Nines	99.99%	52.56 minutes
Five Nines	99.999%	5.26 minutes
Six Nines	99.9999%	31.5 seconds
Seven Nines	99.99999%	3.15 seconds

Storage Silos

Isolate groups of data

- Reduces impact of outages
- Reduces cost of implementation and maintenance
- Be careful to not create Information Silos

Different data has different requirements

- Availability - Server uptime
- Redundancy - RAID level and Replication
- Speed - Disk vs. SSD - Caching
- Backups

Backups

How important is a backup?

- Regulation/Law
- Contract requirements
- Cost of unavailable data
- Cost to replace/replicate
- Cost to lose
- Importance to Business mission

Backups

Where do I backup to?

- Local replication
- Remote/offsite replication
- Local tape
- Offsite tape

For offsite replication

- Uplink speed - 1TB/day ≈ 100Mbps dedicated - could be higher for less bandwidth-efficient protocols

Server Based Storage

When comparing Server Based Storage (SBS) to traditional storage systems (dedicated SAN systems), the following differences should be considered:

- Faster refresh cycle
- Use of commodity parts decreases cost
- SBS usually has a lower performance per unit, but is more cost effective
- When used as DAS:
  - Ability to run applications very close to the data
  - Non-storage-optimized architecture & OS

Checksumming File-Systems

With the large data storage requirements needed today, the chance of encountering an error in the data returned by the drive is becoming significant.

To date, the best method for identifying and correcting these errors is a file-system that performs checksums on the data before it's written and verifies the checksum after the data is read.

Step 1

Select an Operating System.

Operating System Requirements

Select an Operating System that you have adequate support for. Make sure that the Operating System has support for at least one checksumming File-System.

- FreeBSD - ZFS
- Linux - ZFS, Btrfs
- Solaris - ZFS
- Windows - ReFS on top of Storage Spaces

Of the above, ZFS is usually the most robust choice for storage.

Minimize

- Reduce the number of applications/packages that are installed on the file-server.
  - Reduces the installation & update time.
  - May reduce attack surface of server.
  - May reduce system boot time.
- Reduce direct logon access.
  - Reduces potential for privilege escalation.

Step 2

Select Hardware.

Select Compatible Hardware

Most Operating Systems have a Hardware Compatibility List (HCL). This should be checked to verify that the hardware you'd like to use is completely supported.

The Internet is a good source of Compatibility information. Sometimes a device is in the HCL, but the support is poor, or there are issues with the device. These should be avoided, unless they're known to have been corrected.

Select Reliable Hardware

Select hardware that's known to be reliable.

- For this determination, we should consider Uptime, as a percentage. The techniques presented here will provide uptime in the range of 99.99% to 99.999%, when starting with reliable hardware components.
- While there are many opportunities for cost savings on large file servers, unless you have many Peta-Bytes online, it's usually false economy.
- Most large file servers store mission-critical data. This is not a good place to try to save a few dollars.

Select Serviceable Hardware

Select hardware that can be serviced.

For turnkey systems, make sure that the manufacturer has technicians available in your area. Waiting for a technician to be flown out to you adds to downtime, which usually translates to money lost.

For systems built in-house, it's probably wise to keep at least a few critical spares.

Select CPUs

File-systems that support checksums are relatively new, and are not yet fully optimized for the large number of threads that current high-end CPUs have. This is because there are still significant portions of the file-system code that are single threaded.

Because of this, when selecting a CPU, lower latency and higher throughput will be delivered with high clock-rate CPUs. This is usually only available with lower core counts.

Big Memory

File servers use memory to cache reads and buffer some kinds of writes. To adequately perform this buffering, the system needs to have large amounts of memory. For a 1PB file server with multiple 10GbE connections, 256GB is not an unreasonable amount of memory, and 512GB might not be a bad choice, depending on the data access patterns.

If the data is stored on SSDs, this memory could be cut, as a miss on the cache is less costly with SSD storage.

Out of Band Management

In today's Data Center, it's necessary to closely monitor all aspects of a server. Some of the data is easiest to access through an Out of Band Management device (e.g. DRAC, IPMI, LOM).

This also allows the system to be powered up remotely, and provides remote access to the system console, reducing the time to resolve a problem.

Redundancy

The second step for selecting hardware is to design for redundancy. Redundant hardware allows the system to continue operating, even though there are failed devices in the system.

Redundancy - Power

Redundant power supplies became common in servers many years ago. Please don't forget to specify them.

If you have several of the same system on-site, you might want to keep a spare power supply module.

Redundancy - Fans

Systems with redundant fans are relatively rare. The fan controller is usually programmed to spin all fans up to full speed (or at least a higher speed) when one of the fans fails.

This feature can reduce the chances the server will overheat due to a single fan failure.

Dual CPUs

Modern file servers are expected to perform calculations for RAID Parity, Checksum, Encryption, Deduplication and Compression, in addition to the need to track users, directories, files and blocks. In order to achieve the high throughput/low latency demanded in today's Data Center, this will often require that dual CPUs be used.

Redundancy - Drive Controller

Redundant drive controllers are useful to improve reliability, in case of a drive controller or cable failure.

Multiple drive controllers also distributes the I/O load across more PCI-E channels.

Redundancy - Drive Channels

When dealing with large file servers, more channels from the server is better. It keeps the drive tray chains shorter and decreases the overall impact of a failed cable.

Redundancy - Dual-Channel Drives

Drive trays use SAS Expander chips to allow a small number of incoming channels to communicate with many drives.

When an Expander chip has a problem, it can take one or more drives off-line. Use of dual-channel drives (SAS), with dual Expander chips, will improve system reliability.

Redundancy - Network Interface

To improve the reliability of the network connection, and to get access to the bandwidth that a file server can use, it's usually necessary to use multiple network connections. Multiple 10GbE connections isn't unusual, and multiple 25GbE or 40GbE channels are becoming available. Be sure the switches are redundant too.

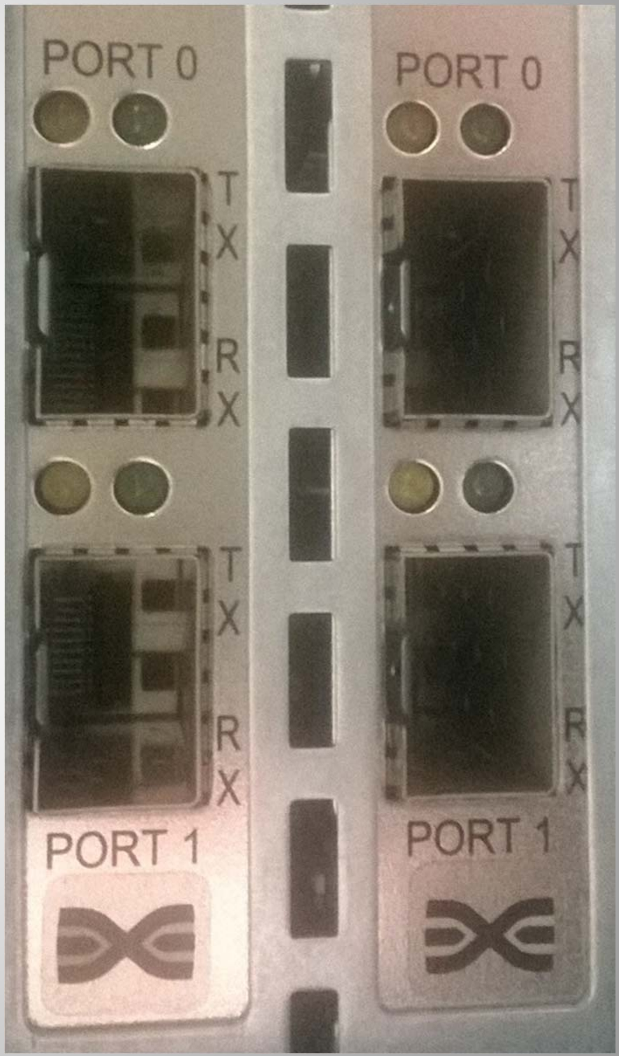

Redundancy - RAID

Due to the large drives that are currently used in most file-systems, it is strongly suggested that a dual-parity RAID level be used. These are RAID-6 and RAID-Z2 (ZFS).

In addition, ZFS has a triple-parity RAID level (RAID-Z3) available. Calculating triple-parity is significantly more complex than calculating dual parity, so this level of protection should be considered only for information that’s very critical, and where the I/O loads are low.

Redundancy - RAID

Dual-parity RAID is critical for two main reasons.

1. It's possible to lose as many as two drives from each array in the system, and still have the data available.
2. When a drive fails, the array is still redundant, allowing bad blocks to be corrected during the reconstruction of the array. With the size of current drives, combined with the amount of data that must be read to perform a reconstruct, the chances of bad data are too high.

Redundancy - Dual Head

The ultimate in redundancy is to have a storage system with dual head nodes. This allows the second head to serve the data if the first head fails. This is called an Active/Passive configuration.

To get a better ROI, it's possible for both nodes to be serving data from different sets of drives in the connected storage, while performing as a secondary server for the drives it doesn't normally serve. This is a dual Active/Passive configuration.

Select Drives

Selecting the Drives is usually simple, but it can be complex in certain extreme cases.

- For frequently accessed files, use SAS SSDs.
- For data archival, including Virtual Tape and backup storage, SAS spinning disks is usually the best choice.
- For clustered file-systems, use SATA.

NOTE: Large SAS SSDs cost about 10 times what a similar sized SAS spinning disks cost. If the customer/user specifies SSD, be sure they're aware of the cost.

Step 3

Purchase the Hardware.

Disk Drives

If purchasing the spinning disks for local integration, the best way to purchase the disks is by the case.

When dealing with spinning disks, the less handling of the drives the better.

Step 3

Configure the Hardware.

Power

Stable power is critical to reliable servers. To achieve this, the following are needed:

- Surge suppression
- Power conditioning
- Battery backup power
- Generator backup power (optional)

Use 240VAC for higher efficiency.

Cooling

As with power, good cooling is critical to a reliable server.

This is an example of a custom modification to improve CPU cooling. By redirecting the heat from one CPU to not enter the other, the CPU temperatures are lower, generating a more reliable system.

Configuration - Drive Tray Cables

There are a few different drive tray cabling topologies. The methods shown at the left will be discussed individually.

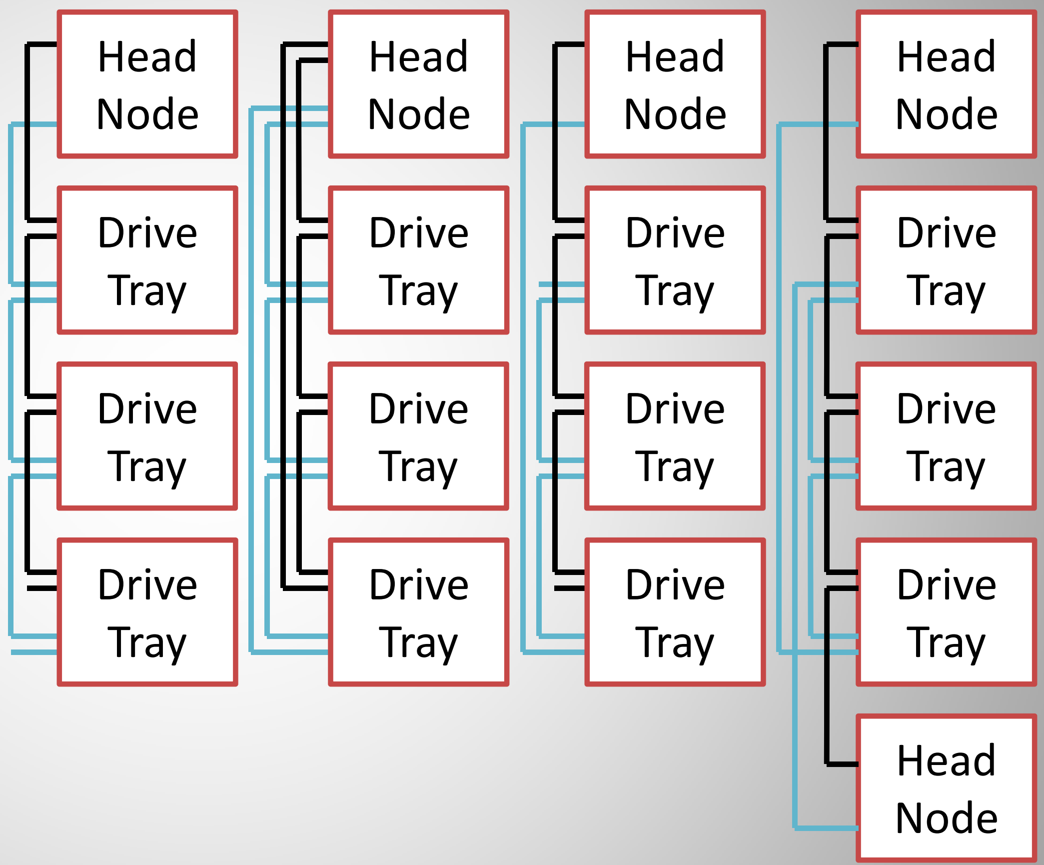

Drive Tray Cables - Traditional

This is the traditional topology for connecting multiple drive trays to a head node. Both the "A" and "B" channels connect to the drive trays in series.

For improved reliability, a second head node can be connected to the end of the drive tray chain.

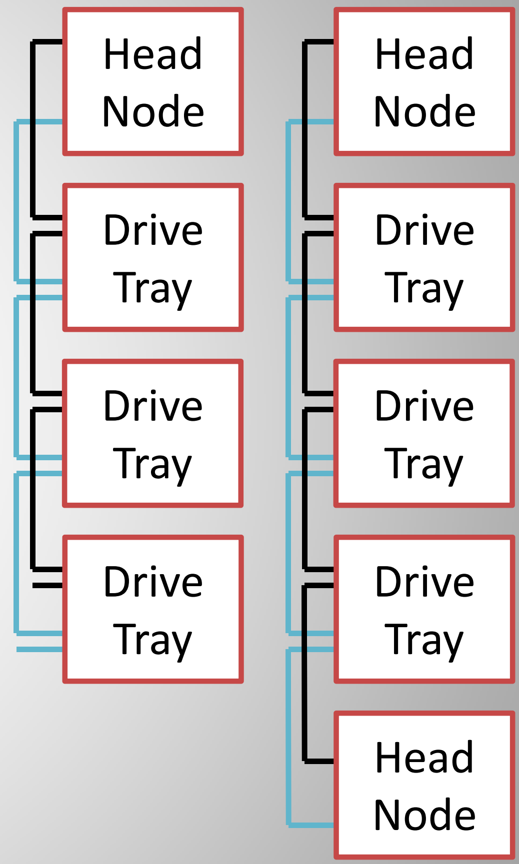

Drive Tray Cables - Loop

This is similar to the Traditional topology, with the addition of a cable from the end of the drive tray chain back to the head node.

In the event of a complete drive tray failure, this additional cable gives continued connectivity to the down-stream drive trays.

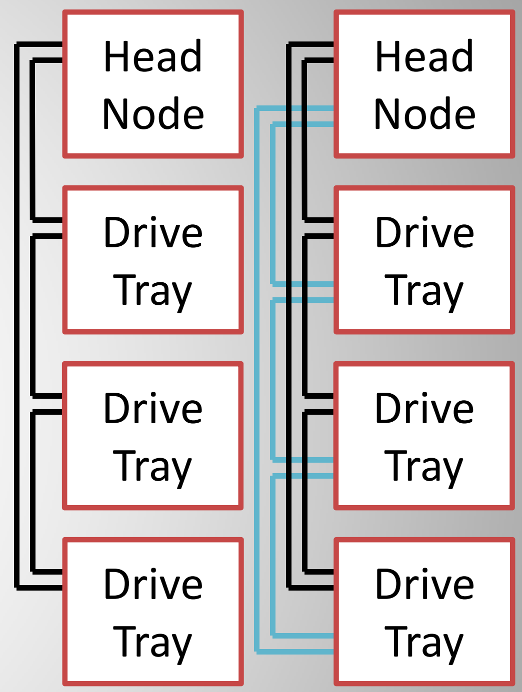

Drive Tray Cables - Improved

This is an improved topology for connection of multiple drive trays to a head node. The “A” channel connects to the drive trays in series, while the “B” channel connects to the drive trays in the opposite order.

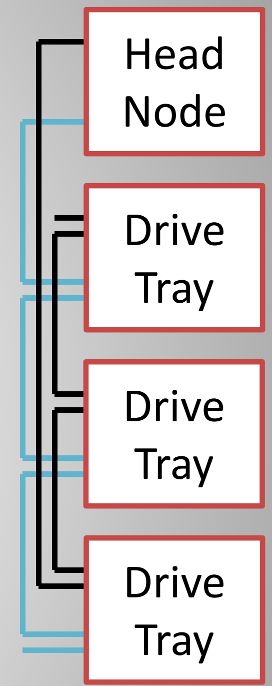

This topology provides the reliability of the Loop topology, while using fewer drive channels.

Drive Tray Cables - Dual Head

This connection topology has the advantages of the improved topology previously discussed, with the addition of a second head node.

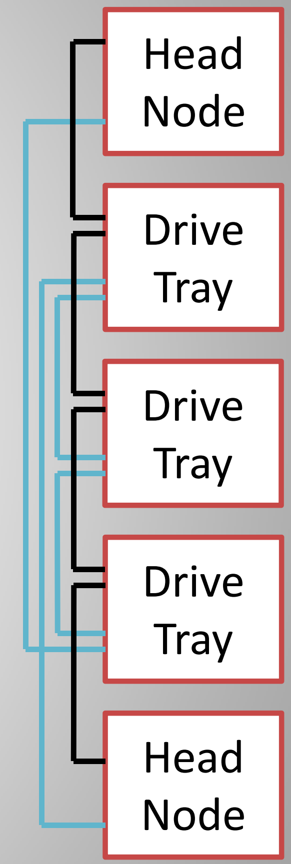

Drive Layout

If at all possible, the drives should be installed into the drive trays so that no drive tray has more than one drive in any array. When this suggestion is followed, complete loss of an entire drive tray will degrade many arrays, but all arrays will remain functional. If a dual-parity RAID level was used, all arrays will remain redundant.

Network Connections

The multiple network connections should be bonded together on the server, and the switch ports should be set to LACP Passive mode.

This will allow the server to receive and deliver additional bandwidth to the network. It will also provide redundancy in the event of a failure.

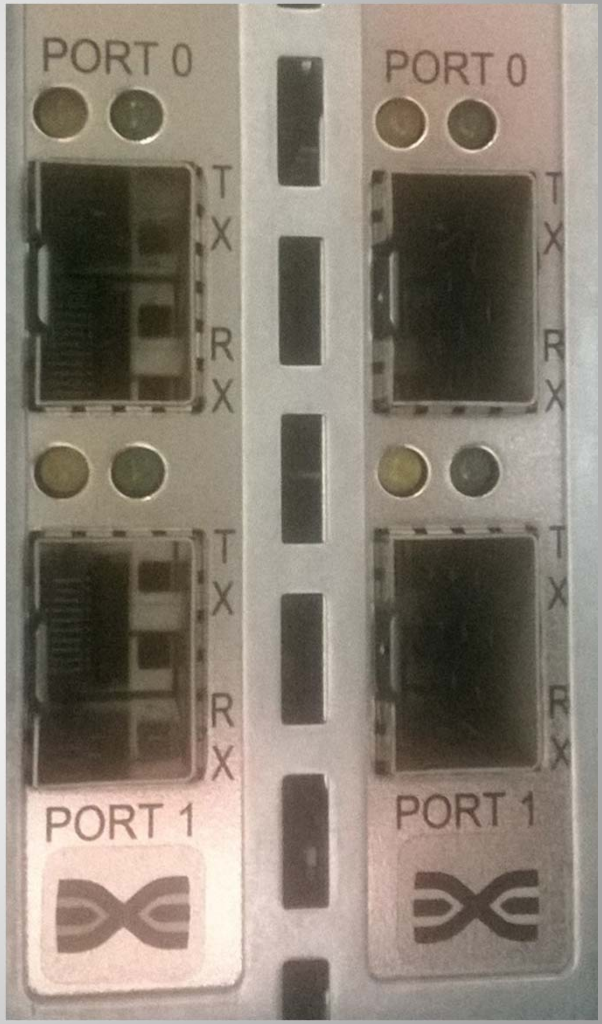

Step 4

Configure the Software.

This section is specific to dual-head configurations.

Dual Head

Getting a dual-head system up and running, in either an Active/Passive, or dual Active/Passive, configuration requires effort beyond that which is normal for a file-server. Here are a few links that could be used to configure an Active/Passive server under Linux:

- https://github.com/ewwhite/zfs-ha/wiki
- http://zfs-create.blogspot.dk/2013/06/building-zfs-storage-appliance-part-1.html
- http://v-optimal.nl/index.php/2016/02/04/ha-zfs/
- http://www.high-availability.com/zfs-ha-plugin/

Example

We'll use this example to tie everything together.

For the sake of simplicity, let's assume that you're the System Administrator, and there are two projects, each individually funded, both of which need a storage server. Both projects are interested in extreme performance over NFS (25GbE) and FC (16Gb).

As neither Dell nor HP wanted free advertising, we'll use Supermicro hardware in this example.

Example - The Head Node

Each Head Node will consist of a 2029U-TR25M:

- 2 x Intel® Xeon® Gold 6126F Processor
- 12 x 16GB DDR4-2666 RDIMM
- 1 x Netlist EV3 16GB NVRAM (ZIL)
- 2 x LSI 9305-16e SAS HBA
- 2 x Qlogic QLE2692-SR-CK dual 16Gb FC HBA
- 1 x Intel XXV710-DA4 dual 10/25GbE adapter
- 2 x Intel S4600 240GB SSD (OS)
- 1 x Supermicro HFI Carrier Card Kit

Example - The Drive Tray

Each drive tray will consist of a SC216BE2C-R741JBOD :

- 24 x Seagate XS3840SE10013 3.84TB SAS SSD

Example - Cables

Several cables will be needed to connect this all together.

- 1 x 1M Omni-Path cable
- 16 x 1M SAS3 cable
- 8 x 1.5M SAS3 cable
- As required - Cat5E Ethernet cable
- As required - 25GbE cable, or fiber with optics
- As required - 16Gb Fiber Channel fiber and optics

Example - Connections

We combine the two systems to achieve a high-availability configuration. The NetList card is used for the ZIL. The Omni-Path connection, with DRBD, is used to mirror the ZIL to the remote systems, and to make the volumes available through both sets of FC interfaces.

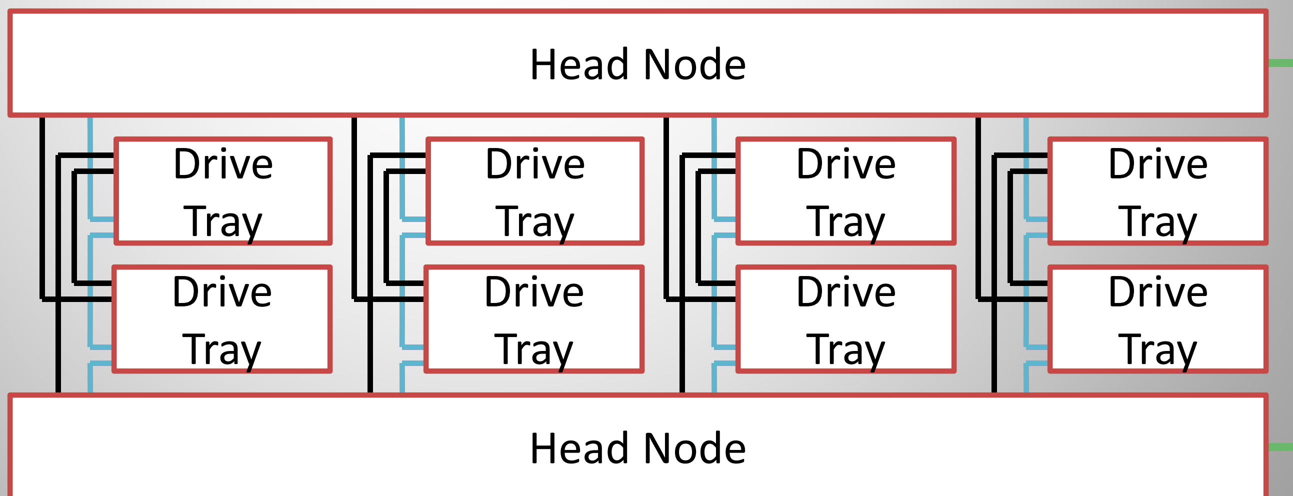

NetList NVRAM

In the example, the NetList NVRAM card is used for the ZIL. When using SSDs for storage, synchronous writes, including those needed for file-system integrity, are a performance bottleneck. To be effective the ZIL needs to be significantly faster than the storage. This can only be accomplished with NVMe or NVRAM drives. The NVRAM drives are large enough to function as a ZIL, and are at least 3x faster than NVMe.

For use with disks, an SSD or NVMe drive is adequate.

Omni-Path

Omni-Path Architecture is Intel's latest HPC mesh interconnect. It uses 16 PCI-E lanes to allow communication between computers and switches. It is also possible to use OPA to communicate directly between two computers without a switch.

This is a 100Gb connection with sub-microsecond latency. When used under DRBD (Protocol B), DRBD is almost as fast as the NetList card and faster than NVMe.

The Path to Six Nines Storage

When working with commodity hardware, four nines is easy and five nines takes some work. To get to six nines you usually need to **rethink the problem**.

Greater complexity often generates lower reliability. Don't depend on complexity; **depend on good design**.

The Path to Six Nines Storage

Play the Game

- Statistics - Get them to work for you
- Scheduled system downtime doesn’t count against you
- Perform recovery operations during scheduled system downtime

The Path to Six Nines Storage

- NFS over UDP - The DNS TTL for the file server will have to be set VERY low
- NFS over TCP - Use sub-interfaces to move the MAC from the Primary server, in addition to DNS above
- FC and iSCSI - The secondary server will need to accept requests and forward them over the OPA link to the Primary server
- Clustered server - Have at least three Storage Manager nodes with 3-way data replication (should deliver at least seven nines)

Feedback

My goal is to improve this presentation. To that end, I would appreciate feedback to:

- ashford@ACCS.com
- www.linkedin.com/in/peterashford