Best Practices and Guidelines for Integrating the Dell EqualLogic FS7600 and FS7610 into an Existing SAN

A Dell EqualLogic Best Practices Technical White Paper

Dell Storage Engineering
February 2013

Table of contents

Acknowledgements

This best practice white paper was produced by the following members of the Dell Storage team:

Engineering: Mike Kosacek

Technical Marketing: Raj Hosamani

Editing: Margaret Boeneke

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to SISfeedback@Dell.com.

SISfeedback@Dell.com

Executive summary

A Dell™ EqualLogic™ FS Series NAS Appliance can be installed into a new or existing EqualLogic iSCSI SAN to provide unified block and file I/O capability. Before an FS7600 or FS7610 is deployed into an existing SAN, it is important to ensure the existing SAN has available resources to support existing servers using block-based I/O as well as clients accessing the Dell Fluid File System (FluidFS). This paper helps the SAN administrator or deployment specialist by presenting best practices for deploying the FS Series NAS Appliance into an existing EqualLogic SAN environment.

1 Introduction

The Dell EqualLogic FS7600 and FS7610 FS Series NAS Appliance add scale-out unified NAS capabilities to EqualLogic iSCSI SANs and enable management of iSCSI, CIFS, and NFS access from a single management console. The FS Series NAS Appliance integrates into an EqualLogic iSCSI SAN and can be deployed along with additional new arrays or it can leverage an existing EqualLogic iSCSI SAN deployment. Unlike many unified storage solutions that only scale in capacity, the EqualLogic FS Series NAS Appliance can scale capacity and performance. When more storage capacity or I/O performance is required, adding more EqualLogic arrays or expanding the Dell Fluid File System (by adding an additional appliance) enables the system to scale out to meet those needs.

The EqualLogic FS Series NAS Appliance incorporates the Dell FluidFS, which is a high-performance scale-out file system capable of presenting a single file-system namespace through a virtual IP address, regardless of the number of NAS controllers in the cluster. While many NAS solutions have strict limits on file share size, the Dell Fluid File System and FS Series NAS Appliance allow a single file system to scale to the capacity of the EqualLogic storage deployment (up to 509 TB usable) and present it as a single name space.

1.1 Audience

This white paper is intended for storage administrators who are planning, implementing, or integrating a Dell EqualLogic FS7600 or FS7610 into an existing EqualLogic iSCSI SAN. Readers should be familiar with general concepts of Dell EqualLogic iSCSI storage as well as Ethernet Local Area Network (LAN) and basic Network Attached Storage (NAS) concepts.

1.2 Terminology

The following terms are used throughout this document.

Table 1 Definition of common terms

Term	Definition
Block Heavy Mix	For the purpose of this paper, a workload which consists of 67% block I/O and 33% file I/O.
Block I/O or Block storage	Describes how an application on a host server accesses data on a local or SAN connected storage system.
Common Internet File System (CIFS) also known as Server Message Block (SMB)	A network protocol typically used by Windows clients to access printers and file shares.
File Heavy Mix	For the purpose of this paper, a workload which consists of 67% file I/O and 33% block I/O.

Term	Definition
File I/O	Describes application or client access to a NFS mount point or CIFS share.
Group	A PS group consists of one or more PS Series arrays connected to an IP network. A group may contain up to 16 arrays.
Local Area Network (LAN)	Network usually used for client connectivity.
Network Attached Storage (NAS)	A network connected, file-level storage device that usually supports protocols such as CIFS and NFS.
Network File System (NFS)	A network protocol typically used by Unix/Linux systems to access NAS.
Pool	A container that each member (array) is assigned to after being added to a group. A pool can have up to eight members.
Storage Area Network (SAN)	A network used exclusively for storage connectivity and usually supporting iSCSI or Fibre Channel access to storage.
Structured data	Data that is managed in a system such as a database.
Unstructured data	Data that is largely unmanaged, such as user files on a share.

For further detail and definitions, see the Dell EqualLogic Configuration Guide.

2 Dell EqualLogic FS Series NAS Appliance

The FS Series NAS Appliance is an active/active appliance cluster that consists of two controllers housed in a single 2U rack space enclosure. Each appliance provides 48 GB of RAM (24 GB in each controller) that is used for the system as well as the mirrored, battery protected cache memory. Up to two FS7600 or FS7610 can work together to provide file storage from one EqualLogic group. Both NAS appliances actively participate in file I/O and automatically distribute new users across available controllers and network ports.

The EqualLogic FS Series NAS Appliance supports all new and existing EqualLogic arrays that are running firmware version 6.0 and later. A NAS cluster with one or two FS7600 or FS7610 appliances can access a single storage pool with up to eight EqualLogic PS Arrays. Each EqualLogic FS7600 controller has four 1 GbE ports for client-side NAS connectivity for NFS and CIFS clients. Four additional 1 GbE ports are used to connect to the EqualLogic iSCSI SAN. Each EqualLogic FS7610 has two 10 GbE ports per controller for client-side NAS connectivity, and two more 10 GbE ports per controller for connectivity to an EqualLogic iSCSI SAN. Controllers in each appliance communicate internally through an internal PCIe backplane for tasks such as write-cache mirroring.

The FS Series NAS appliance also has built-in snapshot capabilities, using a redirect-on-write method. When a snapshot is enabled, any changes to existing data in that snapshot are preserved, and subsequent read/write access is redirected to the new data. This method of snapshots is implemented in the file system (FluidFS) and does not require additional physical disk space for holding snapshots. Instead it utilizes reserved, unused space from the existing NAS file system. Additional features of the FS Series NAS Appliance include replication, Active Directory/LDAP integration, NDMP–based backup support, Network Information Services (NIS), CIFS antivirus integration, and user quotas.

More information on these and other features can be found at www.dell.com/fs7600.

3 Benefits of NAS

The Dell EqualLogic PS Series storage is a leader in the iSCSI SAN market. With the addition of the FS Series NAS Appliance, EqualLogic SANs can also support NFS and CIFS based file access. Management of file access has been integrated into the EqualLogic Group Admin user interface transforming the EqualLogic platform into a robust unified storage platform.

The EqualLogic FS Series NAS Appliance is built on the foundation of the Dell FluidFS. FluidFS is a distributed, scale out file system that provides a single name space across multiple NAS controllers.

One general advantage of a NAS file system is the ability to enable multiprotocol access, the simultaneous access by both Windows (CIFS/SMB) and Unix/Linux (NFS) file clients.

Because of the unique architecture of the Dell FS7600 and FS7610 NAS Appliance, additional benefits are achieved with FluidFS:

- Expansion of NAS volumes without requiring downtime
- Automated load balancing across NAS controllers
- Unified block (iSCSI) and file (CIFS/NFS) access
- Global name space that can scale up to .5 PB today
- Robust, built-in data protection features such as snapshots, replication, and NDMP backup
- Integrated user interface for easily managing both block and file access

4 Lab test objectives

There are a few considerations that need to be made to ensure that the FS7600 and FS7610 can be properly integrated into an existing SAN. While the physical connections are very similar to any other host attached to the EqualLogic SAN, the SAN administrator also needs to consider the requirements of the FS Series appliance such as capacity and pool configuration as well as performance before and after integrating into an existing SAN.

This study investigated the following points:

- The performance impact of adding file I/O to an existing block-based SAN with additional hosts attached
- The performance impact of mixing block and file I/O in the same pool versus separate pools
- The effect of the I/O request size and the MTU size used on the client-side network.

The results presented in this paper can be used as a guideline to help you determine the right answers for your own environment.

5 Test topology and architecture

To study the performance impact of integrating an FS7600 or FS7610 into an existing environment, Vdbench was used to generate a workload of both block and file I/O. Vdbench is an open-source I/O workload generation tool that is available for download on SourceForge.net. Vdbench can be controlled by a script, provides data verification, and allows flexible workload definitions for both block and file I/O. Vdbench requires a JAVA runtime environment, so the Linux OpenJDK Runtime Environment was used in these tests.

Vdbench is available at http://sourceforge.net/projects/vdbench/.

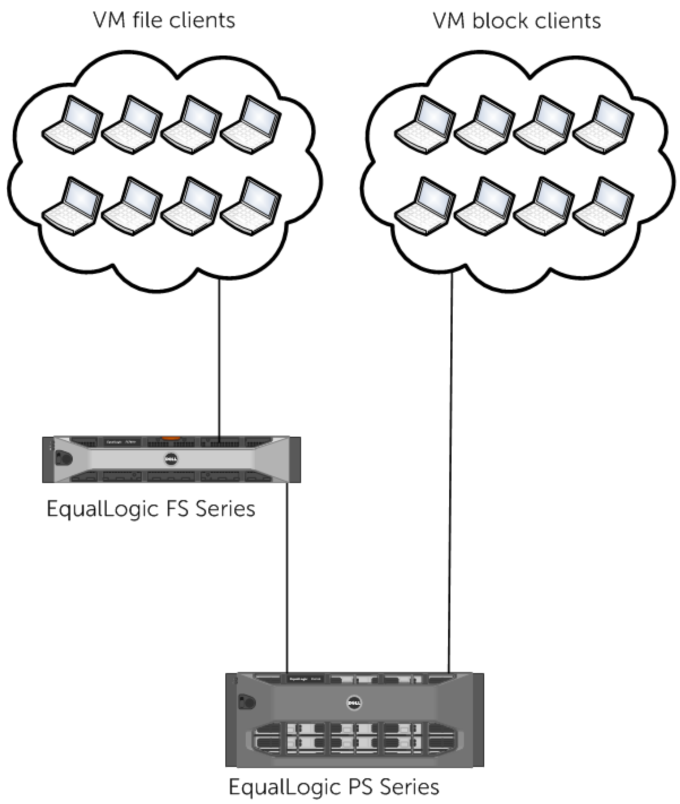

Figure 1 Logical topology

For the workload generation clients, multiple VMware ESXi 5.01 servers were configured with Virtual Machines (VMs) running Linux. The ESXi servers consisted of three Dell PowerEdge servers for the block I/O clients and three for the file I/O clients. Three of the ESXi servers connected to the back-end SAN switches (and to SAN storage) to support VMs that simulated block-attached servers, while the others connected to the front-end client LAN switches to support VMs that simulated NAS-attached clients or servers. For both the 1 Gb and 10 Gb test scenarios, there were 16 clients for block testing and 16 clients for file tests. See Appendix B for details of each configuration.

On the SAN network, jumbo frames (MTU 9000) were configured. This is established best practice for EqualLogic iSCSI SANs and is also required for the FS Series controllers. The standard MTU size of 1500 was used for the client network. Some initial tests were run with both standard and jumbo frames but for the test cases, no significant difference was found so the standard MTU was used for all subsequent test scenarios. The standard MTU size of 1500 is also the most commonly used for typical client networks.

The EqualLogic iSCSI SAN consisted of eight PS Series arrays configured for RAID 10. Initially a single pool with all eight array members was created. Both block hosts and the FS Series controllers accessed volumes in this pool. Later two separate pools were created, each with four array members. The FS Series controllers then accessed volumes in the first pool, and the block hosts accessed volumes in the second pool.

For SAN (block I/O) testing, 16 100GB volumes were created on the SAN and one volume was presented to each of the 16 Linux VMs as iSCSI volumes. Each Linux VM had two virtual NICs for iSCSI SAN access, had the Linux Host Integration Toolkit installed to enable multi-path I/O (MPIO). For file I/O (CIFS and NFS), two file system mount points were created (one each for CIFS and NFS access) on the FS7600 or FS7610 which was then mounted from each of the VM clients.

A Vdbench script was created that generated I/O at the desired block size (or I/O request size) and was run for a specific amount of time. The results were then collected and analyzed.

6 Test methodology

For the block I/O testing, Vdbench was configured to read and write from RAW disks (disks without a file system). A working file size of 99 GB was used for the block clients to use close to the full capacity of the attached volume and also to minimize the effect of client-side caching in the VM.

For the file I/O tests, initially tests were run against both CIFS and NFS mount points. The results were similar, so for the combined block and file testing, only NFS mount points were used for the file clients. Vdbench was configured to create a directory structure that was three directories wide and two directories deep. Each bottom-level directory had sixteen 50 MB files which resulted in a total working set of about 5 GB for each client. Before running the workload test scripts, a separate formatting script was run to create the file and directory structure on the file system.

Each test client used a workload of random I/O consisting of 70% read and 30% write operations. Each test was run with an I/O request size of 4K, 8K, 16K, 32K, 64K, 128K, and 256K. Eight Vdbench worker threads were run for each client VM. This resulted in a client that generated more I/O than a typical user might, however it minimized the number of overall client systems required to generate a significant workload on the SAN. Each workload was run for 30 minutes, and then the results were collected from each Vdbench client and from a separate management system that was running SAN HQ.

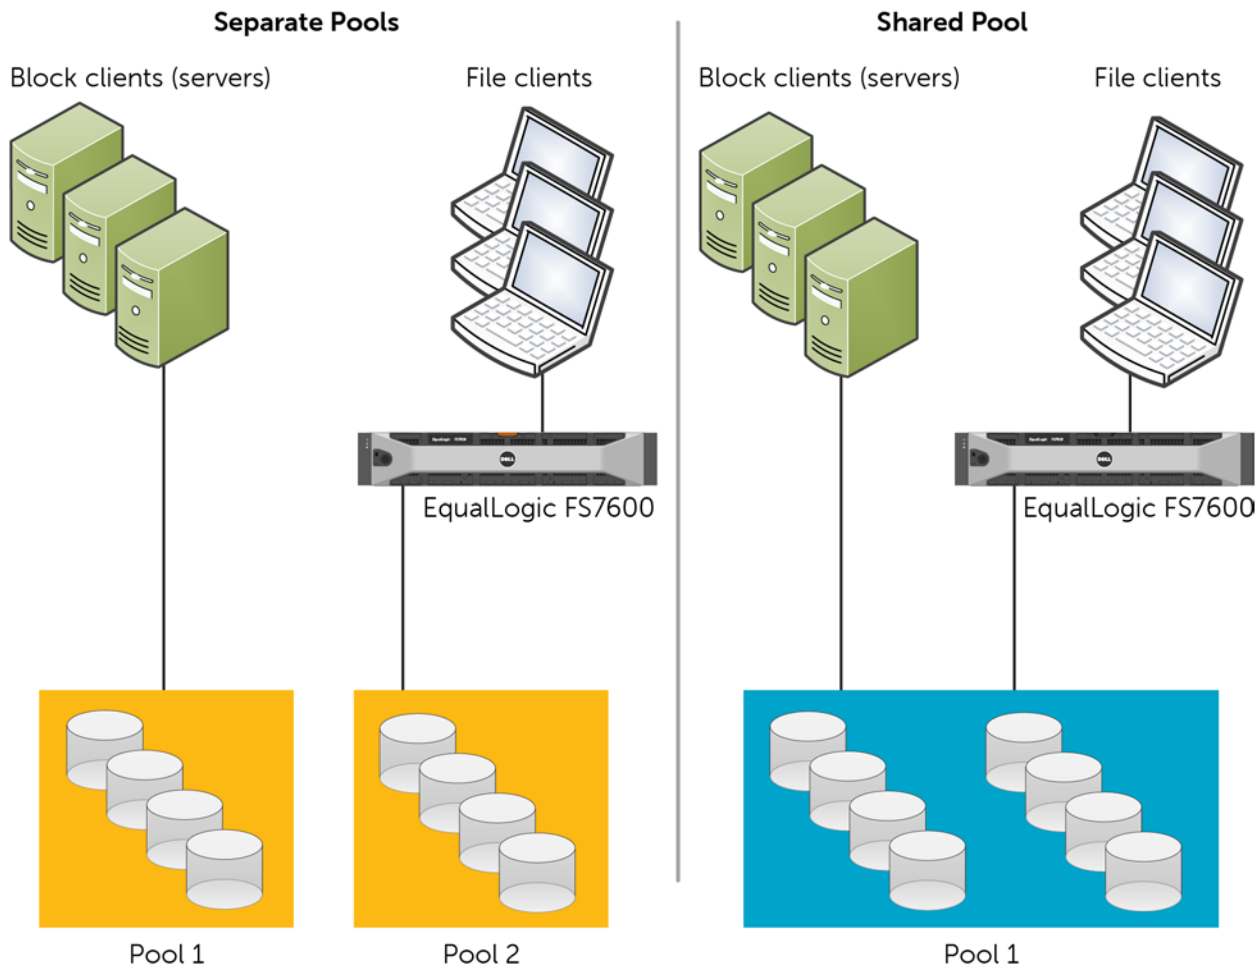

Figure 2 Logical diagram for separate or shared pool

The tests started with only block clients and scaled up from 2 to 16 clients. Then the tests were repeated with CIFS and NFS clients to verify that each set of clients would scale as expected. Next, the file clients were configured for NFS access and tests were run with both block and file I/O at the same time. Then block heavy scenarios were run that consisted of 67% block hosts and 33% NFS file clients. Then the mix was reversed to a file heavy scenario where 67% of the hosts were doing NFS file I/O and 33% of the clients were doing block I/O. These tests were first run with a single, shared storage pool (with eight arrays) then they were repeated with two pools, each containing four arrays (see Figure 2).

7 Test results and analysis

Note: The results provided in this paper are intended for the purpose of comparing the specific configurations used in our lab environment. The results do not portray the maximum capabilities of any system, software, or storage.

7.1 Small block results with shared versus separate pool design

For the combined block and file tests, the tests were run at various block sizes from 4K to 256K. Similar scaling results were achieved over each block size although only the results from particular tests may be shown in this paper. One example of the results is shown in Figure 3, which shows the IOPS performance achieved with the configured workload using an I/O request size of 4K while the FS7600 shared the storage pool with the block clients. As expected, because the number of block and file clients were varied in each test, the block heavy mix shows more total block IOPS (with 16 block clients) versus the file heavy mix which only used eight block clients. Conversely, the file heavy mix shows a higher percentage of IOPS from the file clients.

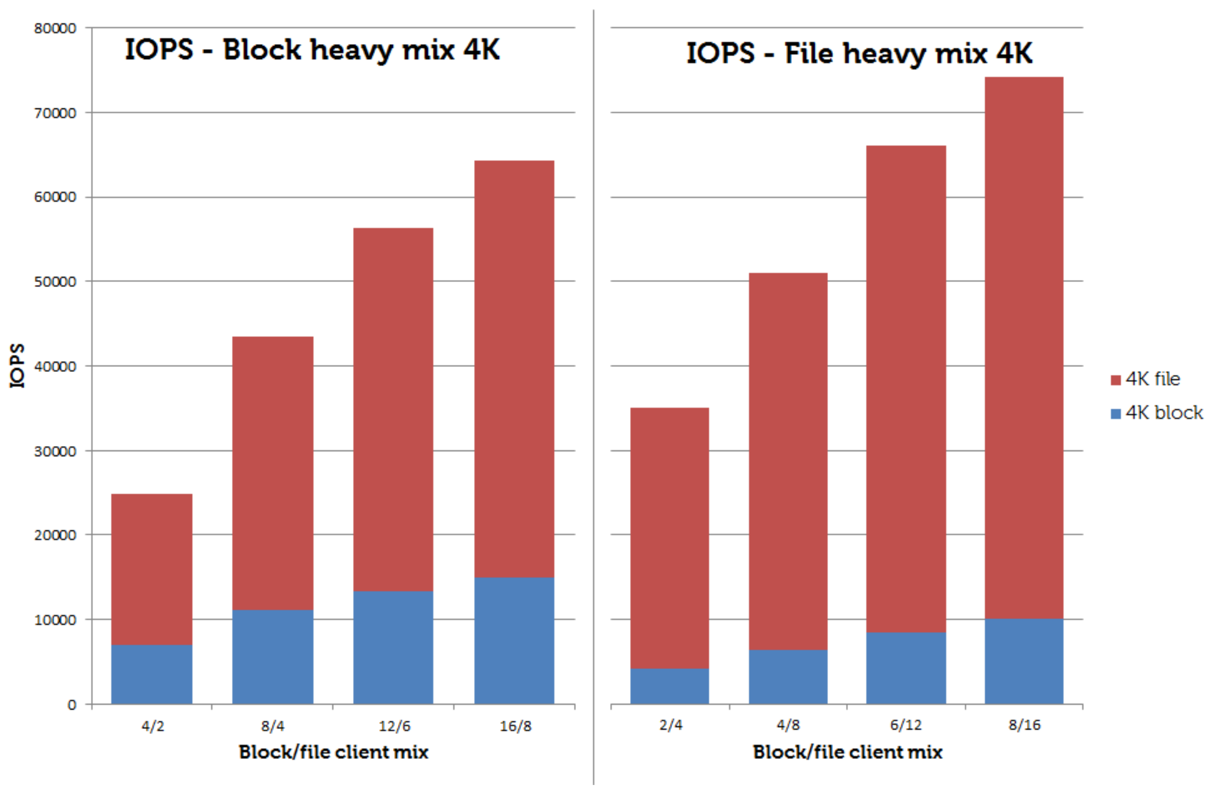

Figure 3 4K IOPS performance with FS7600 in shared pool

However, with the file heavy mix, the overall total number of IOPS (with both file and block combined) is greater than with the block heavy workload mix because of the ability to cache more of the active data on the FS7600.

When the same tests were run with the FS7600 configured in a separate pool (as shown in Figure 4), the block clients were able to scale slightly higher than with the shared pool, even though the separate pool only has half the number of arrays than were available with the shared pool design. This indicates that sharing the pool with the FS Series NAS Appliance can affect the performance of existing block hosts. This occurs because in a shared pool design, volumes from both block and file may reside on the same set of physical disks which can create contention and increase disk latency (or response time) under heavier workloads.

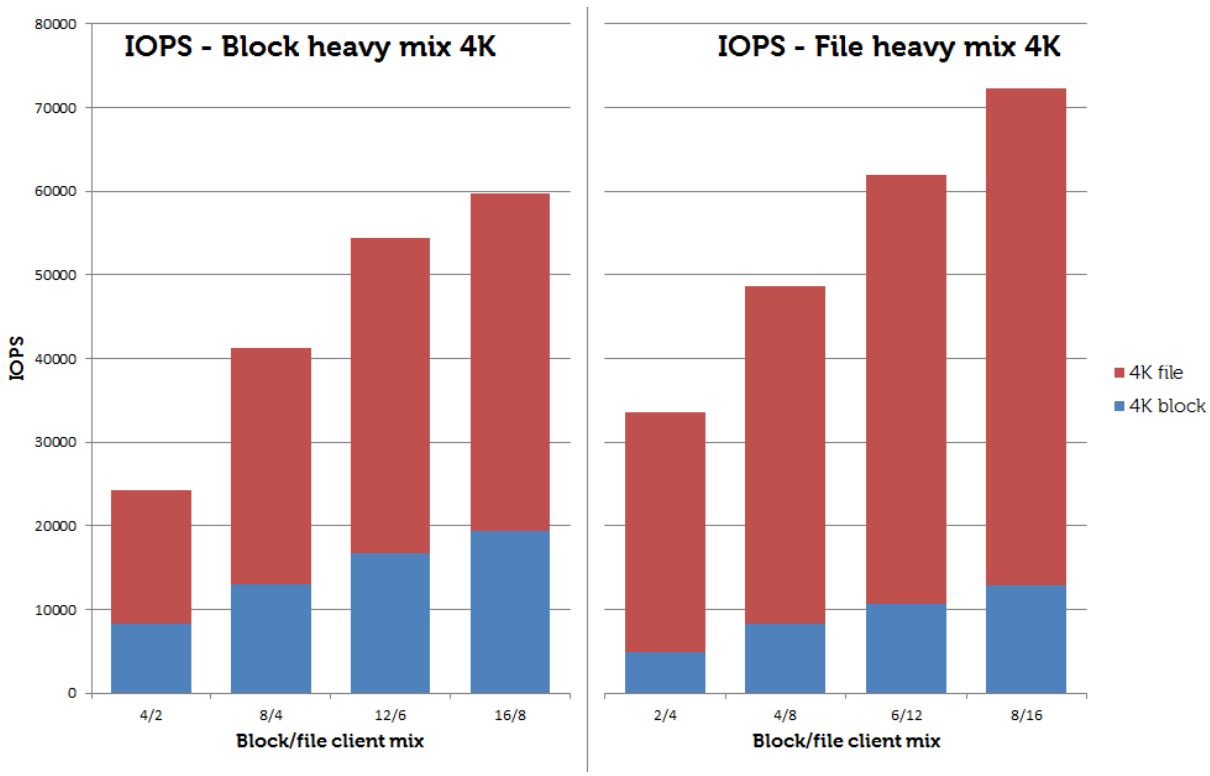

Figure 4 4K IOPS performance with FS7600 in separate pools

The separate pool test results also show that the FS7600 was able to scale similarly to the results seen with a shared pool when the file heavy workload mix was run. In the case of the block heavy mix, the file clients performed better in the shared pool design where more arrays, and therefore more disk spindles, were available.

7.2 Large block results with shared versus separate pool design

For workloads that consist of larger I/O sizes (such as backup and media streaming workloads), the network throughput delivered is a better performance indicator. In Figure 5 the results from the block heavy and file heavy tests are shown when the I/O request size from the Vdbench block and file clients was 128K. For these results, the throughput is shown in MB/sec.

With the block heavy mix, the total throughput numbers exceeded that of the file heavy mix. In both the block heavy and file heavy mix, the portion of throughput resulting from the file clients is similar though, which means the difference was mostly due to the increased block performance from additional block clients (in the block heavy workload mix). Even when the number of file clients doubled in the file heavy mix, the portion of file I/O did not increase, indicating that the file VM clients were limited and unable to provide additional throughput with this workload.

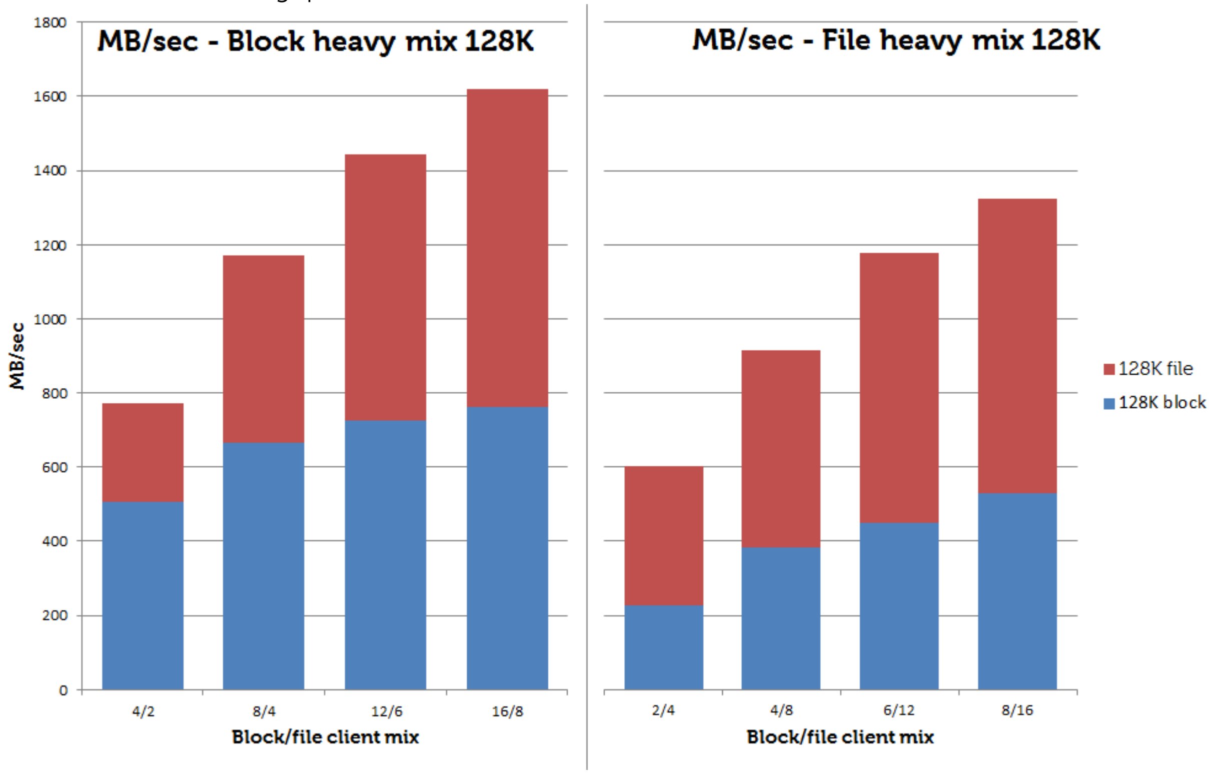

Figure 5 Throughput (MB/sec) with 128K I/O request in shared pool

Looking at the results in Figure 6, the maximum performance of the block I/O in a separate pool exceeded the results achieved when the block volumes shared the pool with the NAS reserve. In some cases increasing the number of disk spindles results in an increase in throughput too, however in this case mixing block and file I/O types affected the performance of the block clients. With larger I/O sizes, contention can occur at the physical disk or in the cache memory on the array members as well.

When comparing the file heavy mix in Figure 5 and Figure 6, the portion of file I/O remains fairly consistent whether using a shared or separate pool design. This is largely due to the presence of additional cache memory on the FS Series NAS Appliances. However, when comparing the block heavy mixes, the file I/O performance is slightly better in the shared pool design, similar to the results demonstrated by the small block workload. In general, when block and NAS volumes are placed into separate pools, the performance is more predictable and scales more evenly.

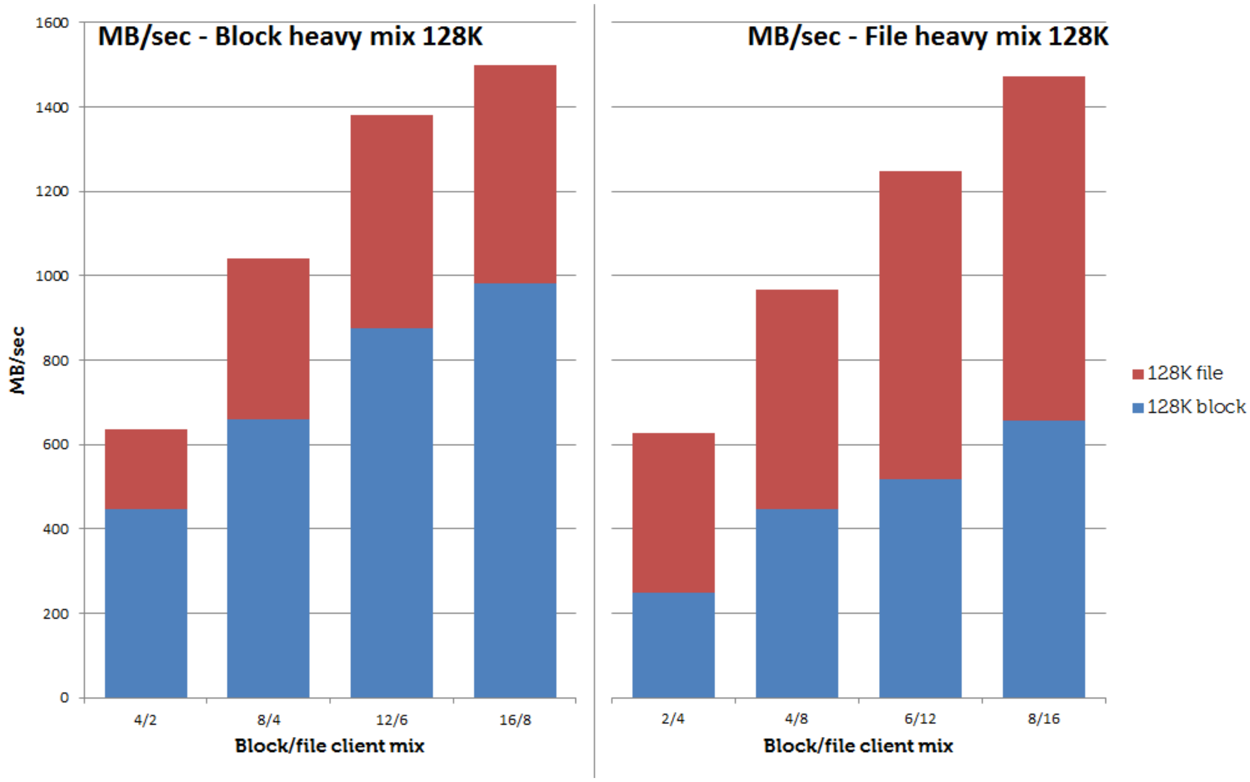

Figure 6 Throughput (MB/sec) with 128K I/O request in separate pool

For workloads consisting of mostly large files (128KB) and with random access patterns, configuring separate pools for block and file workloads delivers better aggregate performance and scalability. If more of the workload results from file clients, or if file clients are used exclusively, then a single shared pool may provide better performance.

8 Planning and design considerations

8.1 Network ports

Each FS7600 controller contains a total of eight 1GbE ports that must be connected for optimal performance. A total of 16 network ports are required for each controller pair. Each FS7610 controller contains a total of four 10GbE ports or a total of eight ports for each controller pair. As a best practice for EqualLogic SANs, Dell recommends that the LAN and SAN should be physically separated.

For maximum availability, it is recommended that both the LAN and SAN have redundancy in the switches. Connections from the FS7600 or FS7610 can be staggered across redundant switches in a stack or LAG or staggered across modules in a chassis so that a single switch (or module) failure does not prevent connectivity from the client to the NAS controller or from the controller to the SAN storage.

Figure 7 Rear view of FS7600 appliance showing network ports

On each FS7600 controller there are four ports for client connectivity and on each FS7610 there are two ports. The LAN ports must be connected to switches to allow CIFS and NFS client connectivity. The LAN ports can be configured for one of two modes, the default being Adaptive Load Balancing (ALB). LACP (IEEE 802.3ad) mode can also be used, but it requires further configuration of the network switches (on the LAN or client side) to support it. A total of eight 1Gb RJ-45 LAN switch ports are required for each FS7600 NAS Appliance (controller pair) or four 10Gb SFP+ ports for each FS7610 NAS Appliance.

On the SAN, there are four 1 Gb ports per FS7600 controller for iSCSI connectivity. The FS7610 requires two SAN connected 10Gb ports for each controller or a total of four for each controller pair. Some internal communication such as heartbeat between controllers and file system replication between FS Series appliances uses the SAN. The SAN switches should not be configured to block IPv6 or IPv6 multicasting as this is needed for some internal operations. Cache-mirroring between NAS controllers in a single chassis occurs over a PCI-express (PCIe) connection that is internal to the FS Series appliance chassis.

Table 2 Network port count requirement for FS Series appliances

FS model	# LAN ports per controller		# SAN ports per controller		Total ports used
	Per controller	Per appliance	Per controller	Per appliance	
FS7600	4 x 1Gb	8 x 1Gb	4 x 1Gb	8 x 1Gb	16 x 1Gb RJ-45
FS7610	2 x 10Gb	4 x 10Gb	2 x 10Gb	4 x 10Gb	8 x 10Gb SFP+

The FS Series controllers do not support VLAN tagging and VLANs are not required on either the LAN or SAN switch. However, if VLANs are used, then switch ports should be configured as untagged so that any tags are removed on port egress (when the packet leaves the switch).

Refer to the EqualLogic FS7600 Installation and Setup Guide or the EqualLogic FS7610 Installation and Setup Guide on www.equallogic.com for specific instructions on cabling the FS Series NAS appliance.

8.2 IP addresses

Each NAS controller requires a unique static IP address on the LAN. This allows Group Manager to manage the individual NAS controllers.

A NAS service IP address is used by clients to access any CIFS shares and NFS exports hosted by the NAS service. When a NAS service is configured, one NAS service IP address is required if clients are on the same subnet as the FS Series NAS and connections are automatically balanced across all available connections.

If the client network is routed, meaning clients are on a different subnet than the FS Series NAS, the client network configuration can be modified to add additional NAS service IP addresses to allow for proper load balancing across multiple client subnets. The total number of NAS service IP addresses required will depend on the bonding mode used for the NAS service:

- For ALB, four IP addresses are required for each NAS controller. For example, in a two-controller NAS appliance, eight NAS service IP addresses are required.
- For LACP (IEEE 802.3ad), one IP address is required for each controller. You need only one IP address for each controller because all the client network interfaces on a controller are bonded together. For example, in a single NAS appliance, only two NAS service IP addresses are needed.

On the SAN, each FS7600 controller requires four IP addresses for the iSCSI interfaces or for each FS7610 controller two IP addresses are needed. One additional IP address is required for management - allowing the NAS service to manage the PS Series storage.

The internal network is used for communication between NAS controllers. FS Series controllers within the same appliance will communicate over their internal PCIe connection, while appliances extend cluster communications to the SAN switches. Each FS Series cluster requires its own unique block of IP addresses for this purpose. When the internal network is configured, you can choose one of the following options:

- 256 IP addresses (/24) – specifies 256 for the IP address block size
- 128 IP addresses (/25) – specifies 128 for the IP address block size
- 64 IP address (/26) – specifies 64 for the IP address block size

In most cases, since the SAN should be comprised of physically isolated network switches, the internal network only needs to be unique from the IP addresses used for the storage interfaces on the FS Series NAS controllers and the storage arrays. The use of private (non-routable) IP address blocks such as 10.0.0.0, 192.168.x.x, and 172.16.0.0 are generally acceptable to use for this purpose and the network block (or subnet mask) chosen has no impact on system performance.

8.3 Pool configuration

Each EqualLogic group allows up to four separate pools. A pool logically separates storage arrays. For example, a volume in Pool A will not share arrays (and therefore physical disk spindles) with arrays in Pool B. Pools can be used to keep volumes on separate arrays or RAID levels, and it can also be used to separate application workloads. For example, if a new FS Series NAS Appliance will be added to an existing SAN, depending on the current workloads on the SAN, it may be desirable to configure the FS Series NAS Appliance to use a separate pool so that it does not impact performance in the existing pool.

Using a separate pool requires at least one new array, or enough free space that an existing array can be moved (evacuated) from the existing shared pool to create a new pool for the FS Series NAS Appliance. If an existing array is moved from another pool, also note that this could affect the performance in the existing pool since the number of spindles available in that pool would be reduced. Careful consideration of the current workload as well as required storage capacity should be taken before removing arrays from a pool that contains existing volumes.

In most cases, the choice will be between adding a new array (or arrays) or to allow the FS Series NAS Appliance to reside in the same pool as the other workloads. Using a separate pool provides for greater separation of the block and file I/O workloads, however it may also prevent some efficiency (in disk space utilization) that may be gained by sharing the same pool. Refer to the following table for general recommendations.

Table 3 Pool configuration best practice summary

Unified workload type		
Application Profile	**Block Heavy**	**File Heavy**
Small block I/O: OLTP database (SQL, Oracle), e-mail (Exchange)	Separate	Shared
Small file I/O: User home directories, general file share		
Large block I/O: Media streaming, Medical Imaging and records, CAD/CAM. Backup & restore	Separate	Shared
Large block I/O: Media streaming, Medical Imaging and records, CAD/CAM, Backup & restore		

Before installing and configuring the FS Series NAS Appliance, the administrator should ensure that there is adequate storage capacity (free space) in the pool to create NAS volumes as well as adequate I/O capacity. If there is not enough free space, then additional arrays must be added to the pool. If there is not adequate I/O capacity to support the new file workload, then the performance of existing block workloads may be decreased as the storage system is forced to share resources between them.

8.4 Volume configuration

During configuration of an FS7600 or FS7610, the administrator is allowed to choose the pool that the NAS Reserve (and corresponding volumes) is created in. This can be an existing pool or a new pool containing at least one array member. All FS Series controller pairs use the same (single) pool and if the pool is shared with existing block volumes, there must be adequate free space available to create the NAS Reserve.

A minimum of 250 GB of free disk space in the pool (for the NAS Reserve) is required for each controller pair and is reserved for the system, so there must also be additional free space (greater than 250GB) allocated to create a NAS file system. Best practices also dictate that a minimum of 5% of total pool space be kept free to allow normal volume load balancing and distribution to occur automatically. The amount of NAS Reserve can be increased at any time (provided there is free pool space) however, once created, it cannot be decreased in size.

If the minimum required free disk space is not available, then additional arrays must be added to the group prior to configuring the FS Series NAS appliance. If additional arrays are added to an existing pool which will be shared by block and file I/O volumes, then it is recommended that the group be allowed time to re-balance existing block volumes across the newly added array(s) before configuring the FS Series NAS. Depending on the number of existing arrays and volumes, this operation could take several hours (or longer if there is active I/O on the SAN). If identical array models (with the same sized disks) are used, the re-balancing will be complete when each member of the pool contains a similar amount of the existing volume data.

The size of the NAS Reserve space that is allocated determines the number of volumes that are created for each controller pair. A minimum of two volumes will be created for each NAS appliance. The maximum volume size supported by the EqualLogic arrays is 15 TB; therefore if the space allocated to the NAS Reserve is greater than 30 TB, an additional pair of volumes will be created. This process continues until the capacity requirement is satisfied.

Volumes that are allocated to the FS Series NAS Appliance are controlled by the appliance, so many operations that are available in Group Manager and CLI access for traditional block volumes are disabled for the volumes that makeup a NAS Reserve. This includes options such as Auto-replication, snapshots, and the ability to online/offline a volume. Instead, snapshot and replication functionality is available through NAS-specific commands in Group Manager or the CLI. This functionality enables system administrators to implement file system aware snapshots and replication.

Once the NAS cluster is configured, Group Manager can be used to monitor statistics such as free and used capacity of pools, NAS Reserve, NAS file systems (containers), Snapshot reserve, and used snapshot space.

8.5 Monitoring an existing SAN

SAN HQ is included with every EqualLogic storage system and is ideal for monitoring EqualLogic SANs. Administrators can collect data from the SAN and review it to ensure that there is sufficient storage and I/O capacity for current needs as well as future growth. Figure 7 shows an example of the SAN HQ user interface and some of the many statistics that it can provide.

Prior to adding an FS Series NAS Appliance to an existing SAN, Dell recommends using SAN HQ to collect statistics about the existing SAN. Administrators should carefully monitor historical disk capacity usage and trends to ensure that there will be sufficient free pool space to integrate the FS Series NAS Appliance. Administrators should examine performance statistics such as the average IOPs over longer periods of time as well as examining any short bursts of activity. Understanding which hosts each type of activity comes from helps the administrator determine optimal configuration for pools, number of arrays needed, and other settings.

Administrators should also look at statistics such as Queue Depth to help understand the average amount of outstanding I/O that is waiting for service. Higher numbers indicate higher workloads from attached hosts and in some cases may indicate the need for additional disk I/O capacity (more physical disk spindles) to satisfy performance needs. Latency is also an important factor that affects the response time of applications. Higher latencies may be an indicator of slower response times for the attached hosts and applications running on them. However it is also important to note that larger I/O request or block sizes will usually result in longer latencies than smaller block sizes. Activities such as backup and restore or replication, which usually involve very large block sizes, may show higher latencies but this typical and acceptable for most environments.

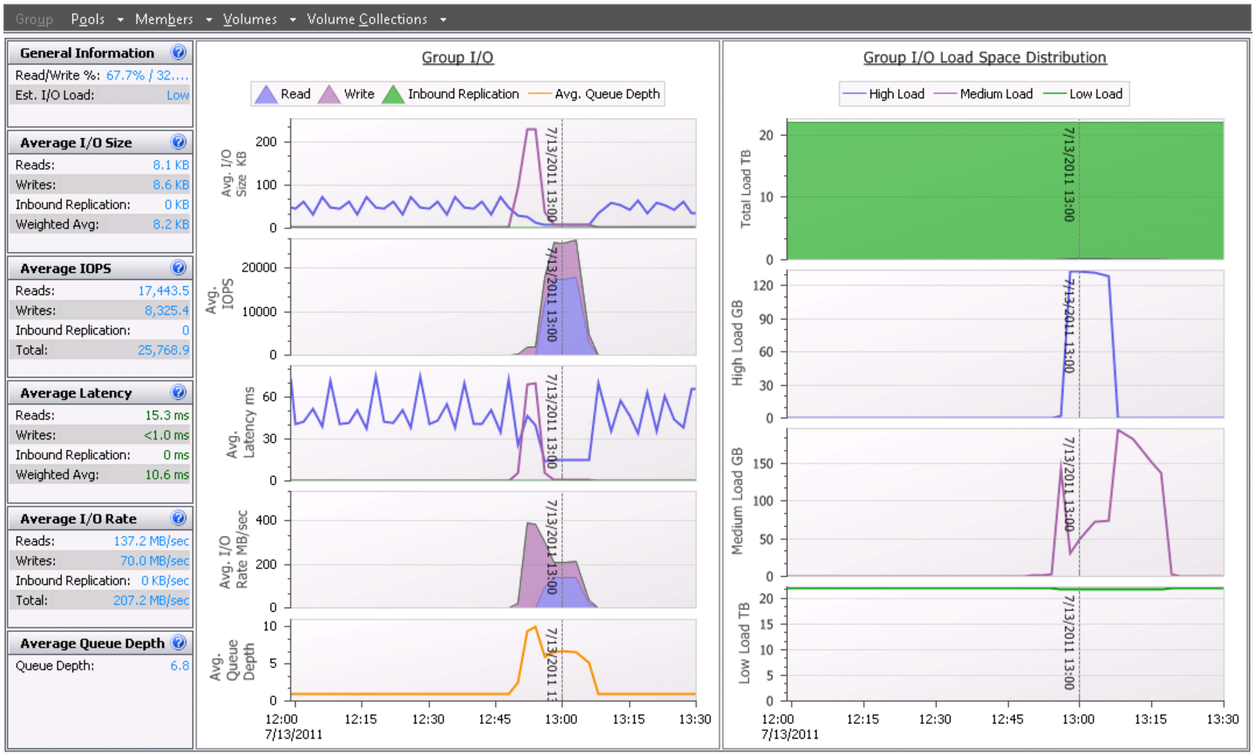

Figure 8 SAN HQ data

9 Summary

The Dell EqualLogic FS7600 and FS7610 can be deployed as part of a new installation or into an existing EqualLogic iSCSI SAN to provide a unified block and file storage solution. When deploying into an existing SAN, administrators should monitor the existing SAN to understand the availability of not only data storage capacity, but also of I/O capacity. Once the existing environment is well understood, the data presented in this paper can then be used to help determine the following:

- How many network ports are required?
- Is there is a need for additional storage resources before deploying an FS Series NAS?
- Will a shared storage pool or separate storage pools perform better?
- Will the addition of an FS Series NAS affect performance of existing systems?

These decisions will help guide the design and deployment of the FS Series NAS Appliance and accompanying EqualLogic iSCSI SAN to achieve optimal performance. Depending on the type of I/O from network clients, the use of the FS Series NAS Appliance may even enable an increase in overall performance capacity for the SAN through caching of network client requests, thus reducing the performance requirement from the storage arrays.

A Appendix – Vdbench script configuration

A.1 Vdbench script used for block client:

Note: The script was customized for each client to access a unique volume.

```
*******************************************************************
*
* 100% random, 70% read block I/O test.
*

hd=default,user=root,shell=vdbench
hd=b1
sd=default,size=99g,threads=8
sd=sd1,host=b1,lun=/dev/eql/b1,openflags=o_direct
wd=default
rd=rd1,sd=sd*,forseekpct=100,rdpct=70,forxfersize=(4k,8k,16k,32k,64k,128k,256k),
elapsed=600,interval=30,iorate=max,pause=60

**********************************************************************
```

A.2 Vdbench script used for NFS file client:

Note: The “anchor” directory must be unique for each client.

```
*******************************************************************
* 100% random, 70% read file I/O test.
*

hd=default
fsd=default,files=16,depth=2,width=3,size=50m
fsd=fsd1,anchor=/nas/nfs/client1

fwd=fwd1,fsd=fsd*,fileio=random,xfersizes=4k,rdpct=70,threads=8
fwd=fwd2,fsd=fsd*,fileio=random,xfersizes=8k,rdpct=70,threads=8
fwd=fwd3,fsd=fsd*,fileio=random,xfersizes=16k,rdpct=70,threads=8
fwd=fwd4,fsd=fsd*,fileio=random,xfersizes=32k,rdpct=70,threads=8
fwd=fwd5,fsd=fsd*,fileio=random,xfersizes=64k,rdpct=70,threads=8
fwd=fwd6,fsd=fsd*,fileio=random,xfersizes=128k,rdpct=70,threads=8
fwd=fwd7,fsd=fsd*,fileio=random,xfersizes=256k,rdpct=70,threads=8

rd=rd1,fwd=fwd1,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync
rd=rd2,fwd=fwd2,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync
rd=rd3,fwd=fwd3,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync
```

```
rd=rd4,fwd=fwd4,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync
rd=rd5,fwd=fwd5,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync
rd=rd6,fwd=fwd6,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync
rd=rd7,fwd=fwd7,elapsed=900,interval=30,fwdrate=max,format=no,pause=30,openflags
=fsync

**********************************************************************
```

B Appendix – Topology and configuration detail

For clients, multiple VMware ESXi 5.01 servers were configured. Three of the servers were connected to the back-end SAN switches (and to SAN storage), while the others were connected to the front-end client LAN switches. On each ESXi server, VMs that ran Linux were created. For both the 1 Gb and 10 Gb test scenarios, there were 16 clients for block testing and 16 clients for file tests.

Table 4 VM configuration

VM configuration			
Purpose	**RAM**	**CPU**	**Operating System**
Block client	2GB	1	Red Hat Enterprise Linux (RHEL) 6.3 32-bit
File client	2GB	1	Red Hat Enterprise Linux (RHEL) 6.3 32-bit

The ESXi servers consisted of three Dell PowerEdge R610 servers for the block I/O clients and three Dell PowerEdge R710 servers for the file I/O clients. For the 1Gb (FS7600) tests, we used the onboard 1 GbE network ports for management connectivity, and then each server also had a pair of quad-port 1 GbE network cards. For the 10Gb (FS7610) tests, we used the onboard 1 GbE network ports for management connectivity, and then each server also had a pair of dual-port 10 GbE network cards.

Table 5 ESXi server configuration

ESXi server configuration				
Type	**RAM**	**CPU**	**# cores per CPU**	**Operating System**
Dell PowerEdge R610	24GB	2	6	vSphere ESXi 5.01
Dell PowerEdge R710	48GB	2	6	vSphere ESXi 5.01

For the 1Gb network, our LAN and SAN switches each consisted of a pair of Dell Force10 S60 switches with stacking modules. For 10Gb, our LAN and SAN switches each consisted of a pair of Dell Force10 S4810 switches each with two 40 Gb QSFP+ ports configured as a Link Aggregation Group (LAG).

Table 6 Network switch configuration

Network switch configuration			
Network Speed	**Switch Model**	**Interconnect type**	**Interconnect**
1Gb	Force 10 S60	Stack	24Gb Stack module
10Gb	Force 10 S4810	LAG	40Gb module

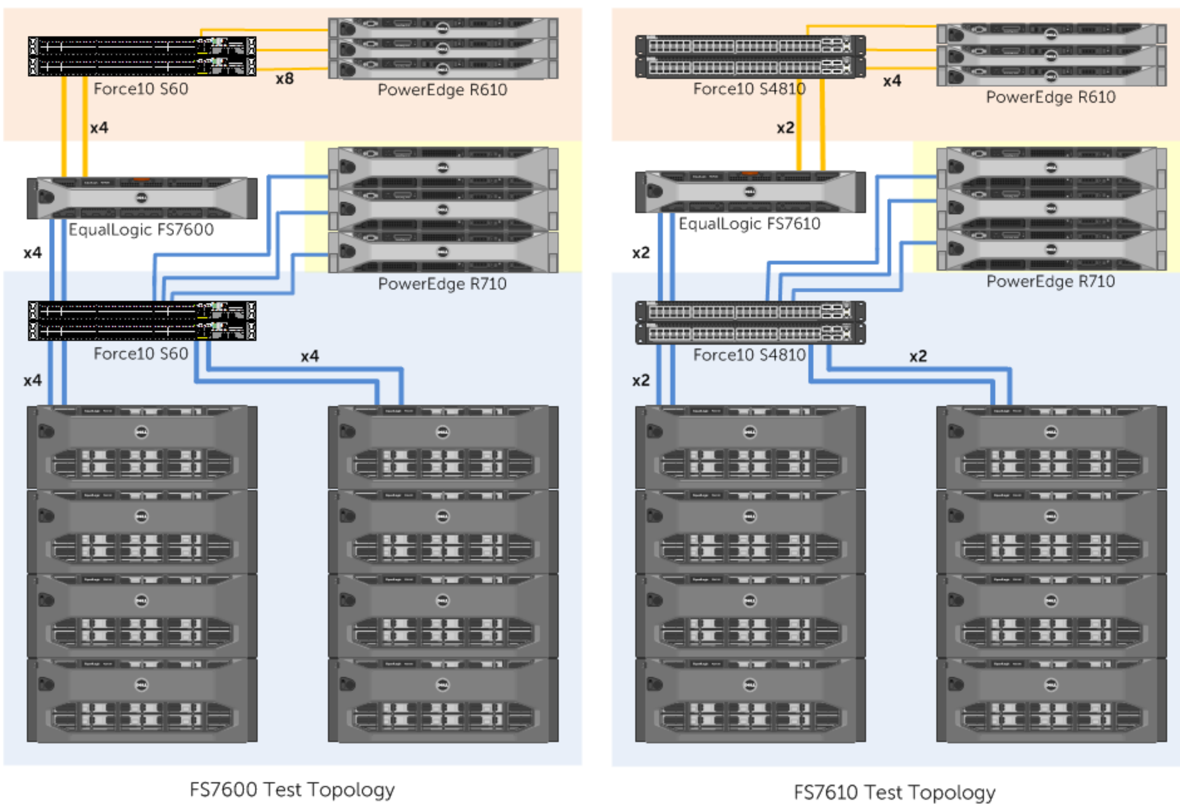

Figure 9 Test topology diagrams

The EqualLogic iSCSI SAN consisted of eight PS6100XV or PS6110XV arrays. Each array member was configured for RAID 10.

Table 7 iSCSI storage configuration

iSCSI Storage Configuration					
SAN Speed	**Array model**	**# of Arrays**	**# Drives per array**	**Drive Type**	**Array Firmware**
1Gb	PS6100XV	8	24	600GB 15K RPM	6.01
10Gb	PS6110XV	8	24	600GB 15K RPM	6.01

For file I/O (CIFS and NFS), we created two file system mount points (one each for CIFS and NFS access) on the FS7600/FS7610 which was then mounted from each of the VM clients.

Table 8 FS Series NAS storage configuration

FS Series NAS Storage Configuration		
NAS Reserve	**CIFS file system**	**NFS file system**
22.2 TB	2TB	2TB

Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and your installations.

Referenced or recommended Dell publications:

- Dell EqualLogic PS Series Network Performance Guidelines: http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229
- Dell EqualLogic website: www.dell.com/equallogic
- Dell FS Series NAS website: www.dell.com/fs7600
- Dell PS Series online help documentation: http://psonlinehelp.equallogic.com/V6.0/groupmanager.htm

For EqualLogic best practices white papers, reference architectures, and sizing guidelines for enterprise applications and SANs, refer to Storage Infrastructure and Solutions Team Publications at:

- http://dell.to/sM4hJT