

HP StoreVirtual 4335 Hybrid Storage and Microsoft SQL 2014 Reference Architecture



A high performance solution for your mission and business critical databases

Table of contents

Executive summary	2
Overview.....	2
Solution components	3
HP ProLiant server	3
HP StoreVirtual Storage	4
HP Networking.....	7
Microsoft Windows Server 2012 R2	8
Capacity and sizing.....	9
OLTP type workload description.....	12
OLTP workload data/results.....	13
Data-mining like workload description	15
Data mining workload data/results	15
Bill of materials	16
Summary.....	17
Appendix.....	17
SQL Server	17
HP Networking.....	18
Database workload tools	18
Glossary	18
For more information	19

Executive summary

Organizations of all sizes are dealing with a growing number of business applications. Small- to medium-sized businesses (SMBs) in particular are looking at highly available, high performing, scalable database solutions to store the mission critical and business critical data created by these applications. Many of today's core business applications deploy on a Microsoft® SQL Server database. Planning, implementing, and maintaining the storage infrastructure is critical to a successful SQL Server database deployment. Key challenges for the infrastructure include the requirement for data availability, efficient capacity utilization, and effective IO performance driven by the increasing dependency on solid database solutions. Database solutions for mission and business critical applications depend on the storage infrastructure to provide highly available data, to accelerate OLTP transaction response, and to shorten data mining query response times.

SMBs are moving increasingly to virtualize SQL Server environments to consolidate multiple database instances, increasing utilization of compute resources. While mission critical data typically require dedicated, redundant server resources and shared storage, HP StoreVirtual 4335 Hybrid Storage, with an innovative AO design, is an outstanding storage platform for database consolidation and shared storage. The HP StoreVirtual innovative AO design provides solid state drive (SSD) performance to SQL databases (DBs) without requiring the database to reside completely on expensive SSDs. HP StoreVirtual 4335 Hybrid Storage provides the scaling capacity and high data availability that Microsoft SQL Server needs to meet demands on the fly, delivering the full-featured shared storage required by mission and business critical environments.

Working in concert, HP StoreVirtual 4335 Storage, Windows® SQL Server, and HP Networks provide the shared high performance storage solution required by a SMB database solution. This reference architecture (RA) demonstrates how HP StoreVirtual 4335 Hybrid Storage provides a highly available, capacity efficient, high performance platform for the variety of SQL DB needs in today's SMB environments.

Target audience


The RA benefits customers, sales, presales, and solution architects, showing that all components, hardware, and software work together.

Overview

SMB customers need high availability (HA) and high performance for Microsoft SQL Server deployments because data access is as important to them as it is to enterprise customers. The objective of this RA is to highlight the advantages of StoreVirtual 4335 storage when used as the platform for multiple SQL database deployments. The RA will show how StoreVirtual 4335 can provide a highly available, capacity efficient and high performance platform for the variety of SQL DB needs in today's SMB environments.

HP StoreVirtual 4335 Hybrid storage delivers solid-state performance at hard disk drive (HDD) costs using both SSD and HDD storage tiers. The SSD tier delivers fast, as low as sub-millisecond, latencies while the HDD tier provides lower cost, SAS disk storage AO functionality in the HP LeftHand OS dynamically manages the distribution of a volumes' capacity between the storage tiers. The RA uses StoreVirtual AO technology to support the high performance requirements of mission-critical SQL DBs. Figure 1 represents the solution components in an efficient 15U rack space.

Figure 1. Reference architecture physical configuration



Solution components

HP and Microsoft provide these major solution components:

- HP ProLiant server series
- HP StoreVirtual 4335 Hybrid Storage
- HP Networking
- Microsoft Windows Server® 2012 R2
- Microsoft SQL Server 2014

HP ProLiant server

The HP ProLiant DL380p Gen8 Server series used for this solution sets the data center standard for 2U 2-socket rack servers with the latest in serviceability, enhanced configuration flexibility, customer-inspired design, and unmatched performance. Benefits include:

- Higher performance, power efficiency, and security via Intel® Xeon® E5-2600 and E5-2600 v2 processors.
- HP SmartMemory with greater memory bandwidth, DIMM count, and new memory health programs.
- Data protection with embedded HP Smart Array P420i RAID controller.
- Provides embedded provisioning tools, active health monitoring, and system maintenance capabilities built using the industry-leading HP Integrated Lights-Out (iLO) Management Engine.
- User-inspired design features—HP SmartDrive, Smart Socket guide, and flat cabling design.
- Easy tool-less access—anticipate your needs and help eliminate common issues that result in downtime of customer-serviceable components.

Note

The HP ProLiant server series also offers the HP StoreVirtual VSA. Gain fully featured software-defined storage right on your server: Visit www8.hp.com/us/en/products/data-storage/software-defined-storage.html.

HP StoreVirtual Storage

A key advantage of the StoreVirtual architecture is the extra availability gained when customer data is stored on more than one node, enabling data mobility across tiers, locations, and between physical and virtual storage. Network RAID of a volume protects multiple copies of data across a storage cluster, removing any single point of failure. Advantages include these features:

- HA and predictable linear performance increase with additional nodes
- No failover management required—all automatic
- Can sustain multiple concurrent failures and still keep data online and accessible to applications
- Space efficient block size of 256 KB
- Simple management interface that includes both configuration wizards and context help
- Five nines high availability
- Affordable multi-site disaster recovery
- Space efficient thin provisioning
- Shared storage
- Central management
- Simplified maintenance
- Non-disruptive scalability for performance and capacity
- Real-time monitoring
- Easy integration and operation

HP StoreVirtual 4335 Hybrid Storage¹

Built with proven ProLiant Gen8 technology in a 1U footprint, the StoreVirtual 4335 delivers higher performance, and lower latency, with less hardware reducing cost, power consumption, and rack space. The StoreVirtual 4335, (shown in figure 2), benefits and features include:

- Lower latency with SSD storage
- Automated sub-volume auto-tiering
- Reduced provisioning complexity

Specifications include:

- 1x Intel® E5-2640 2.5 GHz, six core, 12 threads
- 32 GB RAM—DDR3 1333 MHz
- Drives
 - 3x 400 GB enterprise mainstream SSD
 - 7x 900 GB 10K SFF SAS
- Hardware disk RAID 5 (default)
- HP SmartArray P420i with 2 GB Flashed Backed Write Cache
- 4-port 1 Gb Ethernet host interface
- 2-port 10 Gb Ethernet SFP host interface
- HP ProLiant Integrated Lights-Out 4 Remote Management
- Automated energy optimization capabilities

¹ For more information on the StoreVirtual 4335, visit hp.com/go/StoreVirtual4000.

Figure 2. Image of HP StoreVirtual 4335 Hybrid Storage




Tiered storage

Tiered storage is a grouping of like storage capacity internal to a StoreVirtual 4335 storage server that provides a hierarchical designation based on performance capability.

- Tier 0—Faster, smaller amount of SSD capacity.
- Tier 1—Slower, larger amount of SAS capacity.

Figure 3. LeftHand OS organizes the 4335 SSD capacity into Tier 0 storage and SAS capacity into Tier 1 storage



Adaptive optimization

Accelerates applications with patterns of unpredictable and changing workloads. OLTP like database workloads benefit from AO as pages storing actively hit database tables and indexes execute in SSD space. Virtualized environments are also excellent candidates for deployment on an HP StoreVirtual 4335 Hybrid storage solution. The data set associated with each of these applications frequently has a small percentage of data accessed more frequently. This “hot data” is ideally suited to take advantage of the AO functionality. AO dynamically balances I/O workloads between SSD and SAS drives, making the effective performance of the storage systems approach an all-SSD system in terms of both IOPS and latency. The system also benefits from the capacity provided by the SAS drives driving down the cost of the solution and increasing performance with AO.² Less active workloads characterized as file shares or archives will not significantly benefit from SSD performance. AO manages pages containing these less active workloads, based on their heat map locality, that are not performance sensitive in the HDD Tier 1 space.

StoreVirtual AO functionality overview:

- AO algorithm directs new writes to Tier 0.
- LeftHand OS reserves a small percentage of Tier 0 space for new data written to HP StoreVirtual storage reducing write latency.
- AO granularity is a single LeftHand OS page (256 KB).
- Each storage system maintains an access count for each page forming a “heat map.”
- Moves a page from Tier 1 into Tier 0 if the page is more frequently accessed than other pages in Tier 0.
- AO moves pages out of Tier 0 if the page is less frequently accessed as a page in Tier 1.
- Disabling volume AO will gradually move pages to Tier 1.
- Individual volumes are either “Adaptive Optimization permitted” or “Adaptive Optimization not permitted.”
- A threshold algorithm prevents page thrash between tiers.

² HP StoreVirtual Adaptive Optimization technical paper h20195.www2.hp.com/V2/GetPDF.aspx/4AA4-9000ENW.pdf

Best practice

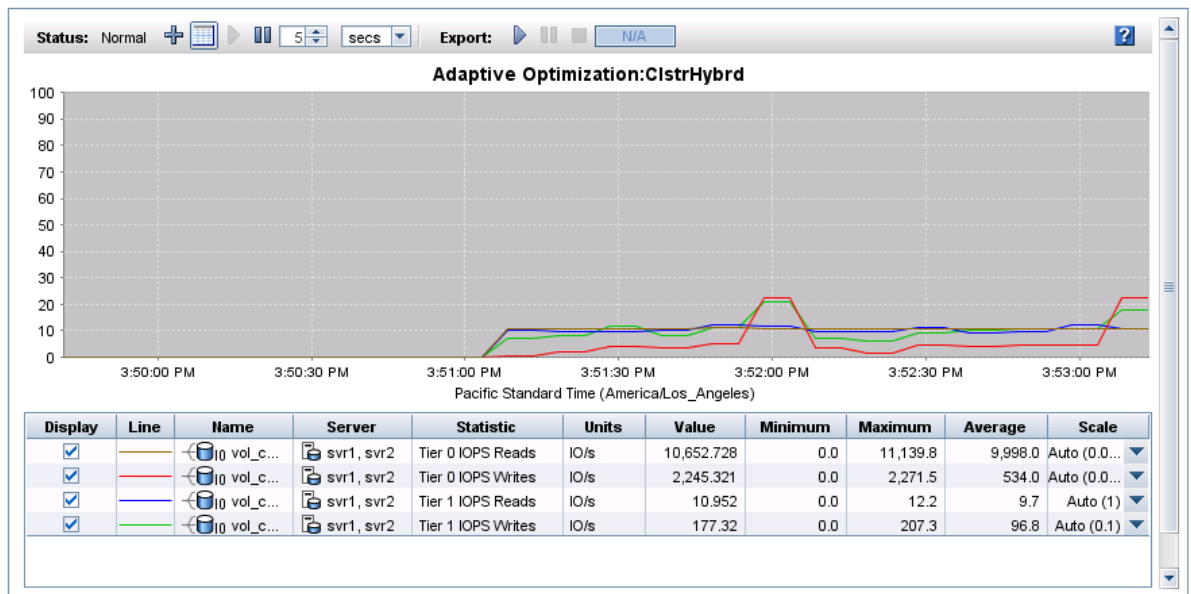
Separate *AO Capable* from *AO Not Capable* nodes into separate StoreVirtual storage clusters.

Note

Achieve the best value of AO with random, unpredictable I/O workloads such as database and virtualization workloads. Workloads without hot spots benefit to a lesser degree from the SSD tier and AO functionality.

The Central Management Console (CMC) contains a new performance monitor display for AO, shown in figure 4. The new panel presents tier and volume specific analysis by utilization, IOPS, and latency.

Figure 4. View of new AO Performance Monitor in CMC

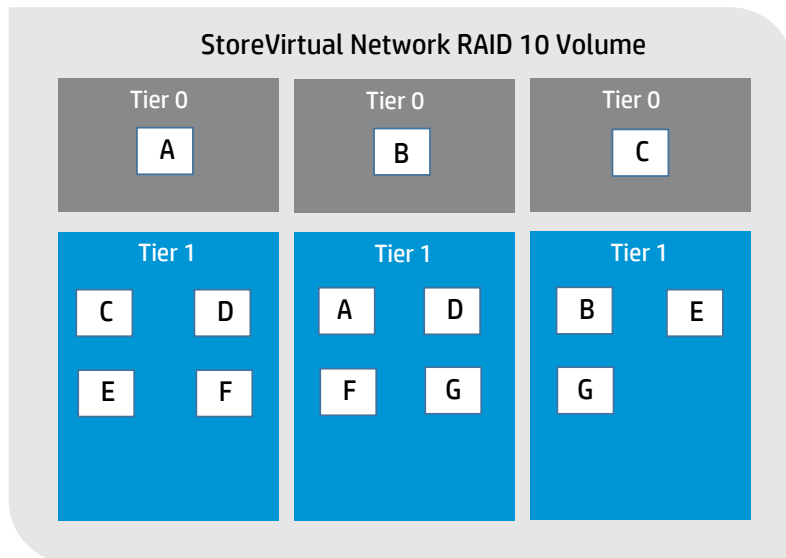


StoreVirtual Network RAID

Figure 5 represents three AO capable storage systems hosting an AO permitted, Network RAID 10 volume. Pages A, B, and C in the volume are “hot” pages and reside in Tier 0 space while their copies reside in Tier 1 space. Pages D through G and their copies, none of which are hot pages, reside in Tier 1 space. The customer benefit to Network RAID 10 is the highly available access to any given page from either of two separate storage systems. The customer benefit of AO is the increase in performance to accessing pages A through C currently in Tier 0 all done automatically and without user intervention.

- Network RAID configured at the volume level.
- Replication of pages is across nodes—not tiers.
- One copy of data might be in either SSD or HDD, the other in HDD.
- There is no need for both copies to be in SSD.

Figure 5. Example of Network RAID 10 and AO



NIC bonding

Network interface bonding provides high availability, fault tolerance, load balancing, and/or bandwidth aggregation for the network interface cards in the storage system. Joining physical NICs into a single virtual interface establishes the bond.

HP Networking

The escalating need to keep mission-critical data continuously available places new and increasing demands on a wide range of customer segments, including SMB environments. Enterprise environments that use HP storage solutions require continual operation with minimal or no downtime for key workloads that include file server, database, messaging, and other lines of business applications. This high level of availability is best achieved with multi-level redundancy built into storage, backup and recovery servers, server clusters, and component connections.

To obtain this reliable level of peak performance, the network for HP StoreVirtual must provide:

- High availability
- High bandwidth
- Low network latency

High availability

Redundant physical network paths are the basic underlying structure for building a highly available network of application servers and iSCSI based storage targets. ProLiant servers running Windows Server 2012 R2 contain a minimum of four network interface ports for SQL Server access and iSCSI traffic. NIC pairs connect to separate top of rack (ToR) switches. Windows Server 2012 R2 manages path redundancy through Windows NIC teaming and Multipath IO.

The redundant network links across devices should provide non-blocking paths. Do not use Spanning Tree Protocol because of the long convergence time involved during network failure events. If meshing or Intelligent Resilient Framework (IRF) technology is not available, then HP recommends Rapid Spanning Tree Protocol.

High bandwidth

Increasing maximum transmission unit (MTU) above the default value of 1500 bytes can increase throughput. The 4335 supports non-default values for MTU size and can benefit throughput. It is a best practice to tune MTU, end-to-end if it benefits throughput. The 4335 supports an MTU value up to 9000 bytes.

End-to-end network flow control manages the rate of data flow between network devices by allowing the receiving device to pause the transmission of data from the sender. Pausing data transmission reduces the risk of dropped network packets.

Low network latency


The network design must also consider network latency. Network latency is a major factor in application and I/O performance. Applications running in virtualized environments have network traffic patterns following an east west direction, compared to applications running on individual servers in traditional non-virtualized data centers. Design for a flat network with emphasis on an east-west network traffic pattern. Minimize the number of network hops and the distance between the application servers.

Intelligent Resilient Framework

HP Intelligent Resilient Framework³ (IRF) technology (figure 6 IRF topology) creates a network fabric with multiple HP 5900 Switch Series to provide data center-class availability and scalability. The HP 5900AF Switch Series with IRF technology virtualizes the physical switches into an IRF fabric. HP IRF provides these advantages.

- Consolidates management of many devices into a single, easy-to-manage, virtual switch.
- Delivers rapid network recovery (measured in milliseconds) from link or device failure to ensure applications and services are uninterrupted.
- Uses all links between switches and servers, allowing scalable “on-demand” performance and capacity to support critical business applications.
- Provides data center class availability and scalability.
- Supports 1: N redundancy; switches that form the IRF fabric appear as one unit in the network.
- Improves management efficiency and streamlines network topologies.
- Enhances disaster recovery by linking installations up to 70 kilometers apart, providing the same fast failover as if they were sitting side by side within the data center.

Figure 6. IRF topology



HP Networking 5900 Series (figure 7 5900AF switch) IRF deployment enables the required highly available network environment.

Figure 7. 5900AF switch



Microsoft Windows Server 2012 R2

The following components harden the database solution against unplanned events and increase the manageability of the environment.

Microsoft Hyper-V Server role

Business critical SQL Server workloads execute as virtual machines (VMs) also referred to as “child partitions”. The role enables virtualization of server hardware into parent and one or more child partitions.

³ Further details on HP Intelligent Resilient Framework at h10144.www1.hp.com/docs/irf/irf.pdf.

Failover clustering feature

Mission-critical database workloads employ the failover clustering feature. Enabled on multiple servers, the Windows Server Failover Clustering (WSFC) feature increases application and service availability. Enterprise applications such as SQL Server 2014 install as failover cluster instances.

NIC teaming

Windows NIC teaming provides bandwidth aggregation across member NIC ports and HA in the event of a NIC port or network interface failure. NIC teaming distributes management operating system and VM network traffic across member NIC ports.

MPIO

Multipath IO enhances availability between Windows servers and SAN using redundant data paths. In this RA, Windows MPIO employs server side iSCSI network management. HP DSM for Microsoft MPIO augments MPIO functionality with iSCSI session management and load balance policies.

Cluster shared volumes (CSV)

Cluster shared volumes enable cluster nodes access to a common shared disk. CSV increases cluster resiliency by having I/O fault detection and recovery over alternate communication paths between the nodes in the cluster.

Microsoft SQL Server 2014 builds on the mission-critical capabilities delivered in the prior release by providing breakthrough performance, availability, and manageability for your mission-critical applications. SQL Server 2014 delivers new in-memory capabilities built into the core database for online transaction processing (OLTP) and data warehousing, which complement existing in-memory data warehousing and Business Intelligence capabilities for the most comprehensive in-memory database solution on the market.

Capacity and sizing

The HP ProLiant DL380p Gen8 Server selected for the RA provide sound data center features and functionality for a 2U 2-socket rack server. With enhanced configuration flexibility, unmatched performance, and leading energy efficient design, the DL380p Gen8 offers the perfect solution for the dynamic compute requirements for virtualized SQL Server workloads.

The ProLiant servers in the reference architecture are configured into two major subsystems; mission critical and business critical. SQL workloads, deemed critical to the business, execute from the mission critical subsystem on a pair of DL380p servers while remaining SQL workloads execute in VMs in the business critical subsystem. Each subsystem supports one or more instances of SQL Server.

Server and database

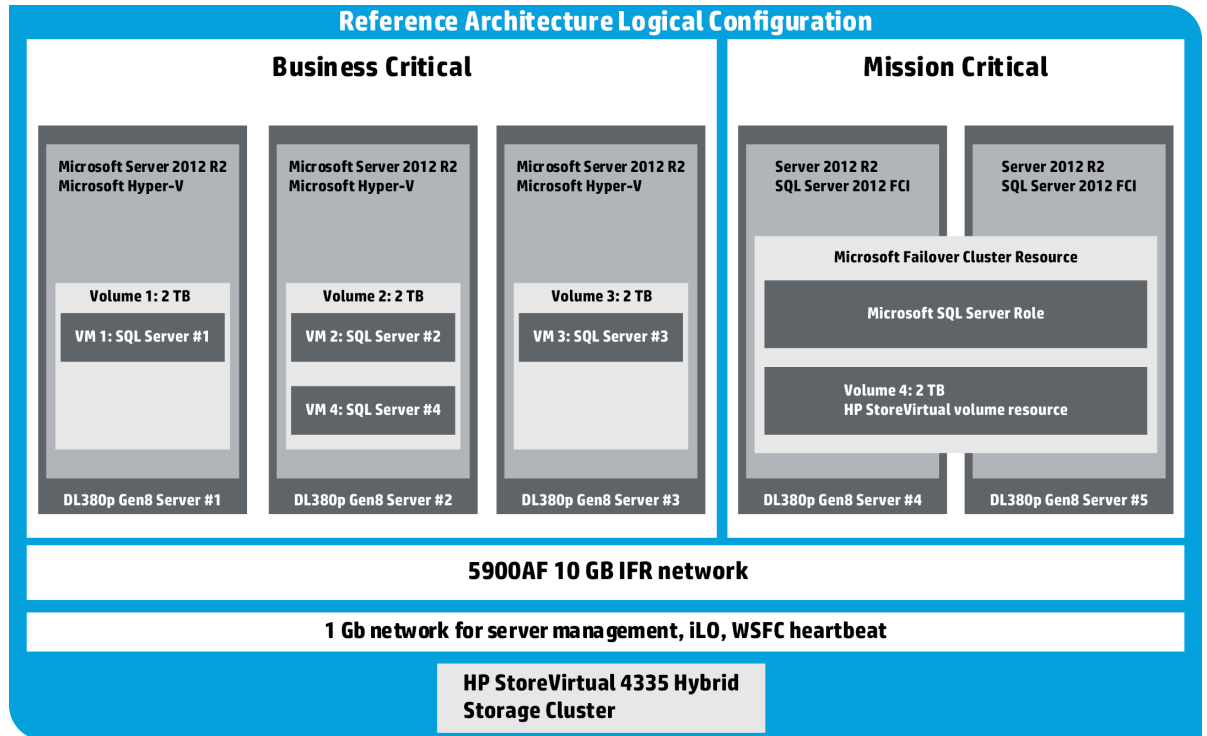
Mission critical SQL server applications run on either ProLiant server in the Microsoft Failover Cluster. The RA specifies a single two TB volume provisioned from the StoreVirtual 4335 by the CMC. Characteristics of the volume are thin provisioning, Network RAID 10, and AO Permitted. The volume is assigned by the CMC to both mission critical servers. The volume is a Microsoft Failover Cluster disk resource for storing SQL Server databases which is shared between the servers. SQL Server instances automatically transition between servers during planned or unplanned service interruptions or downtime events.

The business critical servers use three additional two TB volumes provisioned from the StoreVirtual 4335 cluster; one volume is assigned to each of the business critical servers. Volume characteristics are thin provisioned, Network RAID 10, and AO Permitted. Use cases call for initially setting volume AO to Not Permitted for demonstration purposes only. Changing the AO volume property can act as a function in a larger QoS policy plan. Hyper-V VMs running SQL Server are placed across the servers to balance the business critical workloads. This accommodates SQL Server workloads that may be drained from one server and placed on either of the remaining two servers, allowing planned server downtime to occur without interrupting SQL Server workloads.

Note

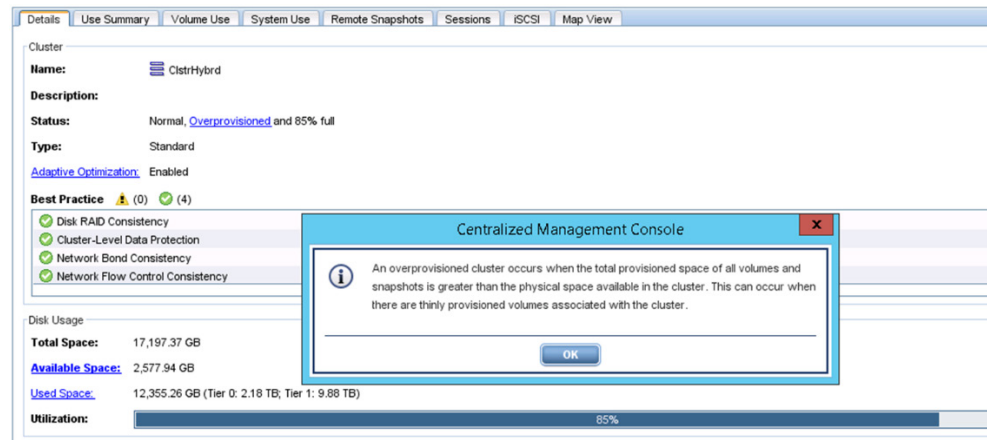
Unplanned server downtime may interrupt any workloads currently provisioned to the business critical subsystem. SQL Server Failover Cluster Instance in the mission critical subsystem mitigates risks of unplanned events. Consider the business requirements carefully when planning and organizing the placement of database workloads into the subsystems.

Figure 8. Reference Architecture logical configuration



Thin provisioning the volumes reduces costs by increasing storage efficiency, allocating space only as data is written without requiring pre-allocation of storage. Overall use and efficiency of the HP StoreVirtual Storage grows, reducing costs and ultimately increasing ROI. Although StoreVirtual can be provisioned any number of ways, this RA uses four thin two TB Network RAID 10 volumes for databases plus additional smaller volumes indirectly associated with the creation of the RA itself. The aggregate size of all volumes slightly overprovisions the capacity of the array and uses approximately 85 percent of the usable capacity after all databases for the use cases are instantiated. The CMC indicates this condition with warnings and context sensitive help, shown in figure 9, Overprovisioning and utilization warning.

Figure 9. Overprovisioning and utilization warning

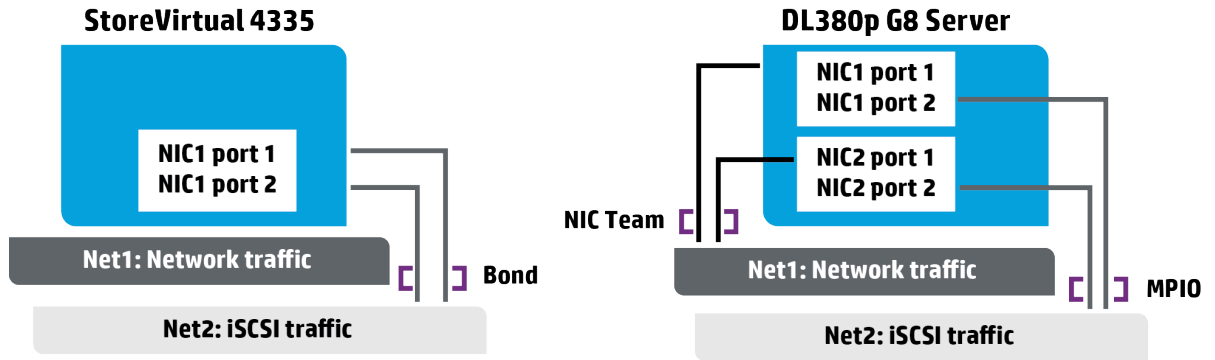


The warnings also notify the storage administrator of the possibility that the LeftHand OS could prevent further writes from occurring if the available space were actually fully utilized by the SQL Server workloads. Additional storage capacity is very easy to add via the CMC and additional storage systems.

Network detail

Server network traffic is handled by a pair of 10 Gb network ports shown in figure 10, Network Adapter Detail. The pair primarily provide for link redundancy and data throughput; they are configured as a Windows NIC team. This RA specifies that the NIC port pair come from separate NC552 PCI Server Adapters to mitigate risk of a single PCI adapter failure.

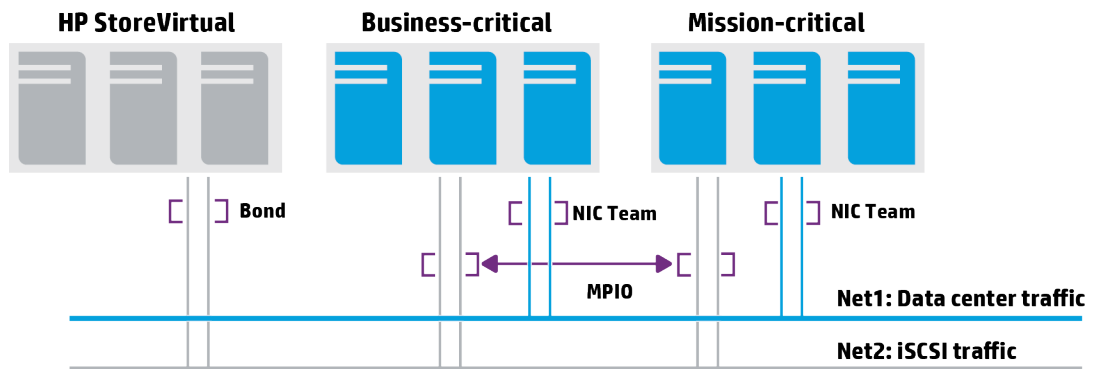
Figure 10. Detail of typical network adapter connections



A second dedicated network manages iSCSI traffic and is isolated from the Ethernet network. Link redundancy provides network performance and mitigates risk of individual link failure. A pair of redundant SAN links, managed by Windows MPIIO, handles failover events on the two paths. The HP StoreVirtual DSM for Microsoft MPIIO provides enhanced MPIIO functionality with additional IO paths to the storage systems. These additional logical paths increase the fault tolerance of the iSCSI network.

The completed environment for the RA (figure 11) contains the clustered StoreVirtual 4335 systems, the two node WSFC, and the three Hyper-V nodes. StoreVirtual traffic is contained within the data center on network Net2, an isolated network (gray), with application and SQL Server traffic contained on network Net1 (blue).

Figure 11. Corporate data center solution with network configuration



Microsoft iSCSI Initiator and MPIIO manages both network paths in the iSCSI network. HP StoreVirtual DSM for MPIIO creates and manages additional iSCSI sessions and load balances across the sessions.

The CMC lists the iSCSI sessions (figure 12) on a given volume for the two servers in the mission critical subsystem. Successful implementation of HP DSM for Microsoft MPIIO is characterized by the additional iSCSI sessions. Together, Microsoft MPIIO, HP DSM, and a round robin policy provide multiple active/active data paths to storage.

Figure 12. DSM iSCSI sessions for a single volume

← iSCSI Sessions (16)

Server	Initiator Node Name	Gateway Connection	Initiator IP:Port	Identifier
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-1 (192.168.20.41)	192.168.20.134/49164	400001370003
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-2 (192.168.20.42), DSM	192.168.20.134/49246	400001370002
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-3 (192.168.20.43), DSM	192.168.20.134/49252	400001370006
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-1 (192.168.20.41), DSM	192.168.20.134/49281	400001370008
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-1 (192.168.20.41)	192.168.20.135/49263	400001370003
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-2 (192.168.20.42), DSM	192.168.20.135/49280	400001370000
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-3 (192.168.20.43), DSM	192.168.20.135/49282	400001370005
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-1 (192.168.20.41), DSM	192.168.20.135/49284	400001370007
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-1 (192.168.20.41)	192.168.20.234/49162	400001370004
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-2 (192.168.20.42), DSM	192.168.20.234/49245	400001370000
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-1 (192.168.20.41), DSM	192.168.20.234/49265	400001370007
svr1	iqn.1991-05.com.microsoft:svr1.corp...	4335-3 (192.168.20.43), DSM	192.168.20.234/49395	400001370005
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-1 (192.168.20.41)	192.168.20.235/49265	400001370004
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-2 (192.168.20.42), DSM	192.168.20.235/49281	400001370002
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-3 (192.168.20.43), DSM	192.168.20.235/49283	400001370006
svr2	iqn.1991-05.com.microsoft:svr2.corp...	4335-1 (192.168.20.41), DSM	192.168.20.235/49265	400001370008

Best practice

HP recommends Multipath IO and the HP DSM with the round robin policy on servers and ALB load balancing policy on HP StoreVirtual storage.

Sizing HP StoreVirtual Storage

The size of the databases used by the OLTP and data warehousing workloads will be between 750 GB and 1 TB. Five StoreVirtual volumes host six databases. Characteristics of the volumes are:

- 2-3 TB
- Network RAID 10
- Thin provisioned

A data protection level of Network RAID 10 will provision twice the configured size of the volume to hold the mirror copy. The six databases will consume up to 12 TB of the 17 TB capacity of the 3-node StoreVirtual 4335 cluster.

A mission critical SQL server instance is installed into a 2-node WSFC. The WSFC provides HA to the SQL application. Although this subsystem is configured with a single instance, the 2-node WSFC on the ProLiant DL380p Gen8 servers support more than one instance of SQL Server.

Business critical SQL Server instances run on Hyper-V VMs hosted by three ProLiant servers. This subsystem provides business continuity for the database applications, allowing them to execute from any of the three ProLiant servers. It also allows for server maintenance windows or planned downtime. The three business critical volumes are sized to hold additional VMs that may move from the other servers within the subsystem.

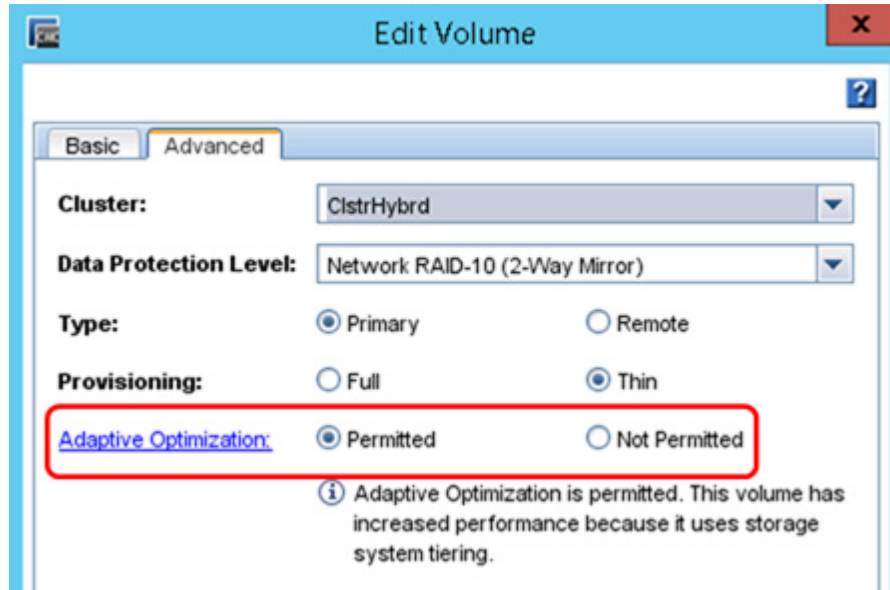
Together, StoreVirtual 4335 and the two server subsystems support a mixed use environment of SQL Server end-user application requirements, providing HA for small and midsize customers.

OLTP type workload description

The use case demonstrates the benefits of AO and the added performance of SSD storage capacity in the StoreVirtual 4335. This demonstration executes an OLTP-type workload in a volume initially provisioned in the slower Tier 1 space, then enabling AO on the volume to allow the LeftHand OS to automatically relocate blocks in the volume to the faster Tier 0 space. The relocation or migration of blocks continues, based on the AO algorithm, which determines which blocks to move between tiers. Only blocks that meet the conditions of the algorithm are moved to Tier 0 space, thereby conserving the valuable Tier 0 space for the most actively accessed data.⁴

⁴ For more information on HP StoreVirtual Adaptive Optimization, visit h20195.www2.hp.com/V2/GetPDF.aspx/4AAA4-9000ENW.pdf.

Figure 13. An example of editing the advanced properties of a volume to permit AO behavior



Note

AO reserves a small amount of Tier 0 space for new writes. As Tier 0 space approaches 100 percent utilization, and if there is continued demand for space, AO migrates pages from Tier 0 to Tier 1 based on the heat map algorithm in the LeftHand OS.

The SQL DB workload consists of OLTP like transactions of 8 KB sized blocks and a random read/write ratio of 80/20 on a single StoreVirtual volume. SQL Server cache constraints are used to adjust the length of IO latency to 35 ms as the database workload runs in Tier 1 space. After IO latency stabilizes at 35 ms, performance metrics is logged during the remainder of the AO use case. Workload and metric logging continue for approximately 10 minutes to establish a T1 baseline for latency before AO is enabled on the volume.

At the end of 60 minutes, performance logging and workload are stopped.

OLTP workload data/results

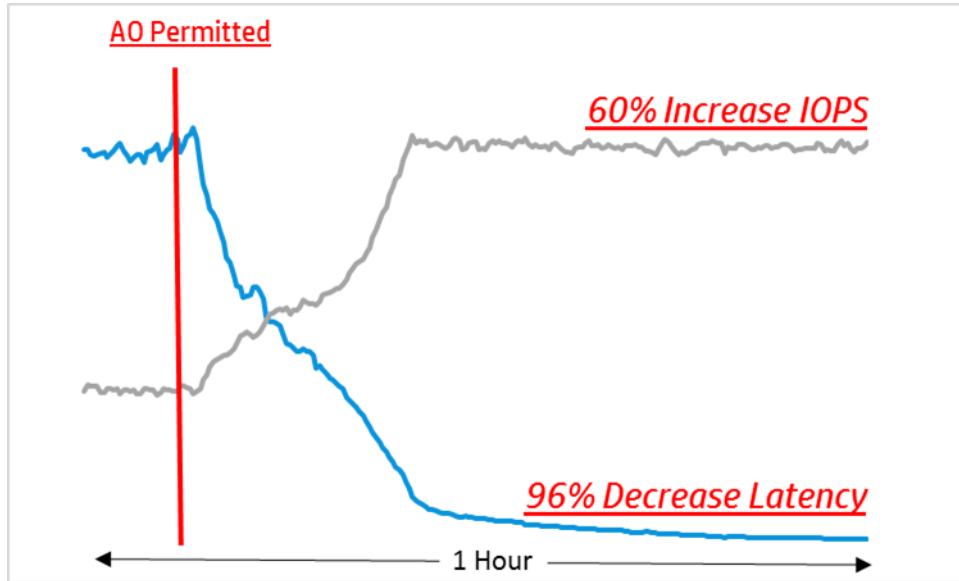
Figure 14 shows IO throughput (gray) and latency (blue) data collected by Windows Performance Monitor.⁵ The collection period was one hour for the OLTP-type workload executing on a single StoreVirtual volume. AO quickly identifies hot pages and moves one of the copies to Tier 0.

Note

It is important to note that after running the OLTP-type workload for one hour, only 22 percent of the 1 TB database occupied T0 space; the remainder continued to occupy T1 space, demonstrating the remarkable efficiency of StoreVirtual AO with workloads characterized by frequent, localized IO accesses.

⁵ Percent change = (delta/original value)*100

Figure 14. Effect of enabling AO on a volume



The same timeframe shows (figure 15 SQL Batch Requests) the performance increase in the Windows Performance Monitor SQL statistic SQL Batch Requests, a measure of Transact-SQL command batches received per second. This statistic is affected by many constraints (such as I/O performance, size of database, size of SQL memory, complexity of requests, and so on). StoreVirtual AO is the only parameter that changes in this use case. Higher batch requests translate to better SQL Server throughput. One final observation is the 60 percent increase in SQL Batch Requests per second just 15 minutes after enabling AO.

Figure 15. SQL Batch Requests

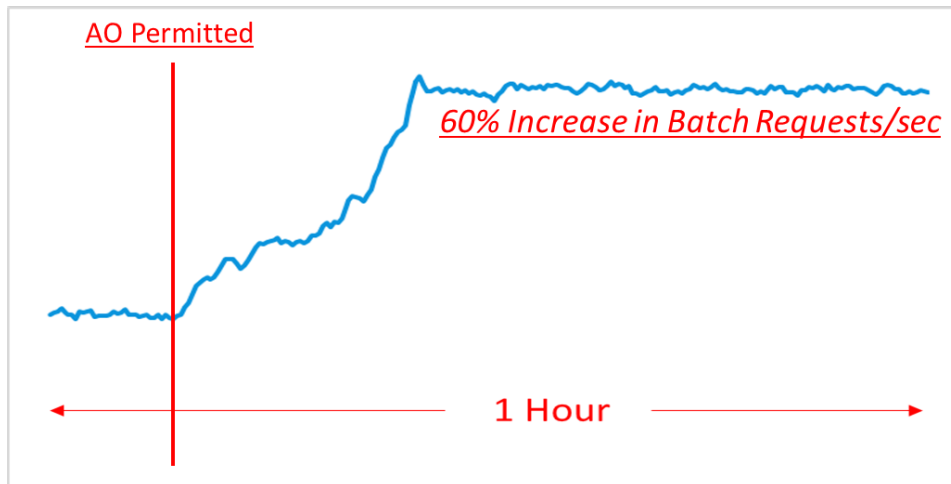
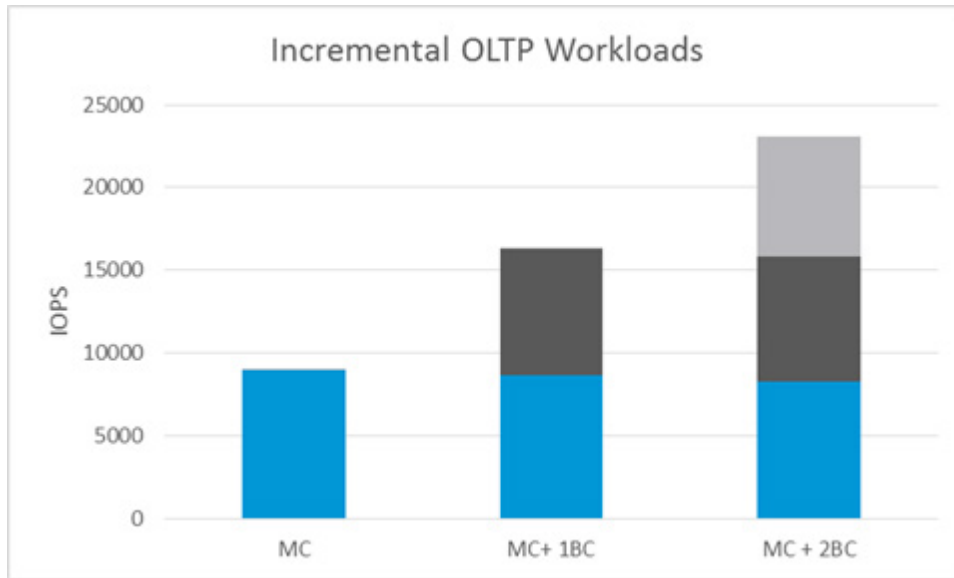


Figure 16 demonstrates that adding business critical workloads has only a small effect on the throughput performance of the mission critical workload. This use case compares backend IO throughput of the 3-node StoreVirtual 4335 cluster with increasing OLTP workloads. In this RA configuration, the rate of throughput performance begins to taper or “knee” beyond the three workloads. The average IO latency is under 2 ms. As total IOPS will vary with any given workload type, determine the optimal business load by monitoring storage performance using the performance tools available in the CMC. Workloads will scale linearly with additional StoreVirtual 4335 Hybrid Storage systems.

Figure 16. Incremental OLTP workloads



Data-mining like workload description

Ad-hoc queries found in data mining workloads extract business information from large datasets. This use case will compare four scenarios that demonstrate AO effectiveness with SQL ad-hoc queries. Data-mining workload IOs are larger than the OLTP IOs, approximately 256 KB, and sequential in nature. A database workload tool, HammerDB,⁶ executes a fixed number of database queries and measures the amount of time to complete the set of queries. A 3-node HP P4500 G2 cluster consisting of 36 15K RPM disks provides a reference point for comparison.

The four scenarios are described as follows:


- P4500 G2-Workload on a three node HP P4500 G2 (for reference purposes)
- 4335 T1-Workload in Tier 1 of the 3- node StoreVirtual 4335 (AO not permitted)
- 4335 T1/T0-Workload on StoreVirtual 4335 data transitioning to Tier 0 (AO permitted)
- 4335 T0–Steady state data resides in both Tier 1 and Tier 0 (AO permitted)

Data mining workload data/results

The results in figure 17 show the performance benefit when using AO permitted for DB volumes running 256 KB block IO reads. The time to run the benchmark is cut by half when compared with the same workload running out of conventional spindle based storage. The P4500 G2 result uses 36 15K RPM disks while the comparative result for the 4335 Tier 1 result uses 21 10K RPM disks.

⁶ HammerDB is an open source database load testing and benchmarking tool available from sourceforge.net/.

Figure 17. Comparison of ad-hoc query completion time



Bill of materials

Table 1. Hardware bill of materials

Quantity	Description
HP ProLiant servers	
5	HP ProLiant DL380p Gen8 2P Intel Xeon CPU E5-2690 @ 2.90 GHz (653200-B21)
3	HP StoreVirtual 4335 Hybrid Storage (F3J70A)
10	HP NC552SFP 10Gb 2-port Ethernet Server Adapter (614203-B21)
HP Networks	
2	HP 5900AF-48XG-4QSFP+ Switch (JC772A)
2	HP 58x0AF Back (power side) to Front (port side) Airflow Fan Tray (JC682A)
2	HP 58x0AF 650W AC Power Supply (JC680A)
Miscellaneous hardware	
1	HP rack and KVM switch with keyboard and monitor
16	HP X240 10G SFP+ to SFP+ Direct Attach Copper Cable
1	HP Networks E3800 Switch (J9575A)
2	HP X240 40G QSFP+ to QSFP+ 1m Direct Attach Copper Cable (JG326A)

Table 2. Tested software

Hewlett-Packard software
LeftHand OS 11.5 (delivered with StoreVirtual storage)
HP StoreVirtual Centralized Management Console
HP StoreVirtual DSM for Microsoft MPIO
Microsoft software
Microsoft Windows Server 2012 R2
Microsoft SQL Server 2014
Windows .NET 3.5

Summary

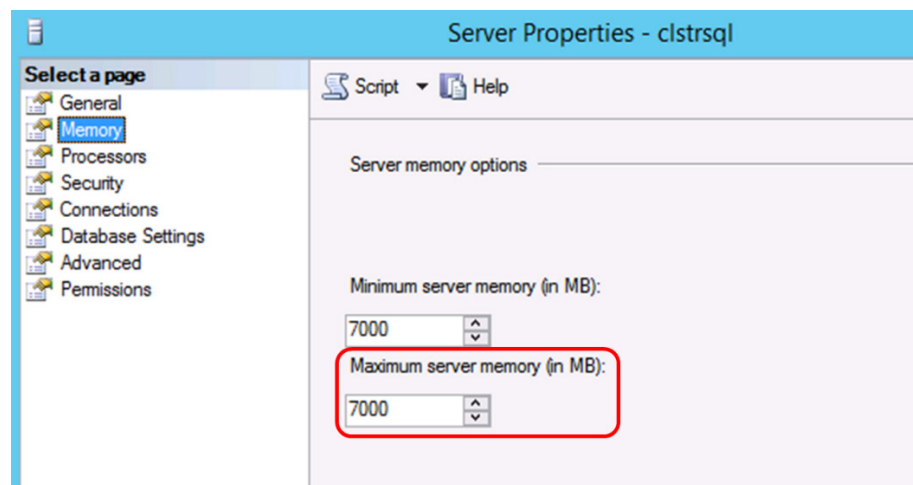
HP StoreVirtual, together with Microsoft SQL Server, delivers the new style of IT, creating a compelling combination for SMBs. Using SQL Server virtualization technologies and affordable but powerful solutions for mission critical databases, this solution provides the scalability and HA to meet the increasing demand for easily managed, affordable storage infrastructure. Pooling internal server storage with Network RAID technology enables scalable capacity and throughput that grow with your business and geographic awareness for stretched clustering to protect your data. Virtualization and clustering technologies manage databases in support of business continuity or even routine maintenance. Combine these innovations with Microsoft SQL Server for a database storage solution worthy of five nine’s availability and reliability.

Appendix

SQL Server

For demonstration purposes only, this RA constricted the amount of cache to the SQL Server instances. The constriction changes IO activity on StoreVirtual and is not a performance recommendation from HP.

Figure A-1. SQL server properties



HP Networking

The network solution recommended is for a basis or starting point only. Consider the networks to be a subset of a larger data center solution for HA or disaster recovery support.

Database workload tools

Two workload tools were used in the use cases, first to create and then to exercise the databases.

- Data mining workload provided by HammerDB, available through hammerora.sourceforge.net/.
- The OLTP brokerage-account type workload is a propriety tool, not available for distribution.

Glossary

Adaptive Optimization (AO): A technology in LeftHand OS that allows data to be stored on storage system media, called Tier 0 and Tier 1, which have different performance characteristics. AO positions data on either storage tier based on previous access history.

Cluster: A grouping of storage systems that resides in a management group. A cluster provides a single pool of storage from which volumes can be created and presented to client servers. The volumes seamlessly span the storage systems in the cluster. The cluster capacity can be expanded by adding additional storage systems to the cluster.

LeftHand operating system: The storage server operating system running on all StoreVirtual storage systems. The OS running on StoreVirtual VSA is identical to that on the dedicated hardware platforms and provides the same features.

Management group: A logical entity for organizing StoreVirtual clusters and storage nodes. It is the top-level container, which can contain one or more clusters or storage nodes.

Network RAID: StoreVirtual technology for maintaining replicated copies of data across a cluster. Network RAID (greater than Network RAID 0) ensures HA of data.

Page: An individual block of data residing on a physical disk. For StoreVirtual devices, the page size is 256 KB.

Storage system: An individual instance of the StoreVirtual LeftHand Operating System running on a single storage server or as a VM (in the case of the StoreVirtual VSA).

Tier: A grouping of like RAID devices internal to a StoreVirtual storage node (server or VM) that is provided a hierarchical designation based on the storage device's performance capability.

- Tier 0—The fastest tier of storage in a hybrid/auto-tiering solution.
- Tier 1—The slower tier of storage in a hybrid/auto-tiering solution.

For more information

HP StoreVirtual Adaptive Optimization white paper

h20195.www2.hp.com/V2/GetPDF.aspx/4AA4-9000ENW.pdf

Additional resources

To download user manuals and other support documentation, go to hp.com/support/storevirtualmanuals

To download the latest compatibility matrix, go to hp.com/go/storevirtualcompatibility

To download software components for LeftHand OS, go to h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=StoreVirtualSW_11.0

To download HP StoreFront Manager for Microsoft Software, go to h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=System_Center

HP Servers and Services

HP ProLiant servers, hp.com/go/Servers

HP Services, hp.com/go/Services

HP Networking

HP Networking, hp.com/go/networking

To help us improve our documents, please provide feedback at hp.com/solutions/feedback.

Learn more at

hp.com/go/StoreVirtual4000

Sign up for updates

hp.com/go/getupdated



Share with colleagues



Rate this document

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, and Windows Server are trademarks of the Microsoft Group of companies. Intel and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries.

