



West Virginia Purchasing Division

2019 Washington Street, East
Charleston, WV 25305
Telephone: 304-558-2306
General Fax: 304-558-6026
Bid Fax: 304-558-3970

The following documentation is an electronically-submitted vendor response to an advertised solicitation from the *West Virginia Purchasing Bulletin* within the Vendor Self-Service portal at ***wvOASIS.gov***. As part of the State of West Virginia's procurement process, and to maintain the transparency of the bid-opening process, this documentation submitted online is publicly posted by the West Virginia Purchasing Division at ***WVPurchasing.gov*** with any other vendor responses to this solicitation submitted to the Purchasing Division in hard copy format.

Header

List View

General Information | [Contact](#) | [Default Values](#) | [Discount](#) | [Document Information](#)

Procurement Folder: 237208

SO Doc Code: CRFQ

Procurement Type: Central Contract - Fixed Amt

SO Dept: 0210

Vendor ID:

SO Doc ID: ISC1700000002

Legal Name: CINCINNATI BELL TECHNOLOGY SOLUTIONS INC

Published Date: 9/7/16

Alias/DBA:

Close Date: 9/14/16

Total Bid: \$222,395.76

Close Time: 13:30

Response Date:

Status: Closed

Response Time:

Solicitation Description:

Total of Header Attachments: 0

Total of All Attachments: 0



Purchasing Division
 2019 Washington Street East
 Post Office Box 50130
 Charleston, WV 25305-0130

**State of West Virginia
 Solicitation Response**

Proc Folder : 237208
Solicitation Description : Addendum #3 SAN Storage Systems (OT1705)
Proc Type : Central Contract - Fixed Amt

Date issued	Solicitation Closes	Solicitation Response	Version
	2016-09-14 13:30:00	SR 0210 ESR09141600000001069	1

VENDOR
000000223374 CINCINNATI BELL TECHNOLOGY SOLUTIONS INC

Solicitation Number: CRFQ 0210 ISC1700000002

Total Bid : \$222,395.76 **Response Date:** 2016-09-14 **Response Time:** 12:18:46

Comments:

FOR INFORMATION CONTACT THE BUYER
 Stephanie L Gale
 (304) 558-8801
 stephanie.l.gale@wv.gov

Signature on File	FEIN #	DATE
--------------------------	---------------	-------------

All offers subject to all terms and conditions contained in this solicitation

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
1	San Storage Solutions	2.00000	EA	\$93,977.060000	\$187,954.12

Comm Code	Manufacturer	Specification	Model #
43210000			

Extended Description : SAN Storage Solutions

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
2	Maintenance and Support				\$34,441.64

Comm Code	Manufacturer	Specification	Model #
81112201			

Extended Description : Initial Year Maintenance and Support

ATTACHMENT A – Pricing Page

Commodity Line Number	Description	Unit of Measure	Quantity	Unit Cost	Extended Cost
Commodity Line 1, Specification 3.1.1	SAN Storage Solutions- Charleston and Clarksburg	Each	2	93977.06	187,954.12
Commodity Line 2, Specification 3.1.2	4 Year Maintenance & Support Warranty	Each	2	17220.82	34,441.64
Total Cost					222,395.76

ATTACHMENT B - Complete Solution Item List Template

Part #	Detailed Description	Unit of Measure	Quantity
K2Q36A	HPE 3PAR StoreServ 8200 2N Fld Int Base	\$ 4,057.45	2
K2P89A	HPE 3PAR 8000 1.92TB SFF SSD	\$ 4,528.39	32
K2P89A 0D1	Factory integrated	\$ -	32
L7B45A	HPE 3PAR 8200 OS Suite Base LTU	\$ 949.02	2
L7B45A 0D1	Factory integrated	\$ -	2
L7B46A	HPE 3PAR 8200 OS Suite Drive LTU	\$ 45.69	32
L7B46A 0D1	Factory integrated	\$ -	32
L7B49A	HPE 3PAR 8200 Replication Suite Base LTU	\$ 494.54	2
L7B49A 0D1	Factory integrated	\$ -	2
L7B50A	HPE 3PAR 8200 Replication Ste Drive LTU	\$ 44.99	32
L7B50A 0D1	Factory integrated	\$ -	32
BD440A	HPE 3PAR 8200 File Persona Ste 1TB LTU	\$ 59.40	2
BD440A 0D1	Factory integrated	\$ -	2
HA114A1	HP Installation and Startup Service	\$ -	1
HA114A1 5XU	HPE Startup 3PAR 8200 2N Fld Int Bas SVC	\$ 2,093.30	2
K2R29A	HPE 3PAR StoreServ RPS Service Processor	\$ 1,536.68	2
BD362AAE	HPE 3PAR StoreServ Mgmt/Core SW E-Media	\$ 5.51	2
BD363AAE	HPE 3PAR OS Suite Latest E-Media	\$ 5.51	2
BD365AAE	HPE 3PAR SP SW Latest E-Media	\$ 1.38	2
H1K92A4	HPE 4Y Proactive Care 24x7 Service	\$ -	1
H1K92A4 WSF	HPE 3PAR Internal Entitlement Supp	\$ -	6
H1K92A4 YNW	HPE 3PAR StoreServ RPS Service Proc Supp	\$ 323.48	2
H1K92A4 YT8	HPE 3PAR StoreServ 8200 2N Base Support	\$ 788.97	2
H1K92A4 YU0	HPE 3PAR 8000 1.92TB SFF SSD Supp	\$ 732.10	32
H1K92A4 YUA	HPE 3PAR 8200 OS Suite Base Support	\$ 2,522.72	2
H1K92A4 YUB	HPE 3PAR 8200 OS Suite Drive Support	\$ 16.89	32
H1K92A4 YUE	HPE 3PAR 8200 ReplicationSuite Base Supp	\$ 1,168.47	2
H1K92A4 YUF	HPE 3PAR 8200 Replication Ste Drive Supp	\$ 24.12	32
H1K92A4 YXA	HPE 3PAR 8200 File Persona Ste 1TB Supp	\$ 47.22	2
QK734A	HPE Premier Flex LC/LC OM4 2f 5m Cbl	\$ 36.00	8
HA124A1	HP Technical Installation Startup SVC	\$ -	1
HA124A1 5Y5	HPE Startup 3PAR 8000 System Reportr SVC	\$ 1,046.65	2
HA124A1 5Y9	HPE Startup 3PAR 8000 Fil Persona St SVC	\$ 3,110.04	2
HA124A1 5Y8	HPE Startup 3PAR 8K Rpl Ste VC-RC-PP SVC	\$ 4,545.45	2
CBTS PS	CBTS Professional Services	\$ 175.00	20
SHIPPING		\$ 545.87	1

Quote No: AR4-00582-00
 Project Name: West Virginia
 Created On: 13 Sep 2016
 Expiration Date: 13 Oct 2016
 Account Manager: Ashley Howard
 Inside Sales: Lisa Lyons lisa.lyons@cbts.net
 CBTS Address: 221 East 4th St., Cincinnati, Ohio, 45202 Fax-513-841-6738



Qty	Part Number	Description	Unit Price	Extended Price
2	K2Q36A	HPE 3PAR StoreServ 8200 2N Fid Int Base	\$ 4,057.45	\$ 8,114.89
32	K2P89A	HPE 3PAR 8000 1.92TB SFF SSD	\$ 4,528.39	\$ 144,908.59
32	K2P89A 0D1	Factory integrated	\$ -	\$ -
2	L7B45A	HPE 3PAR 8200 OS Suite Base LTU	\$ 949.02	\$ 1,898.04
2	L7B45A 0D1	Factory integrated	\$ -	\$ -
32	L7B46A	HPE 3PAR 8200 OS Suite Drive LTU	\$ 45.69	\$ 1,462.13
32	L7B46A 0D1	Factory integrated	\$ -	\$ -
2	L7B49A	HPE 3PAR 8200 Replication Suite Base LTU	\$ 494.54	\$ 989.09
2	L7B49A 0D1	Factory integrated	\$ -	\$ -
32	L7B50A	HPE 3PAR 8200 Replication Ste Drive LTU	\$ 44.99	\$ 1,439.66
32	L7B50A 0D1	Factory integrated	\$ -	\$ -
2	BD440A	HPE 3PAR 8200 File Persona Ste 1TB LTU	\$ 59.40	\$ 118.81
2	BD440A 0D1	Factory integrated	\$ -	\$ -
1	HA114A1	HP Installation and Startup Service	\$ -	\$ -
2	HA114A1 5XU	HPE Startup 3PAR 8200 2N Fid Int Bas SVC	\$ 2,093.30	\$ 4,186.60
2	K2R29A	HPE 3PAR StoreServ RPS Service Processor	\$ 1,536.68	\$ 3,073.36
2	BD362AAE	HPE 3PAR StoreServ Mgmt/Core SW E-Media	\$ 5.51	\$ 11.02
2	BD363AAE	HPE 3PAR OS Suite Latest E-Media	\$ 5.51	\$ 11.02
2	BD365AAE	HPE 3PAR SP SW Latest E-Media	\$ 1.38	\$ 2.77
1	H1K92A4	HPE 4Y Proactive Care 24x7 Service	\$ -	\$ -
6	H1K92A4 WSF	HPE 3PAR Internal Entitlement Supp	\$ -	\$ -
2	H1K92A4 YNW	HPE 3PAR StoreServ RPS Service Proc Supp	\$ 323.48	\$ 646.96
2	H1K92A4 YT8	HPE 3PAR StoreServ 8200 2N Base Support	\$ 788.97	\$ 1,577.93
32	H1K92A4 YU0	HPE 3PAR 8000 1.92TB SFF SSD Supp	\$ 732.10	\$ 23,427.36
2	H1K92A4 YUA	HPE 3PAR 8200 OS Suite Base Support	\$ 2,522.72	\$ 5,045.45
32	H1K92A4 YUB	HPE 3PAR 8200 OS Suite Drive Support	\$ 16.89	\$ 540.55
2	H1K92A4 YUE	HPE 3PAR 8200 ReplicationSuite Base Supp	\$ 1,168.47	\$ 2,336.95
32	H1K92A4 YUF	HPE 3PAR 8200 Replication Ste Drive Supp	\$ 24.12	\$ 772.00
2	H1K92A4 YXA	HPE 3PAR 8200 File Persona Ste 1TB Supp	\$ 47.22	\$ 94.44
8	QK734A	HPE Premier Flex LC/LC OM4 2f 5m Cbl	\$ 36.00	\$ 288.00
1	HA124A1	HP Technical Installation Startup SVC	\$ -	\$ -
2	HA124A1 5Y5	HPE Startup 3PAR 8000 System Reportr SVC	\$ 1,046.65	\$ 2,093.30
2	HA124A1 5Y9	HPE Startup 3PAR 8000 Fil Persona St SVC	\$ 3,110.04	\$ 6,220.09
2	HA124A1 5Y8	HPE Startup 3PAR 8K Rpl Ste VC-RC-PP SVC	\$ 4,545.45	\$ 9,090.89
20		CBTS Professional Services	\$ 175.00	\$ 3,500.00
		Shipping		\$ 545.87
		Total		\$ 222,395.76

**** Freight & Taxes Not Included In Total ****

Client Signature _____

Please fax this bill of materials to 513-841-5072 for purchase.

Statement of Confidentiality

Attention notification regarding: General Terms and Conditions (Warranty (a), (b) & (c))

CBTS is a Dell Enterprise Authorized Reseller and does not manufacture the equipment we resell or develop the software that we license. However, CBTS does pass through all equipment and software warranties as provided by the equipment manufacturer or software publisher.

REQUEST FOR QUOTATION
Storage Area Network Solution Hardware, Software,
Maintenance

2.8 "VMS" means Virtual Machines

2.9 "LUN" means Logical Unit Number

2.10 "WAN" means Wide Area Network

2.11 "MS" means milliseconds

3. GENERAL REQUIREMENTS:

3.1 Mandatory Contract Item Requirements: Contract Item must meet or exceed the mandatory requirements listed below.

3.1.1 Vendor must provide, implement and configure two geographically separated Storage Area Network (SAN) systems (hardware and software) and related controllers, within 60 days of notice to proceed, each with the following specifications.

3.1.1.1 SAN systems must have a minimum of twenty-four (24) terabytes of usable storage before compression and deduplication.

3PAR 8200 is the all flash platform that is recommended as the solution for West Virginia totaling 24TB of usable SSD storage.

3.1.1.2 SAN system must be compatible and interoperable with the following systems and applications:

Refer to document "Cisco Interoperability" for the following applications.

3.1.1.2. 1 Cisco Unified Computing System (UCS), comprised of: Cisco B200 M3 Blade Center, Cisco C220 M3, C240 M3, and C240 M4 Servers, Cisco 6248UP Fiber Interconnect Switch

3.1.1.2.2 Cisco Unified Communications Manager (CUCM) Version 10.x

3.1.1.2.3 Cisco Unity Connection (CUC) Version 10.x

3.1.1.2.4 Cisco Unified Contact Center Express (UCCX)

– Enhanced and Premium Version 10.x

- 3.1.1.2.5** Cisco Integrated Management Controller (CIMC)
- 3.1.1.2.6** Cisco Unified IM Presence (CUP) Version 10.x
- 3.1.1.2.7** Cisco Prime Collaboration Deployment (PCD) Version 10.x
- 3.1.1.2.8** Cisco Prime Collaboration Provisioning (PCP) Version 10.x
- 3.1.1.2.9** Cisco Prime Collaboration Assurance (PCA) Version 10.x
- 3.1.1.2.10** Cisco Expressway Version 8.x
- 3.1.1.2.11** Cisco Jabber
- 3.1.1.2.12** VMware ESXi 5 Enterprise (Version 5.5 U3)
- 3.1.1.2.13** VMware vCenter Version 5
- 3.1.1.2.14** Singlewire Software InformaCast IP Paging Version 11.x
- 3.1.1.2.15** Bridge Communication – Bridge Operator Console Version 3.x

3.1.1.3 Must provide all software and licenses for system management, data backup (including VMs, managed by VMware vCenter) and replication between systems.

Included in Price

3.1.1.4 Must support data compression, deduplication and thin provisioning.

At the heart of every HPE 3PAR StoreServ node there is an HPE 3PAR ASIC that features an efficient, silicon-based zero-detection and hashing mechanism. This unique hardware capability gives HPE 3PAR StoreServ Storage the power to perform inline deduplication and remove allocated but unused space inline and non-disruptively

Reducing capacity requirements by mitigating overprovisioning; using enterprise-class data compaction technologies; and enabling fast, simple, automated space reclamation are essential to the industry-leading efficiency of HPE 3PAR StoreServ Storage systems. Compaction technologies such as thin provisioning, thin deduplication, and thin reclamation offer efficiency benefits for primary storage that can significantly reduce both capital and operational costs with spinning media and SSDs.

HPE 3PAR Thin Deduplication software with Express Indexing relies on the HPE 3PAR ASICs to generate and assign a hash key to each unique incoming write request. Express Indexing, a mechanism that accelerates data signature comparison, is used for ultrafast detection of duplicate write requests in order to prevent duplicate data from being written.

See page 30 of 3PAR StoreServ Architecture

3.1.1.5 Systems must be capable of being expanded (scale out) non-disruptively, up to 100 Terabytes (TB) of total uncompressed storage. Performance and capacity shall scale together.

3PAR 8200 can add up to 9 shelves of disk capacity giving a total of 824TB (approx.) raw storage and 140TB of usable file capacity.

See page 3 of 3PAR quickspecs

3.1.1.6 Must meet the following latency storage performance at all times:

3.1.1.6.1 Host-level kernel disk command latency <4 ms (no spikes above) and

3.1.1.6.2 Physical device command latency <20 ms (no spikes above)

This configuration of 3PAR 8200 all flash gives you 111,500 IOPS at or below .4ms latency.

See WVOT 3PAR configuration document

3.1.1.7 Must have the ability to enable or disable compression on a per-LUN basis.

Enabling and disabling compaction is completed at the Volume level

See page 12 in 3PAR best practice.

3.1.1.8 Systems must support asynchronous replication across WAN links.

New to the HPE 3PAR Remote Copy family with the release of HPE 3PAR OS 3.2.2 is Asynchronous Streaming Remote Copy software, which provides the capability to deploy replication solutions that can provide an RPO measured in seconds.

See 3par Remote Copy documentation.

3.1.1.9 Must have the ability to replace individual drives & power supplies while the system is in production without impacting operation.

HPE 3PAR StoreServ Storage supports switching intelligence to automate isolation of disk devices and drive enclosure SAS adapters from the SAS backend to which they are connected once a specific error threshold is reached. This process isolates the faulty component and protects other devices on that backend from other potential failures.

See page 8 on 3par high availability document

3.1.1.10 The storage array shall have no single point of failure.

HPE 3PAR StoreServ Storage is designed to help minimize operational impact by delivering planned resilience. Full hardware redundancy is built into the system, mitigating all single points of failure. When appropriately configured, HPE 3PAR StoreServ Storage is capable of greater than “five nines” or 99.999 percent availability from a hardware perspective. This is reinforced by autonomic and proactive error checking, and software features that deliver seamless failover/failback to help ensure complete system resilience, even when the unexpected happens.

See page 3 on 3par High availability document

- 3.1.1.11** The storage array must be an all flash solution, however the Operating System and associated software may be stored separately on a hard disk drive.

This configuration of 3PAR 8200 model is all flash using cMLC SSD drives.

- 3.1.1.12** Each system must provide a minimum of 4-ports of 8 Gigabit Fibre Channel (FC) connectivity and be upgradeable to 16 Gigabit FC.

The 8200 model is able to handle a minimum of 4 ports up to 12 ports of 16GB fiber channel, backwards compatible to 8GB

See page 3 of 3par quickspec document.

- 3.1.1.13** Systems must support Fibre Channel over Ethernet (FCoE).

- 3.1.1.14** Must provide 1 and 10 gigabit Ethernet connectivity, which should support both file and block access..

The 8200 model handles both 1 and 10GB connectivity.
(8) 1GB Ethernet; (4) 10GB Ethernet; (4) 10GB FCoE;
(4) 10GB ISCSI.

See page 3 of 3par quickspec document.

- 3.1.1.15** Vendor should provide product data sheets showing all part numbers included in their solution with bid, but must provide upon request. Failure to do so may result in disqualification of bid.

Attached

- 3.1.1.16** Vendor must complete Attachment B – Complete Solution Item List Template in full by entering the part # (column A) for every item which is part of the vendor solution. Vendor must enter detailed description of product (column B), the unit of measure (column C) in which it is sold and the quantity (column D) used for both Charleston and Clarksburg. Vendor should provide with bid, but must provide upon request. Failure to do so may result in

disqualification of bid.

Attached

3.1.1.17 The hardware must be listed on the VMware Hardware Compatibility List for SAN Storage (reference list at <http://www.vmware.com/go/hcl>).

HP 3PAR StoreServ 8000	FCoE	ESXi	6.0 U2, 6.0 U1, 6.0, 5.5 U3
HP 3PAR StoreServ 8000	ISCSI	ESXi	6.0 U2, 6.0 U1, 6.0, 5.5 U3
HP 3PAR StoreServ 8000	FC	ESXi	6.0 U2, 6.0 U1, 6.0, 5.5 U3

3.1.1.18 Under the supervision of Agency personnel, vendor must implement, configure and test functionality and compatibility with applications listed in 3.1.1.2.1 to 3.1.1.2.15.

CBTS Professional Services will continue in conjunction with HP to implement, configure and test functionality as described.

3.1.2 Hardware Maintenance and Support Warranty

3.1.2.1 Vendor must provide a four (4) year maintenance service and support warranty for two systems, 8x5xNBD support with the following features:

3.1.2.1.1 24-hours a day, 7 days a week access for telephone support for troubleshooting technical and configuration issues.

3.1.2.1.2 Next-business-day (NBD) advance hardware replacement. Vendor must ship a replacement within the Next Business Day after failure to resolve technical issues using telephone support. Replacement part must be received between the hours of 8:00AM through 5:00PM Eastern Standard Time, Monday through Friday excluding State and Federal holidays and weekends.

HP Proactive Support includes 24/7 and next business day hardware replacement.

3.1.3 Specification for System Acceptance

STATE OF WEST VIRGINIA
Purchasing Division

PURCHASING AFFIDAVIT

MANDATE: Under W. Va. Code §5A-3-10a, no contract or renewal of any contract may be awarded by the state or any of its political subdivisions to any vendor or prospective vendor when the vendor or prospective vendor or a related party to the vendor or prospective vendor is a debtor and: (1) the debt owed is an amount greater than one thousand dollars in the aggregate; or (2) the debtor is in employer default.

EXCEPTION: The prohibition listed above does not apply where a vendor has contested any tax administered pursuant to chapter eleven of the W. Va. Code, workers' compensation premium, permit fee or environmental fee or assessment and the matter has not become final or where the vendor has entered into a payment plan or agreement and the vendor is not in default of any of the provisions of such plan or agreement.

DEFINITIONS:

"Debt" means any assessment, premium, penalty, fine, tax or other amount of money owed to the state or any of its political subdivisions because of a judgment, fine, permit violation, license assessment, defaulted workers' compensation premium, penalty or other assessment presently delinquent or due and required to be paid to the state or any of its political subdivisions, including any interest or additional penalties accrued thereon.

"Employer default" means having an outstanding balance or liability to the old fund or to the uninsured employers' fund or being in policy default, as defined in W. Va. Code § 23-2c-2, failure to maintain mandatory workers' compensation coverage, or failure to fully meet its obligations as a workers' compensation self-insured employer. An employer is not in employer default if it has entered into a repayment agreement with the Insurance Commissioner and remains in compliance with the obligations under the repayment agreement.

"Related party" means a party, whether an individual, corporation, partnership, association, limited liability company or any other form or business association or other entity whatsoever, related to any vendor by blood, marriage, ownership or contract through which the party has a relationship of ownership or other interest with the vendor so that the party will actually or by effect receive or control a portion of the benefit, profit or other consideration from performance of a vendor contract with the party receiving an amount that meets or exceed five percent of the total contract amount.

AFFIRMATION: By signing this form, the vendor's authorized signer affirms and acknowledges under penalty of law for false swearing (W. Va. Code §61-5-3) that neither vendor nor any related party owe a debt as defined above and that neither vendor nor any related party are in employer default as defined above, unless the debt or employer default is permitted under the exception above.

WITNESS THE FOLLOWING SIGNATURE:

Vendor's Name: CBTS

Authorized Signature: [Signature] Date: 8/23/14

State of OHIO

County of HAMILTON, to-wit:

Taken, subscribed, and sworn to before me this 23rd day of August, 2014.

My Commission expires February 11th, 2021.



JAMIE PITSTICK
Notary Public, State of Ohio
My Commission Expires 02-11-2021

NOTARY PUBLIC [Signature]
Purchasing Affidavit (Revised 08/01/2015)

State of West Virginia
VENDOR PREFERENCE CERTIFICATE

Certification and application is hereby made for Preference in accordance with **West Virginia Code**, §5A-3-37. (Does not apply to construction contracts). **West Virginia Code**, §5A-3-37, provides an opportunity for qualifying vendors to request (at the time of bid) preference for their residency status. Such preference is an evaluation method only and will be applied only to the cost bid in accordance with the **West Virginia Code**. This certificate for application is to be used to request such preference. The Purchasing Division will make the determination of the Vendor Preference, if applicable.


1. **Application is made for 2.5% vendor preference for the reason checked:**
 Bidder is an individual resident vendor and has resided continuously in West Virginia for four (4) years immediately preceding the date of this certification; **or**,
 Bidder is a partnership, association or corporation resident vendor and has maintained its headquarters or principal place of business continuously in West Virginia for four (4) years immediately preceding the date of this certification;
 Bidder is a resident vendor partnership, association, or corporation with at least eighty percent of ownership interest of bidder held by another entity that meets the applicable four year residency requirement; **or**,
 Bidder is a nonresident vendor which has an affiliate or subsidiary which employs a minimum of one hundred state residents and which has maintained its headquarters or principal place of business within West Virginia continuously for the four (4) years immediately preceding the date of this certification; **or**,
2. **Application is made for 2.5% vendor preference for the reason checked:**
 Bidder is a resident vendor who certifies that, during the life of the contract, on average at least 75% of the employees working on the project being bid are residents of West Virginia who have resided in the state continuously for the two years immediately preceding submission of this bid; **or**,
3. **Application is made for 2.5% vendor preference for the reason checked:**
 Bidder is a nonresident vendor that employs a minimum of one hundred state residents, or a nonresident vendor which has an affiliate or subsidiary which maintains its headquarters or principal place of business within West Virginia and employs a minimum of one hundred state residents, and for purposes of producing or distributing the commodities or completing the project which is the subject of the bidder's bid and continuously over the entire term of the project, on average at least seventy-five percent of the bidder's employees or the bidder's affiliate's or subsidiary's employees are residents of West Virginia who have resided in the state continuously for the two immediately preceding years and the vendor's bid; **or**,
4. **Application is made for 5% vendor preference for the reason checked:**
 Bidder meets either the requirement of both subdivisions (1) and (2) or subdivision (1) and (3) as stated above; **or**,
5. **Application is made for 3.5% vendor preference who is a veteran for the reason checked:**
 Bidder is an individual resident vendor who is a veteran of the United States armed forces, the reserves or the National Guard and has resided in West Virginia continuously for the four years immediately preceding the date on which the bid is submitted; **or**,
6. **Application is made for 3.5% vendor preference who is a veteran for the reason checked:**
 Bidder is a resident vendor who is a veteran of the United States armed forces, the reserves or the National Guard, if, for purposes of producing or distributing the commodities or completing the project which is the subject of the vendor's bid and continuously over the entire term of the project, on average at least seventy-five percent of the vendor's employees are residents of West Virginia who have resided in the state continuously for the two immediately preceding years.
7. **Application is made for preference as a non-resident small, women- and minority-owned business, in accordance with West Virginia Code §5A-3-59 and West Virginia Code of State Rules.**
 Bidder has been or expects to be approved prior to contract award by the Purchasing Division as a certified small, women- and minority-owned business.

Bidder understands if the Secretary of Revenue determines that a Bidder receiving preference has failed to continue to meet the requirements for such preference, the Secretary may order the Director of Purchasing to: (a) rescind the contract or purchase order; or (b) assess a penalty against such Bidder in an amount not to exceed 5% of the bid amount and that such penalty will be paid to the contracting agency or deducted from any unpaid balance on the contract or purchase order.

By submission of this certificate, Bidder agrees to disclose any reasonably requested information to the Purchasing Division and authorizes the Department of Revenue to disclose to the Director of Purchasing appropriate information verifying that Bidder has paid the required business taxes, provided that such information does not contain the amounts of taxes paid nor any other information deemed by the Tax Commissioner to be confidential.

Bidder hereby certifies that this certificate is true and accurate in all respects; and that if a contract is issued to Bidder and if anything contained within this certificate changes during the term of the contract, Bidder will notify the Purchasing Division in writing immediately.

Bidder: CBTS

Signed: 

Date: 8/23/14

Title: Director of Sales

*Check any combination of preference consideration(s) indicated above, which you are entitled to receive.

ADDENDUM ACKNOWLEDGEMENT FORM
SOLICITATION NO.:

Instructions: Please acknowledge receipt of all addenda issued with this solicitation by completing this addendum acknowledgment form. Check the box next to each addendum received and sign below. Failure to acknowledge addenda may result in bid disqualification.

Acknowledgment: I hereby acknowledge receipt of the following addenda and have made the necessary revisions to my proposal, plans and/or specification, etc.

Addendum Numbers Received:

(Check the box next to each addendum received)

- | | |
|--|--|
| <input checked="" type="checkbox"/> Addendum No. 1 | <input type="checkbox"/> Addendum No. 6 |
| <input checked="" type="checkbox"/> Addendum No. 2 | <input type="checkbox"/> Addendum No. 7 |
| <input checked="" type="checkbox"/> Addendum No. 3 | <input type="checkbox"/> Addendum No. 8 |
| <input checked="" type="checkbox"/> Addendum No. 4 | <input type="checkbox"/> Addendum No. 9 |
| <input type="checkbox"/> Addendum No. 5 | <input type="checkbox"/> Addendum No. 10 |

I understand that failure to confirm the receipt of addenda may be cause for rejection of this bid. I further understand that any verbal representation made or assumed to be made during any oral discussion held between Vendor's representatives and any state personnel is not binding. Only the information issued in writing and added to the specifications by an official addendum is binding.

CBTS
Company


Authorized Signature

8/23/14
Date

NOTE: This addendum acknowledgment should be submitted with the bid to expedite document processing.



Purchasing Division
 2019 Washington Street East
 Post Office Box 50130
 Charleston, WV 25305-0130

State of West Virginia
 Request for Quotation
 21 – Info Technology

Proc Folder: 237208

Doc Description: SAN Storage Systems (OT1705)

Proc Type: Central Contract - Fixed Amt

Date Issued	Solicitation Closes	Solicitation No	Version
2016-08-08	2016-08-31 13:30:00	CRFQ 0210 ISC1700000002	1

BID RECEIVING LOCATION

BID CLERK
 DEPARTMENT OF ADMINISTRATION
 PURCHASING DIVISION
 2019 WASHINGTON ST E
 CHARLESTON WV 25305
 US

VENDOR

Vendor Name, Address and Telephone Number:

FOR INFORMATION CONTACT THE BUYER

Stephanie L Gale
 (304) 558-8801
 stephanie.l.gale@wv.gov

Signature X

FEIN #

DATE

8/23/14

All offers subject to all terms and conditions contained in this solicitation

ADDITIONAL INFORMATION:

The West Virginia Purchasing Division is soliciting bids on behalf of the West Virginia Office of Technology to establish a contract for the one time purchase of hardware, software, licensing, and support for a Storage Area Network (SAN) solution.

INVOICE TO	SHIP TO
DEPARTMENT OF ADMINISTRATION OFFICE OF TECHNOLOGY 1900 KANAWHA BLVD E, BLDG 5 10TH FLOOR CHARLESTON WV25305 US	WV OFFICE OF TECHNOLOGY BLDG 5, 10TH FLOOR 1900 KANAWHA BLVD E CHARLESTON WV 25304 US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
1	San Storage Solutions	2.00000	EA		# 222,395.76

Comm Code	Manufacturer	Specification	Model #
43210000			

Extended Description :
SAN Storage Solutions

INVOICE TO	SHIP TO
DEPARTMENT OF ADMINISTRATION OFFICE OF TECHNOLOGY 1900 KANAWHA BLVD E, BLDG 5 10TH FLOOR CHARLESTON WV25305 US	WV OFFICE OF TECHNOLOGY BLDG 5, 10TH FLOOR 1900 KANAWHA BLVD E CHARLESTON WV 25304 US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
2	Maintenance and Support				# 34,941.64

Comm Code	Manufacturer	Specification	Model #
81112201			

Extended Description :
Initial Year Maintenance and Support

SCHEDULE OF EVENTS

Line	Event	Event Date
1	Technical Questions Due	2016-08-17

ISC1700000002	Document Phase Final	Document Description SAN Storage Systems (OT1705)	Page 3 of 3
---------------	--------------------------------	---	-----------------------


ADDITIONAL TERMS AND CONDITIONS

See attached document(s) for additional Terms and Conditions

DESIGNATED CONTACT: Vendor appoints the individual identified in this Section as the Contract Administrator and the initial point of contact for matters relating to this Contract.

Ashley Howard , Account Manager
(Name, Title)
Ashley Howard , Account Manager
(Printed Name and Title)
221 E 4th St , Cincinnati, OH
(Address)
513 841 6298 1866 549 0474
(Phone Number) / (Fax Number)
ashley.howard @ cbts.net
(email address)

CERTIFICATION AND SIGNATURE: By signing below, or submitting documentation through wvOASIS, I certify that I have reviewed this Solicitation in its entirety; that I understand the requirements, terms and conditions, and other information contained herein; that this bid, offer or proposal constitutes an offer to the State that cannot be unilaterally withdrawn; that the product or service proposed meets the mandatory requirements contained in the Solicitation for that product or service, unless otherwise stated herein; that the Vendor accepts the terms and conditions contained in the Solicitation, unless otherwise stated herein; that I am submitting this bid, offer or proposal for review and consideration; that I am authorized by the vendor to execute and submit this bid, offer, or proposal, or any documents related thereto on vendor's behalf; that I am authorized to bind the vendor in a contractual relationship; and that to the best of my knowledge, the vendor has properly registered with any State agency that may require registration.

CBTS
(Company)

(Authorized Signature) (Representative Name, Title)
DAN MAZZA - Director of Sales
(Printed Name and Title of Authorized Representative)
8/23/16
(Date)
513-397-7259
(Phone Number) (Fax Number)



Hewlett Packard
Enterprise

HPE 3PAR StoreServ Storage best practices guide

A reference and best practices guide for HPE 3PAR StoreServ Storage

Contents

Typographical conventions.....	4
Advisories.....	4
Introduction.....	5
Audience.....	5
Overview.....	6
New hardware platforms for the StoreServ platform.....	10
What's new in HPE 3PAR OS version 3.2.2 (including MU2)?.....	11
Getting started with HPE 3PAR StoreServ Storage.....	12
FC hosts zoning.....	13
Hosts and host sets.....	14
Provisioning block storage from an HPE 3PAR StoreServ.....	15
Host-based volume managers.....	15
Adaptive Flash Cache.....	15
Considerations for provisioning virtual volumes.....	18
Provisioning file storage from an HPE 3PAR StoreServ.....	19
High availability.....	23
Persistent Ports.....	23
Priority Optimization.....	24
Virtual volumes.....	25
Virtual LUNs (exports) and volume sets.....	26
Remote Copy.....	27
Streaming Asynchronous Replication.....	27
Adaptive Optimization.....	28
Security.....	30
Naming conventions.....	31
Naming convention examples.....	31
Naming conventions with File Persona Objects.....	33
External System Reporter; now EOL (End of Life).....	33
System reporter in the SSMC.....	33
Object selection.....	33
Ongoing management and growth.....	34
Storage Analytics on the Web.....	34
Autonomic rebalance.....	34

Appendix A. Supported host personas.....	35
Appendix B.....	35
Block persona scalability limits.....	35
File persona scalability limits.....	36
Appendix C. File Persona compatibility with HPE 3PAR block features.....	36
Appendix D. Using the StoreServ management console (SSMC) to administer File Persona.....	37
Appendix E. Using the StoreServ management console (SSMC) to manage snapshots.....	43
Summary.....	43

Typographical conventions

This guide uses the following typographical conventions:

Table 1. Typographical conventions

TYPEFACE	MEANING	EXAMPLE
ABCDabcd	Used for dialog elements such as titles, button labels, and other screen elements	When prompted, click finish to complete the installation
ABCDabcd	Used for user input, filenames, commands, paths, and screen output	Start the \OS\Windows\setup.exe
<ABCDabcd>	Used for variables in user input, filenames, paths, and screen output	To add a user to a domain issue the set user—add domain <domainname>:role <username> command
Best practice	Used to highlight best practices for a particular topic section	Best practice: use RAID 6

Advisories

To avoid injury to people or damage to data and equipment, be sure to observe the cautions and warnings in this guide.

Always be careful when handling any electrical equipment.

WARNING!

Warnings alert you to actions that can cause injury to people or irreversible damage to data or the OS.

CAUTION!

Cautions alert you to actions that can cause damage to equipment, software, or data.

Note

Notes are reminders, tips, or suggestions that supplement the procedures included in this guide.

Introduction

Audience

This guide is for system and storage administrators of all levels. Anyone who plans storage policies, configures storage resources, or monitors the storage usage of HPE 3PAR StoreServ Storage should read this guide.

User interfaces

Previously two user interfaces were available for the administration of HPE 3PAR StoreServ; the HPE 3PAR OS CLI software and the HPE 3PAR IMC (InServ Management Console) software. With version 3.2.2 HPE 3PAR introduced the [HPE 3PAR StoreServ Management Console \(SSMC\)](#); all operations for File Persona are accomplished using the CLI or the SSMC.

Units of measure

All units of storage (capacity) are calculated base 2 (x 1,024). Therefore:

- 1 KiB = 1,024 bytes
- 1 MiB = 2²⁰ bytes = 1,048,576 bytes
- 1 GiB = 2³⁰ bytes = 1,024 MB = 1,073,741,824 bytes
- 1 TiB = 2⁴⁰ bytes = 1,024 GB = 1,099,511,627,776 bytes

All units of performance (speed) are calculated base 10 (x 1000). Therefore:

- 1 KB = 1,000 bytes
- 1 MB = 10⁶ bytes = 1,000,000 bytes
- 1 GB = 10⁹ bytes = 1000 MB = 1,000,000,000 bytes
- 1 TB = 10¹² bytes = 1000 GB = 1,000,000,000,000 bytes

Table 2. Related documentation

Complete description of CLI commands	HPE 3PAR CLI Reference
Overview and explanation of HPE 3PAR technology	HPE 3PAR concept guide
Using the Management Console to configure and administer the system	HPE 3PAR Management Console online help
Using the CLI to configure and administer the system	HPE 3PAR CLI Administrator's manual
Identifying storage system components and detailed alert information	HPE Guided Troubleshooting
Using HPE 3PAR Remote Copy software	HPE 3PAR Remote Copy software user guide
Using the HPE 3PAR CIM	HPE 3PAR CIM API Programming Reference
Using the HPE 3PAR 3PARInfo tool	HPE 3PAR 3PARInfo user guide

For identifying storage system configuration specifications and compatibility information, go to the Single Point of Connectivity Knowledge (SPOCK) website at hpe.com/storage/spock.

Overview

HPE 3PAR StoreServ block storage concepts and terminology

The HPE 3PAR StoreServ array is comprised of the following logical data layers:

- Physical disks (PDs)
- Chunklets
- Logical disks (LDs)
- Common provisioning groups (CPGs)
- Virtual volumes (VVs)

The relationship between system data layers is illustrated in figure 1. Each layer is created from elements of the layer above. Chunklets are drawn from physical disks. Logical disks are created from groups of chunklets. Common provisioning groups are groups of logical disks. And virtual volumes use storage space provided by CPGs. The virtual volumes are exported to hosts and are the only data layer visible to hosts.




Figure 1. HPE 3PAR StoreServ system data layers

Physical disks

A physical disk is a hard drive (spinning media or Solid State Drive) located in an HPE 3PAR StoreServ drive enclosure.

Chunklets

Physical disks are divided into chunklets. Each chunklet occupies physically contiguous space on a FC or NL disk. On all current HPE 3PAR StoreServs all chunklets are 1 GB. Chunklets are automatically created by the HPE 3PAR OS, and they are used to create logical disks. A chunklet is assigned to only one logical disk.

Logical disks

A logical disk is a collection of chunklets arranged as rows of RAID sets. Each RAID set is made up of chunklets from different physical disks. Logical disks are pooled together in common provisioning groups, which allocate space to virtual volumes.

The underlying logical disks are automatically created by the HPE 3PAR OS when you create VVs. The RAID type, space allocation, growth increments, and other logical disk parameters are specified when you create a CPG or can be modified later. The HPE 3PAR StoreServ supports the following RAID types:

- RAID 1/RAID 1+0
- RAID 5/RAID 50 (Must be enabled from the CLI for NL drives)
- RAID Multi-Parity (MP) or RAID 6
- RAID 0 (must be enabled from the CLI and provides no data protection from failed drives)

Cage

Cage is a legacy HPE 3PAR term and is interchangeable with “Drive Enclosure”, “Enclosure”, and “Drive Shelf”.

Virtual Copy

Virtual Copy is a legacy HPE 3PAR term and is interchangeable with “Snapshot”.

CPGs (Common provisioning groups)

A CPG is a template for the creation of logical disks that allocate space to virtual volumes on demand. A CPG allows up to 65,536 virtual volumes to share a CPG's assigned resources. You can create Fully Provisioned Virtual Volumes (FPVVs), Thinly Deduped Virtual Volumes (TDVVs), and Thinly Provisioned Virtual Volumes (TPVVs) that draw space from a CPG's logical disks. It is important to note that if no volumes of any type have been created in a CPG, it consumes no space.

VVs (Virtual volumes)

VVs draw their resources from the LDs in CPGs and are exported as LUNs (Logical Unit Numbers) to hosts. Virtual volumes are the only data layer visible to the hosts. You can create clones (previously known as full copies) or snapshots (previously known as virtual copies) of virtual volumes. Clones remain available if the original base volume becomes unavailable. VVs can be created using the CPGs created at installation time or user defined CPGs.

Exporting virtual volumes

For a host to see a VV, the volume must be exported as a LUN. Volumes are exported by creating VV-LUN pairings (VLUNs) on the system. When you create VLUNs, the system produces both VLUN templates that establish export rules, and active VLUNs that the host sees as LUNs as attached disk devices. A VLUN will be created for each path available to the host for each VV exported.

FPVVs (Fully Provisioned Virtual Volumes)

A FPVV is a volume that uses logical disks that belong to a CPG. Unlike TPVVs or TDVVs, FPVVs have a set amount of user space that is allocated for user data. The fully provisioned volume size is allocated and consumed at the time of provisioning, size limits range from 256 MB to 16 TB. The volume size can be increased at any time (provided free space is available) up to the maximum 16 TiB size without any downtime however, the VV size cannot be decreased below the initial allocation.

Note

In previous versions of the HPE 3PAR OS there was a provisioning type termed CPVV (Copy Provisioned Virtual Volume) which simply meant that the provisioned VV had associated snapshot space assigned. As of HPE 3PAR OS version 3.2.2 all volumes created are associated with snapshot space in the same CPG. Using the additional menu options during VV creation in the SSMC or using the `-snp_cpg` option of the `createvv` CLI command, a different CPG can be chosen for snapshot space. Also snapshots are reservationless, if no snapshots are generated, no space is consumed.

TPVVs (Thinly Provisioned Virtual Volumes)

A TPVV is a volume that uses logical disks that belong to a CPG. TPVVs or TDVVs associated with the same CPG draw space from those CPGs' LDs as needed, allocating space on demand in 16 KiB increments for each TPVV. As the volumes that draw space from the CPGs' LDs require additional storage the HPE 3PAR OS automatically creates additional logical disks or expands the size of existing LDs and adds them to the CPG until the CPG reaches a user-defined growth limit, if one has been set, or the system runs out of space.

TDVVs (Thinly Deduped Virtual Volumes)

In addition to the features and functionality of TPVVs, TDVVs go through an additional process before allocating space on a disk. All data writes with a block size 16 KiB or greater have a 32 bit hash (CRC—Cyclic Redundancy Check) generated and the resulting value is compared to a hash lookup table to determine if the data is redundant. If the data is redundant there is only an entry added to the destination volumes lookup table, otherwise it is written to disk. For more information on this, refer to the HPE 3PAR Thin Technologies white paper located [here](#).

Clones (Previously known as Full Copies)

A clone duplicates all the data from a base volume to a destination volume. The base volume is the original volume that is copied to the destination volume. The clone on the destination volume remains available if the original base volume becomes unavailable.

A clone requires the destination volume have usable capacity equal to or greater than the usable capacity of the base volume being cloned. As of HPE 3PAR OS 3.2.1 the clone can be exported immediately after creation, while the data copy continues in the background.

Snapshots (Previously known and licensed as Virtual Copy)

Unlike a clone, which is a block for block duplicate of an entire volume, snapshots preserve a bitmap of a VV at a particular point in time. Updates to VVs are written to SD (Snap Data) space and the bitmap (Snap Admin space) of the VV.

Snapshots for FPVVs, TPVVs, clones and other snapshots are created using copy-on-write techniques available only with the HPE 3PAR snapshot software license (aka Virtual Copy) snapshots for TDVVs are created using ROW (Redirect On Write). Hundreds of snapshots of each virtual volume can be created assuming that there is sufficient storage space available. It is worth noting that snapshots do not consume any space unless data on the base volume has been updated and the original data copied to the SD (snap data) space. Changed data is copied only once regardless of the number of snapshots taken. Snapshots are particularly useful for test/dev environments as they can be created in seconds and exported while not effecting production data. Also, testers and developers can be granted the ability to create snapshots while not having any other administrative privileges, lowering administration requirements (see the [HPE 3PAR CLI](#) guide for correct usage for granting the updatevv privilege). Snapshots can now be updated without the requirement to un-export and re-export the VV. For more information on snapshot/Virtual Copy technologies refer to the white paper [here](#).

Note

Creating snapshots requires an HPE 3PAR Snapshot (aka Virtual Copy) software license.

HPE 3PAR File Persona storage concepts and terminology




Figure 2. HPE 3PAR File Persona logical view

File Provisioning Group

A File Provisioning Group (FPG) is an instance of the HPE intellectual property Adaptive File System. It controls how files are stored and retrieved. Each FPG is transparently constructed from one or multiple virtual volumes (VVs) and is the unit for replication and disaster recovery for File Persona Software Suite. There are up to 16 FPGs supported on a node pair.

Virtual File Server

A Virtual File Server (VFS) is conceptually like a server; as such it presents virtual IP addresses to clients, participates in User Authentication Services and can have properties for such characteristics as user/group Quota Management and Antivirus policies.

File Stores

File Stores are the slice of a VFS and FPG where snapshots are taken, capacity Quota Management can be performed, and Antivirus Scan Services policies customized.

File Shares

File Shares are what provide data access to clients via SMB, NFS, and the Object Access API, subject to the share permissions applied to them.

New hardware platforms for the StoreServ platform

Table 3. 8000 series

	8200	8400	8440	8450
Base VV	32k	32k	64k	64k
Total VV	64k	64k	128k	128k
Max VLUN	128k	128k	256k	256k
FC 16 GB/s Port initiators	256	256	256	256
iSCSI/FCoE Port initiators	256	256	256	256
FC system initiators	2,048	4,096	4,096	4,096
iSCSI/FCoE/System initiators	1,024	2,048	2,048	2,048
Remote Copy Max VVs—Sync	800	800–2 nodes 2400–4 nodes	800–2 nodes 2400–4+ nodes	800–2 nodes 2400–4+ nodes
Remote Copy Max VVs—Asynchronous Periodic	2,400	2,400–2 nodes 6,000–4 nodes	2,400–2 nodes 6,000–4+ nodes	2,400–2 nodes 6,000–4 nodes
Peer Persistence Max VVs	600	600	600	600

Table 4. 20000 series

	20450	20800	20840	20850
Base VV	64k	64k	64k	64k
Total VV	128k	128k	128k	128k
Max VLUN	256k	256k	256k	256k
FC 16 GB/s Port initiators	256	256	256	256
iSCSI/FCoE Port initiators	256	256	256	256
FC system initiators	8,192	8,192	8,192	8,192
iSCSI/FCoE/System initiators	8,192	8,192	8,192	8,192
Remote Copy Max VVs—Sync	800–2 nodes 2400–4+ nodes	800–2 nodes 2400–4+ nodes	800–2 nodes 2400–4+ nodes	800–2 nodes 2400–4+ nodes
Remote Copy Max VVs—Asynchronous Periodic	2,400–2 nodes 6,000–4 nodes	800–2 nodes 2400–4+ nodes	800–2 nodes 2400–4+ nodes	2,400–2 nodes 6,000–4 nodes
Peer Persistence Max VVs	600	600	600	600

What's new in HPE 3PAR OS version 3.2.2 (including MU2)?

Remote Copy enhancements

- Remote copy now supports Asynchronous Streaming; in the industry this is also known as “True Async”.

Asynchronous Streaming allows a source array to acknowledge a host write before the destination array has acknowledged the write back to the source array. Asynchronous Streaming Remote Copy is perfect for environments where very small RPOs are required and environments where synchronous replication is desired but the replication link latencies exceed 10 ms (5 ms for FCIP) where using synchronous replication would result in unacceptable latency.

Note

Streaming Asynchronous Remote Copy is supported on FC and FCIP transports only.

Tuning and Performance enhancements

Improvements in the tunesys (data rebalancing) behavior.

- Tunesys now starts automatically after initiation of the admithw command on 7200 and 8400 platforms, although the balancing algorithm takes care not to interfere with performance; care should be exercised not to execute the command during peak hours.
- More granular reporting of data moved, now expressed in terms of GiB moved and GiB to be moved.
- Tunesys raw speed has been increased.
- Adaptive Optimization (AO) now support “Premium” mode to keep SSD drives filled as much as possible.
- Data deduplication performance enhancements.
- Updating snapshots (updatevv) no longer requires un-export and re-export of virtual volumes.

Security

- Increased number of LDAP servers from one to unlimited.
- Support for subdomains in the same root domain.
- Support for load balancers in front of LDAP (Active Directory) servers.

Resiliency

- Persistent Checksum—Support for end-to-end data integrity on the 8000 and 20000 series platforms.

Efficiency

- New flash cache creation will be done with RAID 0 to increase available space for AFC.

Reporting with System Reporter in SSMC

SSMC comes with a fresh way of reporting capacity and performance data. Using Reports under the System Reporter main menu, users can launch the integrated reporting tool within SSMC. Reports enable users to run historical and real time reports. This new approach offers great advantages over the previous approaches of querying historical data.

System Reporter in SSMC provides the following enhanced features:

- Convenient access to configuration options for selecting systems to include for reporting, specifying sampling parameters, scheduling reports, and generating alerts.
- Extensive selection of reports for obtaining performance and storage utilization statistics on selected objects (i.e., hosts, ports, nodes, physical disks, virtual volumes, etc.).
- Quick access to predefined reports that contain useful statistics for most common types of installations.
- Customization of reports using the standard Web interface that provides specifically selected and formatted reports for specified systems.
- Options for choosing the time and duration for the collection of reporting statistics that can be initiated at a specific time, collected over a period, and/or compared between ranges of periods.

- Capability to isolate and zoom in and out of time periods.
- Performance alerts that can be configured via threshold alerts, once criteria for alert is met, visual notification of the alert is displayed on the SSMC dashboard.
- Ability to edit generated reports, change object definitions.
- Report customization allows hundreds of different report generation.
- The size of the system reporter database is now tunable up to 1 TiB in size allowing for very long data retention.
- Data from the system reporter database can now be exported.

General

- VLAN tagging is now supported for iSCSI connectivity on the 8000 and 20000 platforms.
- IPv6 is now supported for iSCSI connectivity on the 8000 and 20000 series arrays.
- Deduplication can now be enabled and disabled globally using the “setsys” command.

Storage Federation

- Online Import now supports 16 Gb and import priorities.
- Online migration now supports up to four source arrays to one destination array (unidirectional).

Getting started with HPE 3PAR StoreServ Storage

Best practice: In addition to following the recommendations in the HPE 3PAR physical planning guide, it is important to keep the number and type of physical disks as well as the number of drive enclosures as evenly distributed as possible behind node pairs to facilitate maximum performance and load distribution.

Best practice: Tuning CPGs/VVs as a result of a hardware upgrade or changing requirements, use the default tunesys concurrent task level to two in order to limit processing overhead in a production environment.

Setting up your HPE 3PAR StoreServ system ports

Port locations and nomenclature

The HPE 3PAR CLI and SSMC display the controller node, FC, iSCSI, 1 Gigabit and 10 Gigabit Ethernet port locations in the following format: <Node>:<Slot>:<Port>. For example: 2:4:1.

- Node: Valid node numbers are 0–7 depending on the number of nodes installed in your system; when viewing a system from the rear of a cabinet.
- In 7000 and 8000 series arrays nodes are numbered 0–3 from bottom left to the top right when facing the service side (rear) of the nodes.
- In 10000 and 20000 series arrays nodes are numbered 0–7 from bottom left to top right, when facing the service side (rear) of the nodes.
- The 7000 and 8000 series arrays have a single onboard slot in each node, numbered starting at 0.
- The 10000 and 20000 arrays' slots are numbered left to right/top to bottom starting 0 from left to right, bottom to top in a node in the lower enclosure. In the upper enclosure, slots are numbered 0–9 from left to right, top to bottom.
- Port: Valid node port numbers are 1–4 for all add in adapters; counting from the bottom up.
- 7000 and 8000 ports are horizontal and labeled beginning with 1 on the HBA or iSCSI adapter.
- 10000 and 20000 ports are numbered from bottom to top in a node in the lower enclosure. In the upper enclosure, ports are numbered from top to bottom.

Front-end port cabling

Best practice: Each HPE 3PAR StoreServ controller node should be connected to two fabrics. This is to protect against fabric failures.

Best practice: Ports of the same pair of nodes with the same ID should be connected to the same fabric. Example:

- 0:2:3 and 1:2:3 on fabric 1
- 0:2:4 and 1:2:4 on fabric 2

Best practice: Connect Odd ports fabric 1 and even ports to fabric 2 and so forth. Example with a 4-node 7400 with eight host ports:

FABRIC 1	FABRIC 2
0:2:3, 1:2:3, 2:2:3, 3:2:3	0:2:4, 1:2:4, 2:2:4, 3:2:4

Example with a 4-node HPE 3PAR 10400 with 32 host ports:

FABRIC 1	FABRIC 2
0:2:1, 0:2:3, 0:5:1, 0:5:3, 1:2:1, 1:2:3, 1:5:1, 1:5:3, 2:2:1, 2:2:3, 2:5:1, 2:5:3, 3:2:1, 3:2:3, 3:5:1, 3:5:3	0:2:2, 0:2:4, 0:5:2, 0:5:4, 1:2:2, 1:2:4, 1:5:2, 1:5:4, 2:2:2, 2:2:4, 2:5:2, 2:5:4, 3:2:2, 3:2:4, 3:5:2, 3:5:4

FC hosts zoning

Best practice: Use at least two separate SAN switches to enable availability in the event of a switch failure.

Best practice: One initiator to multiple targets per zone (zoning by HBA). This zoning configuration is recommended for the HPE 3PAR StoreServ Storage. Zoning by HBA is required for coexistence with other HPE Storage arrays such as the HPE EVA.

Best practice: Zoning should be done using Worldwide Port Names (WWPN, the WWN of each individual port on HPE 3PAR StoreServ). Port Persistence is not compatible with DID zoning.

Best practice: Hosts should be mirrored to node pairs. For example: zoned to nodes 0 and 1, or nodes 2 and 3. Hosts should **not** be zoned to non-mirrored nodes, such as 0 and 3.

Best practice: Non-hypervisor host

- A single non-hypervisor host port should be zoned with a minimum of two ports from the two nodes of the same pair. In addition, the ports from a host’s zoning should be mirrored across nodes. In the case of hosts attaching with multiple host HBA ports attached to dual switches, each port should be zoned to at least the two-mirrored nodes.
- Non-hypervisor hosts do not need to be connected to all nodes because of the way the volumes are spread on all the nodes.

Best practice: Hypervisor host

- A single hypervisor host should be zoned to a maximum of 4 nodes in two separate node pairs (for example 0&1 and 2&3) to maximize bandwidth.
- Each Hypervisor HBA port should be zoned to a minimum of two ports from each node pair on two node systems; and at least one or more port(s) from at least four nodes on 4, 6, and 8 node systems in order to maximize throughput across multiple node busses.

Table 5. Examples of Valid Zoning

Host Type	Host HBA Port	StoreServ Ports
Non-hypervisor, single HBA single port	HBA1:port 1	0:1:1, 1:1:1
Non-hypervisor, single HBA, two HBA ports to separate switches	HBA1:port 1	0:1:1, 1:1:1
	HBA1:port 2	0:1:2, 1:1:2
Hypervisor, two HBAs, two HBA ports each, connected to two separate switches. Four node StoreServ.	HBA1:port 1	0:1:1, 1:1:1
	HBA1:port 2	0:1:2, 1:1:2
	HBA2:port 1	2:1:1, 3:2:1
	HBA2:port 2	2:2:1, 3:2:2
	HBA2:port 2	2:2:1, 3:2:2

Best practice: Each HPE 3PAR StoreServ system has a maximum number of initiators supported that depends on the model and configuration. In regard to this maximum, 1 initiator = 1 path from a host.

A single HBA zoned with two FC ports will be counted as two initiators.

A host with two HBAs, each zoned with two ports, will count as four initiators.

Best practice: No more than 256 connections (128 if using Persistent Ports) are supported per front-end/host port.

Hosts and host sets

Best practice: When creating hosts, follow the implementation guide for each platform.

Selecting the correct host persona (specifying host operating system) for each host is important. Implementation guides are available for download at the following address: hp.com/go/storage.

Each physical server should have a different host defined, containing the WWNs or iQN for this host.

Best practice for creating a new host (Windows®, SPARC Solaris, VMware®, and Red Hat® Linux®):

1. Install the Host Explorer software in the host if available for the host platform (download here).
2. Zone in all the ports according to the zoning best practices.
3. From the host CLI, execute `tpdhostagent-start`; then `tpdhostagent-push`.
4. This will automatically create the host on the HPE 3PAR StoreServ Storage system.

Best practice for creating a new host manually:

1. Zone in host ports to HPE 3PAR StoreServ using the zoning best practices, one host at a time.
2. For each host, select the host and then create the new host.
3. In the WWN selection screen, select the WWNs associated with the new host.
4. Zone in a single host and then create the host on the HPE 3PAR StoreServ to reduce the possibility of assigning incorrect WWNs to a host. Repeat until all hosts are zoned in.




Figure 3. Creating a host in the Management Console

Best practice: For clusters, create a host set containing all the hosts used by the cluster. This will allow export of shared VVs to all hosts of the cluster in a single operation.

Provisioning block storage from an HPE 3PAR StoreServ

Note

This also includes provisioning block storage for File Persona use.

Host-based volume managers

Use of external volume managers is often unnecessary because of the advanced volume layout algorithms already used by the HPE 3PAR OS. One notable exception to this best practice is large block sequential workloads such as Statistical Analysis Software (SAS) and video streaming as these workloads take advantage of the read ahead algorithm used by the HPE 3PAR OS. The algorithm will instantiate up to 5 read ahead threads per VV, effectively preloading data into data cache when sequential reads are detected.

Adaptive Flash Cache

HPE 3PAR Adaptive Flash Cache is included as part of the HPE 3PAR Operating System Suite version 3.2.1 and later, and is supported in all HPE 3PAR StoreServ Storage arrays that have a blend of solid-state drives (SSDs) and hard disk drives (HDDs).

Note

AFC is not supported on the non-AFC 480 GB SSD drives (**E7Y55A** and **E7Y56A**).

Benefits

- Reduced latency for random read-intensive workloads.
- Responds dynamically providing smart and adaptive data placement based on application and workload demands.
- Enables HPE 3PAR Adaptive Flash Cache across the entire system or select particular workloads to accelerate.

Requirements

- HPE 3PAR OS version 3.2.2 (license is bundled)

Four SSD drives in 7000/8000 series or eight SSD drives in the 10000/20000 series.

Table 6. Supported configurations and maximum flash cache. (7000 and 10000 Series arrays)

OBJECT	7200	7400	7440C	10400 (OLD NODE)	10400 NEW NODE	10800 AFC
AFC System	768 GB	1.5 TB	3 TB	3 TB	4 TB	8 TB
AFC Node Pair	768 GB	768 GB	1.5 TB	1.5 TB	2 TB	2 TB
Min SSD	4	4	4	8	8	8

Table 7. Supported configurations and maximum flash cache. (8000 and 20000 series systems)

OBJECT	8200	8400	8440	20800
AFC System	768 GB	1.5 TB	3 TB	8 TB
AFC Node Pair	768 GB	768 GB	1.5 TB	2 TB
Min SSD	4	4	4	8

Best practice: If SSD drives are not available, use the CLI command `createflashcache sim`. This will simulate how much data could possibly move into the flashcache tier and the possible increase in performance.

Best practice: On systems with four or more SSD drives per node pair available, enable flash cache system wide using the following commands.

`createflashcache <maximum noted in table above>`. If there is insufficient space available for the maximum amount the command will return the maximum amount available.

`setflashcache enable sys:all` (alternative to enabling flash cache for the entire system, VVsets can also be selected.)

For more information, refer to the HPE 3PAR Adaptive Flash Cache white paper [here](#).

Common provisioning groups

Notes

- CPGs primarily are templates for the creation of LDs but have other defining characteristics such as capacity limits.
- If there are no volumes created “in” a CPG, it will consume no space.
- CPGs define:
 - The RAID level for the LDs to be created.
 - Availability level (HA CAGE, HA PORT, or HA MAG).
- Step size as well as other characteristics such as drive geometry.
- CPGs will only be created across drives of the same type (SSD, FC, or NL) and will include drives of different rpms unless otherwise specified using a disk filter.

Best practice: CPGs using Fast Class (FC/SAS) or NL should use RAID 6, for SSD CPGs use RAID 5.

Best practice: When creating CPGs, accept defaults according to performance/capacity requirements.

Exceptions include not enough drive enclosures to achieve HA cage, requiring higher capacity utilization with the use of RAID 5 or providing protection against double disk failure with RAID 6.

Best practice: The number of CPGs should be kept to a minimum. Refer to [Appendix B](#) for pertinent limits.

Best practice: There are cases in which having more CPGs than the minimum will be required.

- Using thin provisioned VVs while using adaptive optimization.
- When using HPE 3PAR Virtual Domain software: a given CPG can only be in one domain.
- When using Adaptive Optimization software: a given CPG can only be in one Adaptive Optimization policy.
- When capacity reporting is required per customer or application, per customer/application CPGs ease capacity reporting.
- When snapshots are heavily used and the snapshot data is kept in a different tier than the source data, use a CPG that matches the production CPG performance characteristics as close as possible to maximize performance.

Best practice: Do not set “growth limits” on CPGs.

If a warning threshold is required, set a “growth warning” (warning in terms of capacity), not an “allocation warning” (warning in percentage).

Best practice: Avoid creation of RAID 0 CPGs as RAID 0 offers no protection from data loss from drive failures (RAID 0 is disabled by default; consult the HPE 3PAR StoreServ CLI manual for instructions for enabling R0 as well as R5 on NL).

Solid-state drive CPGs

Best practice: Solid-state drive (SSD) CPGs should be of the RAID 5 type with a “set size” of 3+1 by default. This will bring superior performance/capacity ratio. If maximum performance is required, use RAID 1.

Best practice: The growth increment should be set to the minimum value, which is 8 GB per node pair.

On two-node systems, set the value to 8 GB, on four-node systems to 16 GB, on six-node systems to 24 GB, and on eight-node systems to 32 GB.

In order to set the CPG growth increment to a lower value than the default, the “show advanced option” box must be checked.

Best practice: Availability should be left to “cage level” availability (the default option) if the system’s configuration allows for it. If not, it should be set to “magazine level” availability. This can be changed using the “advanced options” checkbox in the SSMC.

Other advanced settings such as “preferred chunklets” and “step size” should not be changed from their default values. Also, avoid using disk filters.

Fast class (FC/SAS) CPGs

Best practice: FC CPGs should be of the RAID 6 type by default. This will bring the highest availability for modern high capacity drives. The “set size” (data to parity ratio) can be changed from the default value of 6+2 if the system configuration supports it. If usable capacity is the primary concern use a wider stripe.

Best practice: For applications that have a very high write ratio (over 50 percent of the access rate), create a CPG using RAID 1 if performance (as opposed to usable capacity) is the primary concern.

Best practice: The growth increment should be left to the default value (32 GB per node pair).

Best practice: Availability should be left to “cage level” (the default option) if the system’s configuration allows for it.

If not, it should be set to “magazine level” availability. This can be changed using the “advanced options” checkbox of the Management Console.

Best practice: Leave other advanced setting such as “preferred chunklets” and “step size” to the defaults.

NL CPGs

Best practice: NL CPGs should be RAID 6, which is the default.

The “set size” (data to parity ratio) can be changed from the default value of 8 (6+2) if the system configuration supports it. RAID 5 is not recommended with NL disks.

Best practice: The growth increment should be left to the default value (32 GB per node pair).

Best practice: Availability should be left to “cage level” (the default option) if the system’s configuration allows for it.

Note

HA cage requires that the number of cages behind **each node pair** must equal to or larger than the set size. For example RAID 6 with a set size of 8 (6+2) requires 8 drive enclosures (cages) or more behind each node pair.

Select “magazine level” availability if HA cage is not possible. This can be changed using the “additional setting” checkbox of the SSMC during VV creation.

Other advanced settings such as “preferred chunklets” and “step size” should be left at default values.

Considerations for provisioning virtual volumes**Thinly Deduplicated Virtual Volumes (TDVVs)**

- TDVVs can only reside on SSD storage. Any system with an SSD tier can take advantage of thin deduplication. The option to provision TDVVs is not available in the HPE 3PAR SSMC unless an SSD CPG has been selected.
- The granularity of deduplication is 16 KiB and therefore the efficiency is greatest when the I/Os are aligned to this granularity. For hosts that use file systems with tunable allocation units consider setting the allocation unit to a multiple of 16 KiB.
- Deduplication is performed on the data contained within the virtual volumes of a CPG. For maximum deduplication store data with duplicate affinity on virtual volumes within the same CPG.
- Thin deduplication is ideal for data that has a high level of redundancy. Data that has been previously deduplicated, compressed, or encrypted are not good candidates for deduplication and should be stored on thinly provisioned volumes.
- AO does not support TDVVs.
- When using an HPE 3PAR array as external storage to a third-party array deduplication may not function optimally.

Best practice: Use TDVVs when there is a high level of redundant data, and the primary goal is capacity efficiency.

For more information, refer to the HPE 3PAR Thin Technologies white paper located [here](#).

Thin Provisioned Virtual Volumes (TPVVs)

The use of thin provisioning has minimal performance impact and has the significant operational benefit of reducing storage consumption. However, there are certain workloads and applications thin provisioning may not be of benefit such as:

- Applications that write continuously to new space; an example of this is Oracle redo log files will not benefit from thin provisioning as the space will be consumed until the volume is full.
- Environments that require host encrypted volumes-writing zeroes to a host encrypted volume on a newly created HPE 3PAR StoreServ thin-provisioned volume will cause space to be allocated on the TPVV because the encryption alters the content of the blocks. Applying encryption to thin-provisioned volumes that already contain data or rekeying them also inflates the zero blocks, making the volume consume space as if it was fully provisioned. Attempting to re-thin the volume by writing zeroes to allocated but unused space will not decrease the space utilization. As a result, host encryption and thin provisioning do not cooperate.
- Environments that require SAN encrypted volumes-for example, host-based encryption, encryption by a device in the data path (e.g., SAN switch) also alters the data stream so that blocks of zeroes written by the host are not passed onto the storage. A notable exception is Brocade SAN switches. With the introduction of Fabric OS 7.1.0 the Fabric OS encryption switch can automatically detect if a volume is a thin-provisioned LUN. If a LUN is detected as thin-provisioned, the first-time encryption and rekey are done on the allocated blocks only. This thin provisioned LUN support requires no action by the user.
- Copy on Write file systems that write new blocks rather than overwrite existing data are not suitable for thin provisioning as every write will allocate new storage until the volumes fully allocated. An example of a copy on write file system is Oracle Solaris ZFS.

Best practice: Use TPVVs as a general practice, with exceptions as noted for TDVVs and FPVVs.

For more information, refer to the HPE 3PAR Thin Technologies white paper [here](#).

Note

Copy Provisioned Virtual Volumes (CPVVs), an HPE 3PAR deprecated term, are simply VVs that are associated with a snapshot CPG. As of HPE 3PAR OS 3.2.2 all volumes created are associated with snapshot space, in the same CPG, by default. A different snapshot CPG can be selected when creating a VV by selecting the additional settings option in the SSMC or using the `-snp_cpg` of the `createvv` CLI command.

Fully Provisioned Virtual Volumes (FPVVs)

FPVVs allocate all space and create all LDs at initial provisioning.

- An FPVV provides the highest performance of the three provisioning types.
- With a dynamic optimization license FPVVs can be converted to TPVVs, and further the TPVV can then be converted to TDVVs provided there is available SSD space.
- Workloads that continuously write data in new extents instead of overwriting data or that perform heavy sequential write workloads to RAID 5 or RAID 6 will benefit most from FPVVs.

Best practice: Use FPVVs when the highest performance is the priority.

Provisioning file storage from an HPE 3PAR StoreServ

HPE 3PAR OS version 3.2.1 MU2 introduced the HPE 3PAR File Persona Software Suite comprised of rich file protocols including SMB 3.0, 2.1, 2.0, and 1.0 plus NFSv4.0, and v3.0 to support a broad range of client operating systems. It also includes the Object Access API that enables programmatic data access via a REST API for cloud applications from virtually any device anywhere.

The HPE 3PAR File Persona extends the spectrum of primary storage workloads natively addressed by the HPE 3PAR StoreServ platform from virtualization, databases, and applications via the Block Persona to also include client workloads such as home directory consolidation, group/department shares, and corporate shares via the File Persona—all with truly converged controllers, agile capacity, and unified management.

Note

Make sure that your HPE 3PAR StoreServ Storage is configured and provisioned for physical disk storage and ready to provision storage for File Persona.

HPE 3PAR File Persona has the following expanded and recommended use cases:

- Home directory consolidation
- Group/department and corporate shares
- Custom cloud applications using Object Access API

HPE 3PAR File Persona managed objects

The HPE 3PAR File Persona Software Suite is comprised of the following managed objects:

- File provisioning groups (FPGs)
- Virtual file servers (VFSs)
- File Stores

Best practice: Limit the number of File Share objects created. For example, home directory shares at the department level instead of at the individual user level. This reduces the amount of data cache required while also reducing networking traffic.

SSMC (StoreServ Management Console)

HPE 3PAR offers the new, streamlined SSMC and HPE 3PAR CLI for the management of converged block, file, and object access on HPE 3PAR StoreServ systems. See [Appendix D](#) for basic SSMC navigation for managing the File Persona feature.

Note

Make sure to apply the license for the HPE 3PAR File Persona Software Suite on the system in order to enable the file and object access on the HPE 3PAR StoreServ 7000c, 8000, and 20000 series converged controllers. File Persona will show up in SSMC only after enabling the license on the system. For customers wishing to test File Persona, or any other HPE 3PAR feature, a NFR (Not For Resale) license is available for 180-day or one-year trial periods. FP is licensed on a per host presented TB basis.

SSMC offers two modes of management for File Persona for streamlined management experience. Refer to [Appendix D](#) for the steps to enable the advanced menu.


Normal mode

- Hides more complex configuration and management options
- Uses default values for hidden objects during creation
- Simplifies admin choices for everyday operations



Advanced mode

- Displays more complex configuration and management options
- Allows admin user to specify all values for all objects during creation
- Provides greater control for less commonly performed operations



Best practices for HPE 3PAR File Persona deployment

Following are best practices to keep in mind while deploying File Persona with respect to networking, storage layout, authentication, snapshots, and performance.

Networking

- The HPE 3PAR File Persona supports a dual port 10GbE NIC per node in network bond mode 6 balance-alb or network bond mode 1 (active/passive) or a quad port 1GbE NIC in each node in network bond mode 6 (balance-alb) with the option to use network bond mode 1 (active/passive). Bonding is support only on NICs on the same node.
- File Persona requires an IP address per node in addition to at least one IP address per Virtual File Server.

Best practice: Ensure the File Persona configuration uses the same Network Time Protocol (NTP) server as other servers. In the case of Active Directory it can be an AD server. Kerberos authentication allows a maximum time draft of five minutes, if the drift exceeds this Kerberos authentication will fail. Publically available servers include pool.ntp.org.

Best practice: The same NIC should be used in each node in the system e.g., either 1GbE or 10GbE. It's best to use dual port 10GbE NIC for File Persona to get more network bandwidth. It is also worthy to note that with HPE 3PAR OS version 3.2.2 and beyond, an onboard RCIP (Remote Copy Internet Protocol) port may be used for File Persona.

Best practice: Make sure that there are multiple network connections from node pairs (at least one port per node), preferably connected to at least two network switches for increased availability.

Authentication

Best practice: When configuring authentication order only enable valid authentication providers. For example if Active Directory is not available, leave it disabled.

Best practice: If using Active Directory preconfigure the computer account (for each node running File Persona) before VFS creation to avoid issues with computer account creation.

Storage layout

- The HPE 3PAR File Persona is enabled per node pair for high availability purposes. Upon enablement of File Persona, a RAID 6 volume will be created for each node in the array running the File Persona. These are the system volumes and cannot be used for creating the File Stores to share out to the clients.
- The HPE 3PAR File Persona can share the same CPGs as the block volumes to create the File Provisioning Group (FPG).

Best practice: Use file quotas to ensure that free space remains at 10 percent or greater to avoid problems associated with full file systems including failed writes.

Best practice: Create new CPGs for FPGs or use existing CPGs **except** fs_cpg to enable the greatest administrative flexibility (e.g., remote copy, migration to different performance tiers to isolate the CPGs for file). For File Persona, RAID 6 (FC) and RAID 6 (NL) CPGs provide the best capacity and resiliency.

Best practice: For the best combination of resiliency, and efficient capacity utilization, use RAID 5 for SSD and RAID 6 for FC and NL drives.

Best practice: Create at least one FPG/VFS/file store/file share on each node running the File Persona in the HPE StoreServ array. Distribute users and group/department/corporate shares evenly between the nodes for maximum load distribution.

Note

The type of network interface (whether the onboard interface or an add-on NIC) for the ports used by File Persona must be the same. File Persona cannot be enabled using both the onboard port and an add-on NIC at the same time.

Note

Any CPG being used by an FPG and VFS cannot be a member of a Virtual Domain.

Protocols

The HPE 3PAR File Persona supports SMB 3.0, 2.1, 2.0, 1.0 and NFSv4.0, v3.0 along with Object Access API. This includes advanced SMB 3.0 protocol feature of Transparent Failover, SMB opportunistic locks (oplocks) and leases (file and directory) for all SMB versions; crediting and large maximum transmission unit (MTU) size for SMB 2.x and beyond versions and Offloaded Data Transfer (ODX).

Best practice: For achieving transparent failover, leave continuous availability enabled and use SMB protocol version 3 at a minimum for non-disruptive operations to the clients.

Authentication and authorization

The HPE 3PAR File Persona supports three types of name services—Active Directory, LDAP and Local database for user and group authentication.

Best practice: Use Active Directory (if available) for the most flexible authentication and authorization for deployment scenarios with home directories and corporate/group shares.

Best practice: Unless every user in your environment has an Active Directory User ID (UID) avoid enabling rfc2307.

Best practice: If rfc2307 has to be enabled ensure that all users have UIDs and GIDs defined in AD and then enable rfc2307 support in File Persona before creating any FPGs.

Best practice: In order to maximize flexibility and manageability use Active Directory, LDAP and local in the order for the authentication stack configuration on File Persona via SSMC.

Best practice: After a user directory or a group/department/corporate share has been created, review and set the folder permissions to reflect the appropriate level of access and security for your organization.

Best practice: The File Persona supports Windows continuously available shares functionality and is enabled by default. Windows 8, 8.1 or Server 2012, 2012 R2 clients are required for this functionality. It is recommended that the configuration be verified in a non-production environment before being moved to the production environment.

Snapshots

The HPE 3PAR File Persona provides point-in-time space efficient, redirect-on-write snapshots at the File Store level.

Best practice: Schedule snapshots, with a specific retention count to provide a defined set of recovery point granularities over time. For instance:

- Take daily snapshots and retain 7.
- Take weekly snapshots and retain 4.
- Take monthly snapshots and retain 12.

Best practice: Create a schedule for snapshot clean tasks for every FPG. Snapshot clean tasks ensure that once snapshots are deleted, orphaned blocks are recovered and made available for new files. Create a weekly clean up task and monitor the amount of space returned to the system. If the amount returned is over 10 percent of the total system capacity, schedule clean up tasks every 3 days.

Backup and restore

The HPE 3PAR File Persona supports network share based backup over SMB or NFS protocol and NDMP over iSCSI based backup for the user data in file shares and the system configuration data in each VFS.

Best practice: Make sure to take a weekly backup of system configuration data for each VFS.

Note

For antivirus scanning with File Persona, refer to [Technical overview of HPE 3PAR File Persona Software Suite](#).

Note

[Appendix C](#) covers File Persona compatibility with other HPE 3PAR data services.

High availability

Best practice: Hewlett Packard Enterprise encourages all HPE 3PAR StoreServ Storage customers upgrade to the latest recommended HPE 3PAR OS. Upgrading to the most current HPE 3PAR GA OS ensures that the storage system benefits from the ongoing design improvements and enhancements. For customers participating in the Get 6-Nines Guarantee Program, the program will identify the latest HPE 3PAR OS version that is covered under the Guarantee program.

Best practice: Size the system appropriately so that all workloads and applications dependent on the HPE 3PAR system can perform as needed under the conditions of a node being down. This may occur during an unplanned controller node failure or planned maintenance of a controller node. In no situation should the maximum limits of the system as defined in this document and product specifications be exceeded.

In systems with four or more nodes, a resilience feature called Persistent Cache is automatically enabled. The Persistent Cache feature ensures that no storage controller node is placed into performance limiting “cache write thru” mode as a result of a losing its partner in the node pair. Any node that loses its adjacent node can dynamically form a mirrored cache relationship with another storage controller node. This limits the performance impact of unplanned downtime or controller node maintenance.

Persistent Ports

Persistent Ports functionality is supported for HPE 3PAR OS 3.1.2 and later only (with functionality restrictions on HPE 3PAR OS 3.1.2. Starting with HPE 3PAR OS 3.1.3 support for FCoE connected hosts and iSCSI connected hosts has been added, and the ability to detect an array node suffering “loss_sync” (a physical layer problem occurring between the HPE 3PAR controller node and the switch it is connected to) has been added. There is no Persistent Ports support on versions of the HPE 3PAR OS versions prior to 3.1.2.

For HPE 3PAR StoreServ FC host ports, the following requirements must be met:

- The same host port on host facing HBAs in the nodes in a node pair must be connected to the same FC fabric and preferably different FC switches on the fabric (for example, 0:1:1 and 1:1:1).
- The host facing HBAs must be set to “target” mode.
- The host facing HBAs must be configured for point-to-point connection (no support for “loop”).
- The FC fabric being used must support NPIV and have NPIV enabled.

For HPE 3PAR StoreServ ports FCoE host ports, the following requirements must be met:

- The same Converged Network Adapter (CNA) port on host facing HBAs in the nodes in a node pair must be connected to the same FCoE network and preferably different FCoE switches on the network (for example, 0:1:1 and 1:1:1).
- The FCoE network being used must support NPIV and have NPIV enabled.

For HPE 3PAR StoreServ iSCSI host ports, the following requirements must be met:

- The same host port on host facing CNAs in the nodes in a node pair must be connected to the same IP network and preferably different IP switches on the fabric (for example, 0:1:1 and 1:1:1).

Persistent Ports configuration considerations

Persistent Ports requires that corresponding “Native” and “Guest” host ports on a node pair be connected to the same FC fabric or IP network and the switches they are connected to must support and be configured for NPIV in the case of FC and FCoE. This means that for a minimum configuration to provide Persistent Ports functionality, where the node pair is connected to redundant FC SAN fabrics, each node in a node pair must have at least two FC host ports cabled with one port connected to each fabric.

Best practice: Ensure that the same <slot>:<port> on each node in the node pair are connected to the same FC fabric as this is a requirement to enable the Persistent Port functionality.

For more information on Persistent Ports, refer to the HPE 3PAR StoreServ Persistent Ports white paper [here](#).

Drive enclosure availability

Best practice: When creating CPGs, select HA cage availability if possible. There are three levels of availability that can be selected with HPE 3PAR StoreServ.

In the HPE 3PAR StoreServ 7000 and 8000 and 20000 series storage systems drive magazines consist of only a single drive.

Note

The term “Cage” is equivalent to the terms “Drive enclosure”, “Enclosure”, and “Drive Shelf”.

HA is an acronym for High Availability.

- HA CAGE means that no two members of the same RAID set can be in the same drive enclosure. For RAID 1 and 5, this means, for example, a RAID 5 set size of 4 (R5 3+1) requires four drives enclosure connected to each node pair. With the double drive failure protection of RAID 6 the number of drive enclosures required is reduced by half.
- HA MAG means that no two members of the same RAID set are in the same drive magazine. This allows a wider stripe with fewer drive enclosure; for example, a RAID 6 stripe size of 6+2 (set size 8) would be possible with only four-drive enclosure, provided each enclosure had at least two drive magazines.
- HA PORT applies only to daisy-chained drive enclosure. When this level of availability is selected, no two members of the same RAID set can be in drive enclosure that are dependent on one another for node connectivity. For example, in a system in which there are eight drive enclosure and four of the drive enclosure are connected to another drive enclosure for node access, HA PORT would only allow RAID 6 6+2 (set size 8) in order to prevent the loss of one drive enclosure causing a loss of data access.

Note

HA Port only applies to the 7000 and 8000 series HPE 3PAR arrays. The 10000 and 20000 series array enclosures directly connect to the controller nodes.

Priority Optimization

In HPE 3PAR OS 3.1.2 MU2, [Priority Optimization QoS](#) (Quality of Service) offered the ability to set a maximum limit of IOPS or bandwidth available to a VVset.

Since HPE 3PAR OS 3.1.3 priority levels (high, normal, and low), latency goals, and minimum goals along with the ability to set QoS rules against Virtual Domains have been available.

Care must be taken when setting QoS rules, system performance has to be understood, and historical performance data has to be taken into account when setting maximum limits, minimum goals, and latency goals.

A good understanding of the applications that reside on the VVs is also equally important.

Table 8. Priority Optimization capabilities by version

CONTROL TYPE	MINIMUM HPE 3PAR OS VERSION	DESCRIPTION	DEPENDENCIES AND BEST PRACTICES
Max. limit	3.1.2 MU2 and later	Maximum threshold for IOPS and/or bandwidth for a QoS object.	Maximum limit has no dependencies on the other control types.
Min. goal	3.1.3+	Minimum floor for IOPS and/or bandwidth below which HPE 3PAR Priority Optimization will not throttle a QoS object.	When a minimum goal is set on an object, the user must also configure a maximum limit on the same object within the same rule. Minimum goal will be ignored if the system has no rules with latency goal set.
Latency goal	3.1.3+	The Service Time target the system will try to achieve for a given workload in a QoS object.	This control type requires other rules in the system with minimum goal to be set as the latency goal algorithm needs direction on which workload to target and throttle. The order in which these will be throttled is provided by the priority levels.
Priority level	3.1.3+	Precedence order for QoS subsystem to throttle workloads to meet performance levels.	High priority should be used against critical applications, lower priority on less critical applications.

Best practice: The HPE 3PAR Priority Optimization white paper should be followed when configuring QoS rules. It is available for download [here](#).

Best practice: Maximum limit, this is the maximum amount of IOPS or bandwidth, or both, which a given VVset or domain is allowed to achieve. Best practice is to use the System Reporter data in order to quantify volume performance and set maximum limits rules accordingly.

Best practice: Minimum goal, this is the minimum amount of IOPS or bandwidth, or both, below which the system will not throttle a given VVset or domain in order to meet the latency goal of a higher priority workload. The minimum goal should be set by looking at the historical performance data and by understanding what is the minimum amount of performance that should be granted to the applications that reside in that VVset. The volumes in the VVset may use more IOPS/bandwidth than what is set by the minimum goal, but will be throttled to the given limit as the system gets busier. The performance may also go below the minimum goal; this may happen if the application is not pushing enough IOPS or if the sum of all minimum goals defined is higher than the I/O capability of the system or a given tier of storage.

Best practice: Latency goal—this is the svctime the system will goal to fulfill for a given QoS rule. In order for the goal to work, rules with a “minimum goal” specification must exist so the system can throttle those workloads. Reasonable latency goal should be set; this can be done by looking at historical performance data. The latency goal will also be influenced by the tier the volume resides on; below are some guidelines:

Table 9. Latency goals by drive type

TIER	LATENCY GOAL GUIDELINES
All NL	>= 20 ms
All FC	>= 10 ms
All SSD	>= 1 ms
Sub-tiered volumes (Adaptive Optimization)	Guideline of the middle tier in a 3-tier Adaptive Optimization configuration or of the lowest tier in a 2-tier Adaptive Optimization configuration

When a latency goal is set on a QoS object in a given tier, users should also create QoS rules with a minimum goal and lower priority on other objects that reside in the same tier. This will allow QoS to throttle VVsets that share the same tier towards their minimum goal if there is contention of resource on a that tier of storage in order to meet the latency goal of the higher priority workloads.

Best practice: Use VVsets for setting QoS rules as well as exporting VVs.

HPE 3PAR Adaptive Optimization and QoS interoperate, but the HPE 3PAR Adaptive Optimization data migration process may impact latency goal settings. For example, when a workload is partially migrated from FC drives into NL ones after the collection and analysis of the I/O access pattern by HPE 3PAR Adaptive Optimization, depending on the host access I/O pattern this may have a negative impact on latency and trigger workload to be throttled towards their minimum goal. This is an expected behavior if the QoS rule was modeled after the FC performance.

Best practice: Three priority levels exist: high, normal, and low. As the system gets busier it will start targeting lower priority workloads and throttling their performance in order to meet higher priority workloads latency goals. High priority level should be used against critical applications, lower priority on less critical applications.

For more information on Priority Optimization, refer to the HPE 3PAR Priority Optimization white paper [here](#).

Virtual volumes

Best practice: Zero detect should be enabled on TPVVs that are periodically “zeroed out”. Zero detect is enabled by default in HPE 3PAR OS version 3.1.2 and later.

Best practice: Thin provisioned virtual volumes can have an allocation warning, but should never have an allocation limit. Do not set an allocation limit, not even 100 percent.

Best practice: Virtual volumes should have both a “user CPG” and “copy CPG” selected. A “copy CPG” is required to use any of the following features:

- Clones (Previously known as Full Copies)
- Virtual copy (snapshots)
- Remote copy (remote replication)

Best practice: When choosing a snapshot CPG for VVs being replicated via Remote Copy choose the same tier of storage as the Base VVs and not a lower tier of storage.

Note

As of HPE 3PAR OS version 3.2.2 VVs automatically have snapshot space associated in the same CPG. During VV creation a separate snapshot CPG can be selected using the “additional settings” option in the SSMC or the `-snp_cpg` switch of the `createvv` CLI command.

The user CPG and copy CPG can be different CPGs. This can be useful if snapshots are to be created in a different class of service from the source data.

Unless this is specifically required, use the same CPG for user CPG and copy CPG.

Virtual LUNs (exports) and volume sets

Best practice: Virtual volumes should be exported to host objects, not to ports (port presents) for all hosts.

Best practice: “Matched sets” (a combination of export to hosts and exports to ports) should not be used.

Limiting the number of paths for a given host should only be done by zoning the host with only specific front-end ports.

Best practice: Boot from SAN virtual volumes should be exported with LUN 0.

Best practice: Use volume sets when exporting multiple virtual volumes to a host or host set.

Volume sets allow multiple volumes to be exported to the same host or host set in one operation. When exporting a volume set to a host or host set, the user selects which LUN ID will be used for the first virtual volume of the volume set, and the LUN IDs are incremented for every virtual volume in the volume set.

Best practice: Individual volumes (belonging only to one host) should be exported with LUN IDs starting from 1. Volume sets should be exported with LUN IDs starting from 10.

Best practice: Always leave a gap of LUN IDs between two volume sets. When adding a new virtual volume to a volume set that is already exported, the next LUN ID needs to be free or the operation will fail.

Note

As of HPE 3PAR OS version 3.2.2 this gap is created automatically.

Therefore, it is important that if multiple volume sets are exported to the same host/host set, there are gaps left for adding later virtual volumes to the volume set.

Example: If exporting two volume sets containing 10 virtual volumes each to a host set:

- Export the first volume set with LUN ID 10.
 - LUN IDs 10 to 19 will be used.
- Export the second volume set with LUN ID 30.
 - LUN IDs 30 to 39 will be used.

Remote Copy

Note

Remote Copy initial synchronization has been modified to throttle bandwidth utilization in order to not impact system performance.

Best practice: The **Remote Copy user guide** should be followed strictly when configuring Remote Copy. It is available for download at hpe.com/support. Particular attention should be paid to latency tolerance.

Best practice: Remote Copy FC ports should only be in one zone with one Remote Copy FC port from another HPE 3PAR StoreServ system.

Do not zone hosts or other devices with the Remote Copy FC other than the destination Remote Copy FC port.

As of HPE 3PAR OS version 3.1.3 has substantially improved topology flexibility.

Best practice: All virtual volumes belonging to the same application should be added to the same Remote Copy group. Virtual volumes that have logical links include:

- Virtual volumes used by the same application (data and log volumes, for example)
- Virtual volumes used by a Logical Volume Manager (LVM) volume group
- Virtual volumes that contain virtual disks of the same VM

A Remote Copy group can contain up to 300 virtual volumes.

Best practice: Do not add virtual volumes that have no logical link (host or application) to the same Remote Copy group. This will give the best granularity of Remote Copy failover, by allowing a failover of only one host or application.

Best practice: Source and destination virtual volumes must have the same size.

Best practice: While source and destination virtual volumes do not need to be of the same type of disk or RAID level, performance must be considered before mixing different types of disks/RAID levels on the source and destination systems.

When using Remote Copy in synchronous mode, it is not advised to replicate FC virtual volumes to NL virtual volumes, as the reduced performance of the NL disks might impact the primary virtual volume performance as the source system waits on the destination system to acknowledge the write.

Best practice: When using Remote Copy in combination with Virtual Domains, the source and destination virtual volumes must be in domains of the same name on the source and destination systems. However, they do not need to be of the same type (type of provisioning, RAID, or disk).

Best practice: In case of complete communication failure between two HPE 3PAR StoreServ systems, wait for a period of low activity before restarting the Remote Copy groups.

Best practice: If the hosts on the secondary sites are powered on, do not export the secondary (destination) virtual volumes unless using a geographically distributed cluster with automatic failover, such as CLX, Metrocluster, or Geocluster, where this is specifically required.

Streaming Asynchronous Replication

Asynchronous Streaming Remote Copy is perfect for environments, where very small RPOs are required. Or for environments where synchronous replication is desired but replicating the link latencies are so large it will result in unacceptable write I/O response times for the data to be replicated synchronously.

Best practice: Asynchronous Streaming Remote Copy solutions generally require replication link speed to be sized within 95–99 percent of the maximum data generation rate.

Best practice: Ensure the latency is under 10 ms; as data is placed in data cache at the source array until writes are acknowledged by the destination array higher latencies can result in source cache utilization of up to twenty percent and the source array to stop the replication group.

Note

Streaming Asynchronous Remote Copy is supported on FC and FCIP transports only.

Adaptive Optimization

Best practice: The following combinations are acceptable within the same [Adaptive Optimization](#) configuration (policy):

- SSD, FC/SAS, and NL
- FC/SSD and FC
- FC/SAS and NL

Using different RAID levels within the same policy is acceptable.

Note

Configurations with SSD and NL tiers only are not recommended; four controller configurations require the same type of drives behind controller pairs.

Best practice: When configuring two or three tier solutions containing SSDs, if region density data is not available for sizing the SSD tier assume the SSD tier will only provide the following IOPS per SSD drive:

- 200 GB SSD ~ = 550 IOPS
- 400 GB SSD ~ = 1000 IOPS¹
- 480 GB SSD ~ = 1100 IOPS²
- 920 GB SSD ~ = 2150 IOPS³

With specific regard to HPE 3PAR AO, the increase in the estimated number of IOPS on larger drives is not because of a difference in technology, but rather the increased probability that “hot” data regions are on the larger SSDs vs. the smaller SSDs.

Best practice: Always size the solution assuming the NL tier will contribute 0 percent of the IOPS required from the solution.

Best practice: Configurations that only contain SSD and NL are not recommended unless this is for a well-known application with a very small ratio of active capacity compared to the total usable capacity (1–2 percent).

For more information on Adaptive Optimization, refer to the Adaptive Optimization white paper [here](#).

Best practice: When using thin provisioning volumes along with Adaptive Optimization, select a CPG using FC disks for the user CPG of the thin provisioning volumes. This means that when new data is written, it will be on a mid-performance tier by default which can then be distributed appropriately.

Best practice: Ensure that the default tier (FC) has enough capacity and performance to accommodate the requirement of new applications until data is migrated to other tiers.

When new data is created (new virtual volumes or new user space for a thin volume), it will be created in the FC tier, and Adaptive Optimization will not migrate regions of data to other tiers until the next time the Adaptive Optimization configuration is executed.

It is therefore important that the FC disks have enough performance and capacity to accommodate the performance or capacity requirements of new applications (or applications that are in the process of being migrated to HPE 3PAR StoreServ) until the moment when the regions of data will be migrated to the other tiers.

^{1,2,3} Estimated IOPS for SSDs used in an Adaptive Optimization configuration. This is not reflective of how many IOPS and SSD can deliver.

Best practice: If SSDs are used in Adaptive Optimization configurations, no thin provisioning volumes should be directly associated with SSD CPGs. The thin provisioning volumes should only be associated with FC CPGs.

This will ensure that SSD capacity is consumed by Adaptive Optimization and will allow this capacity to be safely used to 95 percent or even 100 percent.

To help ensure that no TPVV is associated with an SSD or NL CPG, run the “showcpg” command and confirm that only the FC CPG reports a TPVV value greater than 0.

In the following example, only the FC CPG has TPVVs associated with it:

```

-----[MB]-----
      --Volumes--  -----Usage---  -----  Usc  -----  Snp  -----  Adm  -----
Id Name      Domain Warn% VVs      TPVVs  Usc  Snp  Total      Used      Total  Used  Total
4 AO_01_SSD  -      -      0      0      0   0   1201792    1201792   11648  0    16384
8 AO_01_FC   -      -      469    469    469  0   11665920   11665920  8560128 0    120832
12 AO_01_NL  -      -      0      0      0   0   29161088   29161088  729472  0    325632
    
```

Best practice: All CPGs used in an Adaptive Optimization configuration should have the same level of availability. Using a CPG with “magazine level” availability in an Adaptive Optimization configuration with CPGs with “cage level” availability will mean that all virtual volumes will have an effective availability equivalent to “magazine level”.

Best practice: CPGs for SSDs should use RAID 5, with the minimum growth increment supported (8 GB per node pair). CPGs for FC disks should use RAID 6, with the default growth increment.

CPGs for NL disks should use RAID 6, with the default growth increment.

Refer to the “[Common provisioning groups](#)” section for details of CPG best practices.

Best practice: Schedule the different Adaptive Optimization configurations to run at the same time, preferably at night. Adaptive Optimization will execute each policy in a serial manner but will calculate what needs to be moved at the same time.

Best practice: It is preferable not to set any capacity limit on the Adaptive Optimization configuration level, or on the CPG (no allocation warning or limit).

This will allow the Adaptive Optimization software to make excellent use of the different tiers available in the system. In HPE 3PAR OS version 3.1.2 and later, use the SSMC to set the CPG capacities.

Best practice: If a capacity limit is required for a given tier, set a capacity-warning threshold (not limit) on the CPG itself through the SSMC or the CLI. Adaptive Optimization will not attempt to move more data to a CPG than the capacity warning set on the CPG.

Best practice: Always ensure that at least one of the CPGs used by the Adaptive Optimization configuration does not have any growth warning or limit.

Ideally, this should be the CPG to which the virtual volumes are linked; the CPG in which new user space will be created if needed.

Best practice: Use a simple Adaptive Optimization configuration model as often as possible. For most applications, use generic Adaptive Optimization configurations that:

- Use all the tiers available in the system.
- Run during the days of the workweek only (for example, Monday–Friday).
- Execute once a day, preferably at night.
- Use a measurement/hours of 24.
- Use a mode of “balanced”.

For well-known applications that require a high level of performance, use tailored Adaptive Optimization configurations that:

- Preferably use all the tiers available in the system.
- Execute immediately at the end of the high-activity period.
- Use a measurement/hours that only covers the length of the high-activity period.
- Use a mode of “performance”.

For test environments where performance is not a requirement, use an Adaptive Optimization configuration that:

- Uses only FC and NL tiers.
- Run during the days of the workweek only (for example, Monday–Friday).
- Executes once a day, preferably at night.
- Uses a measurement/hours of 24.

Best practice: Do not mix Adaptive Optimization with any other application or process that moves data on a LUN or between LUNs.

Any application or process that moves data on a LUN or between LUNs, such as VMware Storage DRS, should be disabled, as they might conflict with each other.

Only Adaptive Optimization, being the lowest level (storage level) should be used.

Security

Best practice: Change the password for the “3paradm” user.

The password of the “3parcim” user can also be changed if the CIM/SMI-S service is enabled.

Do not change the password or remove the accounts for the “3parsvc”, “3parservice”, “3paredit”, or “3parbrowse” users. These are randomly generated at the time of the initialization of the system and are required for communication between the service processor and the HPE 3PAR StoreServ system.

Best practice: Create a different user for each system administrator that will use the system.

Alternatively, configure the system to use active directory and make sure all users use their own accounts to log in to the system.

Best practice: When scripting, use the lowest privilege level required.

If a script requires only read access to the system, use a browse account. If a script doesn’t need to remove objects, use a create account.

Naming conventions

A concise naming convention is important to effectively managing an HPE 3PAR system. A naming convention should be consistent and descriptive. HPE 3PAR systems are case sensitive. Preferably, use CamelCase to ensure ease of readability and a descriptive name. When using the bundled call home feature, use of a clear naming convention allows support personnel to ameliorate support calls more efficiently and by adding descriptive comments, further enable support case resolution.

Use hierarchical names that allow multiple objects of the same classification group to begin with the same characters. Examples:

```
prd.unix.datavg01.vv
```

```
prd.unix.appvg02.vv
```

Use a naming convention and a suffix that allows all objects of the same type to be grouped together when sorting on the name field and that allows effective search of all objects when using patterns (question mark [?] or asterisk [*]).

Examples: showcpg *.cpg.*

Define the naming convention early in the implementation. (See the following examples.)

Naming convention examples

Hosts

Host names support up to 31 characters.

Host name will be of the form <TYPE>.<OS>.<HOST>.<OBJECT TYPE>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <OS> can be win for Windows hosts, vmw for VMware hosts, lin for Linux hosts, sol for Solaris hosts, aix for AIX hosts, hpux for HP-UX hosts, etc.
- <HOST> is the host name.
- <OBJECT TYPE> is one of VV, cpg, VVset, etc.

Examples:

- prd.win.server1.vv
- prd.hpux.server2.vv
- dev.lin.server3.vv

Host sets

Host set names support up to 31 characters.

Host set name will be of the form <TYPE>.<OS>.<CLUSTER NAME>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <OS> can be win for Windows hosts, vmw for VMware hosts, lin for Linux hosts, sol for Solaris hosts, aix for AIX hosts, hpux for HP-UX hosts, etc.
- <CLUSTER NAME> is the name of the cluster.

Examples:

- prd.win.sqlcluster1.vvset
- prd.vmw.esxcluster2.vvset
- dev.lin.cluster3.vvset

Virtual volumes

VV names support up to 31 characters.

For standalone servers, VV name will be of the form <TYPE>.<OS>.<HOST>.<FS NAME>.<ID>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <OS> can be win for Windows hosts, vmw for VMware hosts, lin for Linux hosts, sol for Solaris hosts, aix for AIX hosts, hpux for HP-UX hosts, etc.
- <HOST> is the host name.
- <FS NAME> is the file system name or drive letter.
- <ID> is an optional ID for volume groups containing multiple LUNs.

Examples:

- prd.win.server1.e.vv
- prd.hpux.server2.datavg1.01.vv
- dev.lin.server3.data1.vv

For clusters, VV name will be of the form <TYPE>.<OS>.<CLUSTER>.<RES NAME>.<ID>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <OS> can be win for Windows hosts, vmw for VMware hosts, lin for Linux hosts, sol for Solaris hosts, aix for AIX hosts, hpux for HP-UX hosts, etc.
- <CLUSTER> (optional) contains the cluster name.
- <RES NAME> is the resource name of the cluster object mapped to this virtual volume.
- <ID> is an optional ID for volume groups containing multiple LUNs.

Examples:

- prd.win.cluster1.sql1.m.vv
- prd.vmw.datastore1.vv
- prd.aix.datavg1.01.vv

Virtual Volume sets

VVset names support up to 31 characters.

VVset name will be of the form <TYPE>.<OS>.<HOST/CLUSTER>.<FS NAME>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <OS> can be win for Windows hosts, vmw for VMware hosts, lin for Linux hosts, sol for Solaris hosts, aix for AIX hosts, hpux for HP-UX hosts, etc.
- <HOST/CLUSTER> is the host name or cluster name.
- <FS NAME> is the OS name for the group of LUNs, volume group, and datastore group.

Examples:

- prd.win.hyperv.vmdata1.vvset
- prd.vmw.esx.datastoregrp2.vvset
- dev.lin.server3.datavg.vvset

Naming conventions with File Persona Objects

Note

File Provisioning Groups and Virtual File Servers do not support the use of periods (.) in their name, use an underscore instead.

Virtual File Servers should use <Department name or abbreviation> and the suffix _vfs

File Provisioning Groups should use <Department name> and the suffix _fpg.

External System Reporter; now EOL (End of Life)

HPE 3PAR System Reporter has been traditionally run from an external system prior to HPE 3PAR version 3.1.2. The External System Reporter (ESR) is now EOL (End of Life) support is available until June 30, 2016. For more complete information on System Reporter and its use with the SSMC (StoreServ Management Console), refer to the use guide which can be found [here](#).

On node System Reporter data can be queried via CLI, HPE 3PAR Management Console, and [HPE 3PAR StoreServ Management Console \(SSMC\)](#).

Best practice: For users that are not interested in Scheduled Reports the best practice is to use SSMC for System Reporter capabilities.

Best practice: For users that don't have needs to control how long their historical performance data is retained, it's recommended to use the reporting functionalities from HPE 3PAR StoreServ Management Console. This provides historical reporting by using the data stored on the internal database that resides in the HPE 3PAR OS. It will also provide an indication of how long the data is retained.

System reporter in the SSMC

Object selection

There are two selections available. By default, All is selected—which enables the user to chart all objects for the category selected. For example, if the user selects to chart real-time exported volume performance, the choices for object selection is all, or Filter by objects. Using the default of All then charts performance of all volumes created on the array. Selecting to filter by objects allows the user the ability to choose which object is charted. A user can choose to monitor an object by the user of a virtual volume, host, or port number. The user may highlight several objects and use the Add radial button at the bottom. If the user needs to select multiples items that are not all displayed, the user can highlight one object, depress Add + radial button, scroll down the list of menu items and add additional objects by either using Add + for multiple objects, or Add if only one additional object is added.

A user may also search for objects to chart by using the search bar identified by the magnifying glass icon. Using Filter by objects, the user must also choose plotting style. The two choices are:

- Individual series for each selected objects
- Aggregated series for all selected objects

Individual plot for each selected object is applicable only when Filter by objects is selected. For all objects, this option is not valid.

Best practice—Use the search bar to identify objects to enter rather than scrolling through menu items.

Best practice—Rule of thumb. Only display one rule at a time, combining rules does not allow the user to differentiate the outcome of the object.

Appendices D and E have examples of how to get started. For more information on the HPE 3PAR SSMC click [here](#).

Ongoing management and growth

Storage Analytics on the Web

Hewlett Packard Enterprise introduces StoreFront Remote—storefrontremote.com/ a Storage Analytics solution built on the HPE 3PAR Central call home infrastructure.

Best practice: Configure your HPE 3PAR Storage Array to send configuration information to HPE using the HPE 3PAR Service Processor. This enables Hewlett Packard Enterprise to offer meaningful analytics about your Storage devices, advisories to continuously optimize your Storage infrastructure and proactive support.

Best practice: Register your HPE 3PAR StoreServ arrays in SF Remote by logging in using your HPE Passport ID and by following the steps listed in the portal once you log in. Note that you'll be able to view your devices in SF Remote 24–48 hours after you register.

Best practice: Use System Access Groups in SF Remote to grant access rights to other users.

Best practice: Use the Column filter options and the “Status Filters” to create your combination of filters and save it as a “Custom Filter” for reuse and easy retrieval.

Best practice: Use the “notes” feature to tag Devices with comments and share it with other if you wish.

Autonomic rebalance

Starting with HPE 3PAR StoreServ 3.1.3, the rebalancing of data after hardware upgrades is now an integrated feature, not requiring a Dynamic Optimization license. After purchasing and installing new hardware, the rebalancing can be started by selecting the HPE 3PAR StoreServ in the SSMC and selecting “tune system”. The process is automatically invoked after the admithw command is issued on the 7200 and 8200 platforms.




Figure 4. Autonomic rebalancing [tunesys] selection

Autonomic rebalancing (tune system) can also be started from the CLI by typing tunesys.

Depending on the amount of hardware added to a system and the degree to which the system may have been unbalanced before, the re-leveling process can take several hours to several days or longer to complete. It is also possible that if the system was near capacity before the hardware upgrade it may take multiple runs of tunesys to level the system out. Other variables will also affect how long tunesys takes to complete including the speed of the drives affected by the re-leveling, how close to capacity the previously installed drives are at upgrade, how busy the system is, etc. The autonomic rebalancing process will give priority to servicing host I/O.

Best practice: Execute tunesys during off peak hours. Leave the leveling tolerance at the default values.

Appendix A. Supported host personas

PERSONA_ID	PERSONA_NAME	PERSONA_CAPS
1	Generic	UAREpLun, SESLun
2	Generic-ALUA	UAPreLun, SESLun, RTPG, ALUA
7	HP-UX-legacy	VolSetAddr, LunOSCC
8	AIX-legacy	NACA
9	Egenera	Softinq
10	ONTAP-Legacy	Softinq
11	VMware	SubLun, ALUA
12	OpenVMS	UAREpLun, RTPG, SESLun, LunOSCC
13	HP-UX	UAREpLun, VolSetAddr, SESLun, ALUA, LunOSCC
15	Windows Server®	UAPeLun, SESLun, ALUA, WSC

Appendix B

Block persona scalability limits

MODEL	MAXIMUM VVS
7200	1,024
7400/7450	2,048
10400	4,096
10800	8,192
20400	4,096
20840/20850	8,192

MODEL	MAXIMUM VLUN
7200	65,536
7400/7450	65,536
10400	131,072
10800	131,072
20400/20840/20850	131,072

File persona scalability limits

ITEM	LIMIT
File Provisioning Groups per node pair	16
VFS per FPG	1
File Stores per VFS	16
Max size per File Provisioning Group	32 TB
Max aggregate capacity of all FPGs, per node pair	64 TB
Max SMB Shares per node pair	4000
Max NFS Shares per node pair	1024

Appendix C. File Persona compatibility with HPE 3PAR block features

HPE 3PAR DATA SERVICES	FILE PERSONA
Adaptive Optimization	✓
Dynamic Optimization	✓
Thin Provisioning	✓
Thin Persistence	✗ ⁴
Thin Conversion	✓
Thin Copy Reclamation	✓
Adaptive Flash Cache	✓
Thin Deduplication	✓
Virtual Domains	✗
Priority Optimization (QoS)	✗
Virtual Copy	✓
Remote Copy	✓
Peer Persistence	✗
Peer Motion	✗

⁴ File Persona supports thin built-in zero detect.

Appendix D. Using the StoreServ management console (SSMC) to administer File Persona

HPE StoreServ OS version 3.2.1 MU2 introduces the File Persona for SMB and NFS access to the StoreServ array. As of the release of 3.2.1 MU2 HPE 3PAR has begun a transition to unified/converged administration of both block and file storage. This new management tool is the “StoreServ Management Console” (SSMC).

The SSMC can be downloaded from the HPE software depot.

1. Upon launching the SSMC the first time, set the admin password and add StoreServs to be managed.
2. After adding the StoreServs to be managed with their username and passwords, use the action menu to accept the SSL certificate.
3. After the initial configuration, login to the console using the StoreServ login (not the admin login).
4. Login and then select the down karat next to the HPE 3PAR StoreServ logo.



Figure 5. StoreServ management console dashboard

5. Select the file persona menu group and then “Persona Configuration”.



Figure 6. Main menu for SSMC functions

6. Select “Configure File Persona”




Figure 7. Configure File Persona

7. On the file persona configuration page, click on the settings gear to set the IP addresses of each individual node as well as the subnet mask for all addresses.
8. Note that the creation of the File Persona will take at least 5 minutes on a mostly idle system.




Figure 8. File Persona network configuration

- Configure IP addresses for each node, in figure 10 the network addresses are configured for node pair 0, 1.

Figure 9. Node network address configuration

- Lower on the file persona configuration page, scroll down and set the IP address of a DNS server. Optionally, configure active directory services for authentication and authorization.

Figure 10. File Persona setting

- Press create to complete the step.

12. Next select “Create virtual file server”




Figure 11. Create virtual file server

13. Create a Virtual File Server, supplying a name, select a CPG and assign an IP address for the SMB share. Note that selecting the advanced options in the upper right hand corner, will allow selection of an FPG, if FPGs have been created. Note that it will take several minutes to create the VFS depending on factors such as the FPG size and the current utilization of the array.




Figure 12. VFS block storage settings

- After the Virtual File Server has been created shares need to be added to make the data accessible. It is important to verify at this point that an authentication and authorization set (Active Directory for example) have been set. Writing data to the VFS share and changing access and authorization afterwards can result in loss of access to data.




Figure 13. Create File Share

- Give the share a name and select a VFS/file store combination.




Figure 14. Name the File Share

16. Set the permissions on the share.




Figure 15. Set the permissions on the share

17. Verify connectivity to the share by browsing to Windows explorer to browse to the share.




Figure 16. Verify connectivity to the share

Note

Shares on the VFS can also be managed using the Shared folders MMC plugin.

Appendix E. Using the StoreServ management console (SSMC) to manage snapshots

1. Highlight the VFS that contains the volumes to snapshot and from the action menu select “Create file snapshot”.
2. Create a schedule for snapshots. In this example, the volumes will be snapped every hour for 1 week. As the 168 mark is reached the array will continue to generate snapshots, always removing the oldest before creating the newest.

The screenshot shows the 'Create File Snapshot' dialog box with the following details:

- General Tab:**
 - System: 3.2.1.128
 - Virtual file server: VFS_test
 - File store: home
 - Tag: home
 - Retention: 168
- Schedule Tab:**
 - Schedule name: Hourly
 - Alert: Generate if task failed
 - Schedule pattern: Hourly
 - Starting at: 0 Minutes after the hour

Buttons: OK, Cancel

Figure 17. Create scheduled snapshots

3. In addition, create a schedule to maintain weekly snap shots for a period of 3 months. The schedules are available by selecting schedules from the main drop down menu in the upper left of the SSMC interface.

Summary

HPE 3PAR StoreServ Storage is the last storage architecture you will ever need—regardless of whether your organization is a small- or medium-sized business or a large global enterprise. With a range of models to meet the needs of small to large data centers running key business applications up through enterprise-wide deployments of mission-critical applications and beyond, HPE 3PAR StoreServ Storage has you covered. It's storage that offers the effortless performance and flexibility you need to accelerate new application deployment and support server virtualization, the cloud, IT as a service (ITaaS), or whatever else your future may hold. It's one of the most advanced storage platforms, which is here to help you master unpredictability—effortlessly, without exception, and without compromise.

Learn more at

hpe.com/us/en/storage/3par.html



Sign up for updates

★ Rate this document



© Copyright 2012–2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Windows Server and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle is a registered trademark of Oracle and/or its affiliates. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.



**Hewlett Packard
Enterprise**

HPE 3PAR StoreServ Architecture

Contents

HPE 3PAR StoreServ hardware architecture overview	5
HPE 3PAR ASIC	5
Full-mesh controller backplane	6
Active/Active vs. Mesh-Active	6
System-wide striping.....	7
Controller node architecture	7
Drive enclosures	10
Highly virtualized storage operating system.....	11
Fine-grained approach to virtualization.....	11
Multiple layers of abstraction	11
Logical disks.....	12
LD Ownership.....	13
Common provisioning groups.....	13
Virtual volumes.....	13
VLUNs and LUN masking.....	14
Metadata handling.....	14
System-wide sparing.....	14
Flash-optimized innovations.....	15
Workload-centric Storage Personas	16
HPE 3PAR Block Persona.....	16
HPE 3PAR File Persona.....	16
Multi-tenant architecture benefits.....	17
Tier-1 resiliency.....	17
Hardware and software fault tolerance.....	18

Advanced fault isolation and RAID protection	18
Controller node redundancy.....	19
Data integrity checking	19
Memory fencing.....	20
Persistent technologies.....	20
HPE 3PAR Persistent Checksum	20
HPE 3PAR Persistent Cache	21
HPE 3PAR Persistent Ports.....	22
HPE 3PAR Remote Copy.....	22
HPE 3PAR Virtual Domains	23
Data encryption.....	23
Maintaining high and predictable performance levels.....	24
Load balancing	24
Mixed-workload support.....	24
Storage quality of service.....	25
Adaptive Flash Cache	25
Express Writes.....	26
Performance benefits of system-wide striping.....	26
Bandwidth and communication.....	27
Data transfer paths.....	27
Sharing and offloading of cached data.....	28
Pre-fetching.....	29
Write caching.....	29
Fast RAID 5.....	29
HPE 3PAR RAID MP (Fast RAID 6)	29
Data compaction.....	30
HPE 3PAR ASIC for inline data reduction.....	30
HPE 3PAR Thin Deduplication and Thin Clones.....	31
HPE 3PAR Thin Provisioning.....	31
HPE 3PAR Thin Persistence	31
HPE 3PAR Thin Conversion.....	32
HPE 3PAR Thin Copy Reclamation	32
HPE 3PAR Virtual Copy.....	32
Autonomic storage management.....	33
Ease of use.....	33
HPE Smart SAN	34

HPE StoreFront Remote Management (SFRM).....	35
Data optimization.....	35
Autonomic rebalancing.....	36
HPE 3PAR Storage Federation.....	36
HPE 3PAR Peer Motion.....	37
HPE 3PAR Online Import.....	38
HPE 3PAR Peer Persistence.....	38
Proactive support.....	40
Summary.....	41
For more information.....	41

Modern, Tier-1 Storage for the New Style of Business

HPE 3PAR StoreServ Storage provides a single product family to meet your primary storage needs—regardless of whether you are a midsize enterprise experiencing rapid growth, a large enterprise looking to support IT as a Service (ITaaS), or a global service provider building a hybrid or private cloud.

HPE 3PAR StoreServ Storage supports true convergence of block, file, and object access while offering the performance and flexibility that you need to accelerate new application deployments and support server virtualization, the cloud, ITaaS, and your future technology initiatives. It's a storage platform that allows you to spend less time on management, gives you technically advanced features for less money, and eliminates trade-offs that require you to sacrifice critical capabilities such as performance and scalability. With HPE 3PAR StoreServ Storage, you can serve unpredictable and mixed workloads, support unstructured and structured data growth, and meet block, file, and object access storage needs from a single capacity store.




Figure 1. HPE 3PAR StoreServ portfolio

The all-flash HPE 3PAR StoreServ models—20850, 20450, and 8450—offers massive performance scalability from hundreds of thousands of IOPS to millions of IOPS all within a sub-millisecond response time. These all-flash systems uniquely offer performance without any compromises, be it the enterprise class Tier-1 data services (such as 3-data-center replication) or resiliency using their scalable (>2-nodes) architecture. These systems are well suited for massive workloads consolidation, and should be the first choice for customers building their next generation (all-flash) data centers. Furthermore, HPE 3PAR StoreServ 8200 and 8400 All-Flash Starter Kit are designed to offer an incredibly affordable entry point for flash performance without limiting scalability, functionality, or data mobility.

The converged flash HPE 3PAR StoreServ models—20800, 20840, 8440, 8400, 8200—leverage a flash-first approach that offers flash performance with the additional flexibility to add spinning media (HDD) as a tier. In an all-flash configuration, these systems also deliver hundreds of thousands to millions of IOPS all within a sub-millisecond response time. 20800, 20840, and 8440 are hyper-scalable models, supporting multi-PB raw capacity and scalability, whereas 8400 and 8200 balance performance and scalability and offer a very economic starting point for customers looking to take advantage of the HPE 3PAR architecture. Furthermore, HPE 3PAR StoreServ 20800 Starter Kit is designed to offer incredibly affordable entry point for Tier-1 Enterprise Flash Array that can offer true enterprise class consolidation.

It is important to remember that all HPE 3PAR StoreServ models are built on a single architecture, run the exact same HPE 3PAR Operating System and offer a common set of enterprise data services, be it the 8200, the 20800 or previous generations of HPE 3PAR StoreServ. HPE 3PAR StoreServ arrays can natively replicate and federate amongst each other without the need for any external replication or virtualization appliance.

We had created the best storage arrays leveraging the proven 3PAR architecture, which were flash optimized further to offer our customers investment protection and well, everyone took notice. With proven and affordable performance for large file processing, large database queries and on-demand video workloads, HPE 3PAR StoreServ 20850 had set world record in the recent SPC-2 benchmarks and is #1 array across board for performance and #1 array for price performance across all All-Flash Arrays. And counting.

This white paper describes the architectural elements of the HPE 3PAR StoreServ Storage family that deliver Tier-1 resiliency across midrange, all-flash, and tiered storage arrays, as well as by supporting a single operating system and rich set of data services across the portfolio.

HPE 3PAR StoreServ hardware architecture overview

Each HPE 3PAR StoreServ Storage system features a high-speed, full-mesh passive interconnect that joins multiple controller nodes (the high-performance data movement engines of the HPE 3PAR StoreServ Architecture) to form a cache-coherent, Mesh-Active cluster. This low-latency interconnect allows for tight coordination among the controller nodes and a simplified software model.

In every HPE 3PAR StoreServ Storage system, each controller node has a dedicated link to each of the other nodes that operates at 4 GiB/s in each direction. In an HPE 3PAR StoreServ 20800 Storage system, a total of 56 of these links form the array's full-mesh backplane. In addition, each controller node may have one or more paths to hosts—either directly or over a storage area network (SAN). The clustering of controller nodes enables the system to present hosts with a single, highly available, high-performance storage system. This means that servers can access volumes over any host-connected port—even if the physical storage for the data is connected to a different controller node. This is achieved through an extremely low-latency data transfer across the high-speed, full-mesh backplane.

The HPE 3PAR StoreServ Architecture can be scaled from 1.2 TiB to 6 PB of raw capacity, making the system deployable as a small, remote office system, or a very large centralized system. Until now, enterprise customers were often required to purchase and manage at least two distinct primary architectures to span their range of cost and scalability requirements, well that changes now. HPE 3PAR StoreServ Storage is the ideal platform for virtualization and cloud computing environments. The high performance and scalability of the HPE 3PAR StoreServ Architecture is well suited for large or high-growth projects, consolidation of mission-critical information, demanding performance-based applications, and data lifecycle management. High availability is also built into the HPE 3PAR StoreServ Architecture through full hardware redundancy. Controller node pairs are connected to dual-ported drive enclosures. Unlike other approaches, the system offers both hardware and software fault tolerance by running a separate instance of the HPE 3PAR Operating System on each controller node, thus facilitating the availability of customer data. With this design, software and firmware failures—a significant cause of unplanned downtime in other architectures—are greatly reduced.

HPE 3PAR ASIC

The HPE 3PAR StoreServ 20000 and 8000 systems use the fifth and latest generation of the HPE 3PAR ASIC, the HPE 3PAR Thin Express ASIC. The HPE 3PAR Thin Express ASIC is engineered and designed for solid-state performance. The ASIC enables the new 20000 and 8000 series to deliver up to 5X improvement in system bandwidth and faster XOR operations. It works in parallel with the CPU, evenly processing the I/O workload across the node Active—Mesh scale—out architecture, ensuring lower latency resulting in better system bandwidth.

The HPE 3PAR ASICs also features a uniquely efficient, silicon-based zero-detection and deduplication mechanism that gives HPE 3PAR StoreServ Storage systems the power to perform inline deduplication and remove allocated but unused space with minimal impact to performance. The HPE 3PAR ASICs also deliver mixed-workload support to alleviate performance concerns and cut traditional array costs.

Transaction- and throughput-intensive workloads run on the same storage resources without contention, thereby cutting array purchases in half. This is particularly valuable in virtual server environments, where HPE 3PAR StoreServ Storage boosts virtual machine density so you can cut physical server purchases.

The Thin Express ASIC also enables Persistence Checksum that delivers T10-PI (Protection Information) for end-to-end data protection (against media and transmission errors) with no impact to applications or host operating systems.

Full-mesh controller backplane

Backplane interconnects within servers have evolved dramatically over the years. Most, if not all, server and storage array architectures have traditionally employed simple bus-based backplanes for high-speed processor, memory, and I/O communication. Parallel to the growth of SMP-based servers, significant investments were also made to switch architectures, which have been applied to one or two enterprise storage arrays.

The move from buses to switches was intended to address latency issues across the growing number of devices on the backplane (more processors, larger memory, and I/O systems). Third-generation full-mesh interconnects first appeared in the late 1990s in enterprise servers. Figure 2 shows the full-mesh backplane of an HPE 3PAR StoreServ 20000 Storage system.




Figure 2. Full-mesh backplane of an HPE 3PAR StoreServ 20000 Storage system

The HPE 3PAR StoreServ full-mesh backplane is a passive circuit board that contains slots for up to four or eight controller nodes, depending on the model. As noted earlier, each controller node slot is connected to every other controller node slot by a high-speed link (4 GiB/s in each direction, or 8 GiB/s total), forming a full-mesh interconnect between all controller nodes in the cluster—something that Hewlett Packard Enterprise refers to as a Mesh-Active design. These interconnects deliver low-latency, high-bandwidth communication and data movement between controller nodes through dedicated point-to-point links and a low overhead protocol that features rapid inter-node messaging and acknowledgment. It’s important to note that, while the value of these interconnects is high, the cost of providing them is relatively low. In addition, a completely separate full-mesh network of serial links provides a redundant low-speed channel of communication for exchanging control information between the nodes.

The HPE 3PAR StoreServ 20000 features an eight node capable backplane that supports two to eight controller nodes. HPE 3PAR StoreServ 8000 Storage systems feature either a dual-node or quad-node-capable systems that is essentially an equivalent of what was used in erstwhile enterprise class arrays which offer the same high-speed links between nodes.

Active/Active vs. Mesh-Active

Most traditional array architectures fall into one of two categories: monolithic or modular. In a monolithic architecture, being able to start with smaller, more affordable configurations (i.e., scaling down) presents challenges. Active processing elements not only have to be implemented redundantly, but they are also segmented and dedicated to distinct functions such as host management, caching, and RAID/drive management. For example, the smallest monolithic system may have a minimum of six processing elements (one for each of three functions, which are then doubled for redundancy of each function). In this design—with its emphasis on optimized internal interconnectivity—users gain the Active/Active

processing advantages of a central global cache (e.g., LUNs can be coherently exported from multiple ports). However, these architectures typically involve higher costs relative to modular architectures.

In traditional modular architectures, users are able to start with smaller and more cost-efficient configurations. The number of processing elements is reduced to just two, because each element is multifunction in design—handling host, cache, and drive management processes. The trade-off for this cost-effectiveness is the cost or complexity of scalability. Because only two nodes are supported in most designs, scale can only be realized by replacing nodes with more powerful node versions or by purchasing and managing more arrays. Another trade-off is that dual-node modular architectures, while providing failover capabilities, typically do not offer truly Active/Active implementations where individual LUNs can be simultaneously and coherently processed by both controllers.

The HPE 3PAR StoreServ Architecture was designed to provide cost-effective single-system scalability through a cache-coherent, multi-node clustered implementation. This architecture begins with a multifunction node design and, like a modular array, requires just two initial controller nodes for redundancy. However, unlike traditional modular arrays, enhanced direct interconnects are provided between the controllers to facilitate Mesh-Active processing. Unlike legacy Active/Active controller architectures—where each LUN (or volume) is active on only a single controller—this Mesh-Active design allows each LUN to be active on every controller in the system, thus forming a mesh. This design delivers robust, load-balanced performance and greater headroom for cost-effective scalability, overcoming the trade-offs typically associated with modular and monolithic storage arrays.

System-wide striping

Through a Mesh-Active design and system-wide striping, the HPE 3PAR StoreServ Architecture can provide the best of traditional modular and monolithic architectures in addition to massive load balancing.

The HPE 3PAR StoreServ Mesh-Active design not only allows all volumes to be active on all controllers, but also promotes system-wide striping that autonomously provisions and seamlessly stripes volumes across all system resources to deliver high, predictable levels of performance. System-wide striping of data provides high and predictable levels of service for all workload types through the massively parallel and fine-grained striping of data across all internal resources (disks, ports, loops, cache, processors, etc.). As a result, as the use of the system grows—or in the event of a component failure—service conditions remain high and predictable. Unlike application-centric approaches to storage, HPE 3PAR StoreServ Storage provides autonomic rebalancing that enables the system to evenly balance and use all available physical resources. This is particularly important with hardware upgrades since existing data should be rebalanced and stripped across new available resources. On HPE 3PAR StoreServ Storage, this is done without service disruption or preplanning.

For flash-based media, fine-grained virtualization combined with system-wide striping drives uniform I/O patterns by spreading wear evenly across the entire system. Should there be a media failure, system-wide sparing also helps guard against performance degradation by enabling a many-to-many rebuild, resulting in faster rebuilds. Because HPE 3PAR StoreServ Storage autonomically manages this system-wide load balancing, no extra time or complexity is required to create or maintain a more efficiently configured system.

A detailed discussion of resource allocation, including the system's virtualized tri-layer mapping methodology, is provided in the section [“Highly virtualized storage operating system.”](#)

Controller node architecture

An important element of the HPE 3PAR StoreServ Architecture is the controller node and it is a powerful data movement engine that is designed for mixed workloads. As noted earlier, a single system, depending on the model, is modularly configured as a cluster of two to eight controller nodes. This modular approach provides flexibility, a cost-effective entry footprint, and affordable upgrade paths for increasing performance, capacity, connectivity, and availability as needs change. In addition, the minimum dual-controller configuration means that the system can withstand an entire controller node failure without impacting data availability. Controller nodes can be added in pairs to the cluster non-disruptively, and each node is completely hot-pluggable to enable online serviceability.

Unlike legacy architectures that process I/O commands and move data using the same processor complex, the HPE 3PAR StoreServ Storage controller node architecture separates the processing of control commands from data movement, which helps ensure that CPU bandwidth is available for control processing and is not used for bulk data transfer. This innovation eliminates the performance bottlenecks of existing platforms that use a single processing element to serve competing workloads, for example online transaction processing (OLTP) and data warehousing workloads.

The HPE 3PAR ASIC within each controller node performs parity calculations on the data cache. The Zero-Detect mechanism built into the ASIC allows a hardware-assisted fat-to-thin volume conversion in conjunction with HPE 3PAR Thin Conversion software that enables users to take “fat” provisioned volumes on legacy storage and convert them to “thin” provisioned volumes on the HPE 3PAR StoreServ Storage system, this takes place inline and non-disruptively during the migration. This Zero-Detect capability also removes streams of zeroes present in I/O prior to writing data to the back-end storage system in order to reduce capacity requirements and prolong SSD life span. The HPE 3PAR ASIC is also a crucial element of the system’s ability to perform inline, block-level Thin Deduplication with Express Indexing (see the “Thin Deduplication with Express Indexing” section for more details).

HPE 3PAR StoreServ 20800 offers scalability in its truest sense not only with support for 1920 drives and 6 PB raw capacity but also with a port count of up to 160 FC ports for host connectivity and for consuming other advanced data services like replication, host persona and federation. Each of these ports is connected directly on the I/O bus, so all ports can achieve full bandwidth up to the limit of the I/O bus bandwidths that they share.

Each HPE 3PAR StoreServ 20000 series controller node can have a maximum of the following ports per node:

- (20) Quad-port 16 Gbps Fibre Channel adapter
- (10) Dual-port 10 Gbps iSCSI or Fibre Channel over Ethernet (FCoE) converged network adapter
- (6) 10 Gbps Ethernet adapter for file and object access services using File Persona

With all the host ports available on an 8-controller configuration, HPE 3PAR StoreServ 20000 series systems offer abundant multiprotocol connectivity. For back-end connectivity, each node can have up to (3) Quad-port 12 Gbps SAS Adapters.




Figure 3. Controller node design of the HPE 3PAR StoreServ 20000 Storage system

Each HPE 3PAR StoreServ 8000-series controller node has two built-in 4 port 16 Gbps Fibre Channel ports and one PCIe expansion slot. This slot can hold one of the following adapters:

- Quad-port 16 Gbps Fibre Channel adapter
- Dual-port 10 Gbps iSCSI or Fibre Channel over Ethernet (FCoE) converged network adapter
- Dual-port 10 Gbps Ethernet adapter for file and object access services using HPE 3PAR File Persona
- Quad-port 1 Gbps Ethernet adapter for file and object access services using HPE 3PAR File Persona

With up to 24 ports available on a quad-controller configuration, HPE 3PAR StoreServ 8000-series systems offer abundant multiprotocol connectivity. For back-end connectivity, each node has two built-in 2 x 4-lane 12 Gbps SAS ports.




Figure 4. HPE 3PAR StoreServ 8000 Controller Node

Across all 8000-series models, this controller node design extensively leverages commodity parts with industry-standard interfaces to achieve low costs and keep pace with industry advances and innovations. At the same time, the HPE 3PAR ASICs add crucial bandwidth and direct communication pathways without limiting the ability to use industry-standard parts for other components. Processor specifications by HPE 3PAR StoreServ model are shown in table 1.

Table 1. Processor specifications by HPE 3PAR StoreServ Storage system model

MODEL	CPUS	CONTROLLER NODES	TOTAL ON-NODE CACHE	TOTAL CACHE INCLUDING OPTIONAL FLASH CACHE
8200	Intel 6 core * 2.2 GHz/per controller node	2	Up to 64 GiB	832 GiB
8400	Intel 6 core * 2.2 GHz/per controller node	2 or 4	Up to 128 GiB	1664 GiB
8440	Intel 10 core * 2.4 GHz/per controller node	2 or 4	Up to 384 GiB	8384 GiB
8450	Intel 10 core * 2.4 GHz/per controller node	2 or 4	Up to 384 GiB	Not applicable
20800	2 * Intel 6 core * 2.5 GHz/per controller node	2, 4, 6, or 8	Up to 1.8 TiB	33.8 TiB
20840	2 * Intel 8 core * 2.5 GHz/per controller node	2,4,6, or 8	Up to 3.6 TiB	51.6 TiB
20850	2 * Intel 8 core * 2.5 GHz/per controller node	2, 4, 6, or 8	Up to 3.6 TiB	Not applicable
20450	2 * Intel 8 core * 2.5 GHz/per controller node	2 or 4	Up to 1.8 TiB	Not applicable

Drive enclosures

Another key element of the HPE 3PAR StoreServ Storage system is the drive enclosure or drive chassis, which serves as the capacity building block within the system. This section looks in more detail at the different drive enclosures supported by HPE 3PAR StoreServ Storage systems.

Table 2. Capacity building block for HPE 3PAR StoreServ 20000 systems

	HPE 3PAR STORESERV 20800	HPE 3PAR STORESERV 20840	HPE 3PAR STORESERV 20850	HPE 3PAR STORESERV 20450
Controller Nodes	2, 4, 6, 8	2,4,6,8	2, 4, 6, 8	2, 4
Max capacity (raw)	6000 TiB	6000 TiB	3932 TiB	1966 TiB
Max usable file capacity	512 TiB	512 TiB	512 TiB	256 TiB
Hard drives	8-1920	8-1920	NA	NA
SSDs	6-1024	6-1024	6-1024	6-512

Table 3. Capacity building block for HPE 3PAR StoreServ 8000 systems

	HPE 3PAR STORESERV 8450	HPE 3PAR STORESERV 8440	HPE 3PAR STORESERV 8400	HPE 3PAR STORESERV 8200
Controller Nodes	2, 4	2, 4	2, 4	2
Max capacity (raw)	1843 TiB	3000 TiB	2400 TiB	750 TiB
Max usable file capacity	2-256 TiB	2-256 TiB	2-256 TiB	2-128 TiB
Hard drives	NA	6-960	6-576	6-240
SSDs	6-480	6-480	6-240	6-120

Note: For further details please refer to [SPOCK](#).

Highly virtualized storage operating system

HPE 3PAR StoreServ Storage uses the same highly virtualized storage operating system across all models—including high-end, midrange, hybrid, and all-flash arrays. To help ensure performance and improve the utilization of physical resources, the HPE 3PAR Operating System employs a tri-level mapping with three layers of abstraction (including HDD, chunklets, LDs).

Fine-grained approach to virtualization

The tri-level abstraction methodology imposed by the HPE 3PAR Operating System relies on a fine-grained virtualization approach that divides each physical disk into granular allocation units referred to as chunklets, each of which can be independently assigned and dynamically reassigned to different logical disks that are used to create virtual volumes.

Multiple layers of abstraction

As shown in figure 5, the physical disk abstraction layer breaks physical drives of any size into a pool of uniform-sized, 1 GiB chunklets. The fine-grained nature of these chunklets eliminates underutilization of precious storage assets.

Complete access to every chunklet eliminates large pockets of inaccessible storage. This fine-grained structure enhances performance for all applications as well, regardless of their capacity requirements. For example, while a small application might only allocate a small amount of physical capacity, this capacity will be virtualized and striped across dozens or even hundreds of drives. With this approach, even a small application can leverage the performance resources of the entire system without provisioning excess capacity.

The first layer of abstraction employed by the OS breaks media devices into 1 GiB chunklets to enable higher utilization and avoid stranded capacity. This fine-grained virtualization unit also enables mixed RAID levels on the same physical drive, thereby eliminating dedicated RAID groups and seamlessly supporting new media technologies such as SSDs.

The second layer of abstraction takes the 1 GiB chunklets created from abstracting physical disk capacity and creates logical disks (LDs) striped across the system's physical drives and implementing specified RAID levels. Multiple chunklet RAID sets from different PDs are striped together to form an LD. All chunklets belonging to a given LD will be from the same drive type. LDs can consist of all NL, FC, or SSD chunklets. There are no mixed-type LDs, although Fast Class (Fibre Channel or SAS) LDs, may consist of both 10K and 15K drive chunklets. The association between chunklets and LDs allows LDs to be created with template properties based on RAID characteristics and the location of chunklets across the system. LDs can be tailored to meet a variety of cost, capacity, performance, and availability characteristics. In addition, the first- and second-level mappings taken together serve to parallelize work massively across physical drives and their Fibre Channel or SAS connections. LDs are divided into "regions," 128 MB of contiguous logical space from a single LD.

The third layer of abstraction maps LDs to Virtual Volumes (VVs), with all or portions of multiple underlying LDs mapped to the VV. VVs are the virtual capacity representations that are ultimately exported to hosts and applications as virtual LUNs (VLUNs) over Fibre Channel, iSCSI, or FCoE target ports. A single VV can be coherently exported through as few as two ports or as many as ports as desired (no fewer than two, one from each of two different nodes as a minimum). This layer of abstraction uses a table-based association—a mapping table with a granularity of 128 MB per region and an exception table with a granularity of 16 KB per page—as opposed to an algorithmic association. With this approach, a very small portion of a VV associated with a particular LD can be quickly and non-disruptively migrated to a different LD for performance or other policy-based reasons, whereas other architectures require migration of the entire VV. This layer of abstraction also implements many high-level features such as snapshots, caching, pre-fetching, and remote replication.




Figure 5. Virtualization with a tri-level mapping methodology that provides three layers of abstraction

One-stop allocation, the general method employed by IT users for volume administration, requires minimal planning on the part of storage administrators. By an administrator simply specifying virtual volume name, RAID level, and size, the HPE 3PAR Operating System autonomically provisions VVs at the moment that an application requires capacity. This process is also known as “just-in-time” provisioning. Contrast this to traditional architectures where the storage administrator must assign physical disks to RAID sets when the array is installed, which can be difficult or impossible to change later on and makes it difficult to respond to changing requirements.

The three-layer abstraction implemented by the HPE 3PAR Operating System can effectively utilize any underlying media type. This means that HPE 3PAR StoreServ Storage is able to make the most efficient use of SSDs through massive load balancing across all drives to enable ultra-high performance and prolong flash-based media life span.

Logical disks

There are three types of logical disks (LDs):

- User (USR) LDs provide user storage space to fully provisioned VVs.
- Snapshot data (SD) LDs provide the storage space for snapshots (or virtual copies), thinly provisioned (TPVV) and thinly deduplicated (TDVV) virtual volumes.
- Snapshot administration (SA) LDs provide the storage space for metadata used for snapshot and TPVV and TDVV administration.

As mentioned earlier, RAID functionality is implemented at the LD level, with each LD mapped to chunklets in order to implement RAID 1+0 (mirroring + striping), RAID 5+0 (RAID 5 distributed parity + striping), or RAID MP (multiple distributed parity, with striping).

The HPE 3PAR Operating System will automatically create LDs with the desired availability and size characteristics. In addition, several parameters can be used to control the layout of an LD to achieve these different characteristics:

- **Set size:** The set size of the LD is the number of drives that contain redundant data. For example, a RAID 5 LD may have a set size of 4 (3 data + 1 parity), or a RAID MP LD may have a set size of 16 (14 data + 2 parity). For a RAID 1 LD, the set size is the number of mirrors (usually 2). The chunklets used within a set are typically chosen from drives on different enclosures. This helps ensure that a failure of an entire loop (or enclosure) will not result in data becoming unavailable until the drive enclosure is repaired. It also helps ensure better peak aggregate performance because data can be accessed in parallel on different loops.
- **Step size:** The step size is the number of bytes that are stored contiguously on a single physical drive.
- **Row size:** The row size determines the level of additional striping across more drives. For example, a RAID 5 LD with a row size of 2 and set size of 4 is effectively striped across 8 drives.

- **Number of rows:** The number of rows determines the overall size of the LD given a level of striping. For example, an LD with 3 rows, with each row having 6 chunklets' worth of usable data (+2 parity), will have a usable size of 18 GiB (1 GiB/chunklet x 6 chunklets/row x 3 rows).

Note

Hewlett Packard Enterprise recommends that the storage administrator allow the array to decide the best combination of step size, row size, and row number.

LD Ownership

Every LD has an "owner" and a "backup" owner. Using the traditional layout, chunklets from any given PD are owned by a single node with the partner node as the backup owner; thus every node creates LDs from the PDs it "owns." Express Layout alters ownership for flash drives. If the set size configured requires more than 50 percent of the drives behind the node pair then the LD will be created using chunklets from PDs behind the node pair, allowing each node to create LDs larger than traditionally possible. This allows smaller flash systems to create larger set sizes, reducing RAID overheads and improving capacity efficiency.

Common provisioning groups

A common provisioning group (CPG) creates a virtual pool of LDs that allows VVs to share the CPG's resources and allocate space on demand. You can create fully provisioned VVs, thinly provisioned VVs and thinly deduplicated VVs that draw space from the CPG's logical disk pool.

CPGs enable fine-grained, shared access to pooled logical capacity. Instead of pre-dedicating logical disks to volumes, a CPG allows multiple volumes to share the buffer pool of LDs. For example, when a TPVV is running low on user space, the system automatically assigns more capacity to the TPVV by mapping new regions from LDs in the CPG to the TPVV. As a result, any large pockets of unused but allocated space are eliminated. Fully provisioned VVs cannot create user space automatically, and the system allocates a fixed amount of user space for the volume at the time it is created.

Virtual volumes

There are two kinds of VVs: "base volumes" and "snapshot volumes." A base volume can be considered to be the "original" VV and is either a fully provisioned virtual volume, a thinly provisioned virtual volume or a thinly provisioned deduplicated virtual volume. In other words, it directly maps all the user-visible data. A snapshot volume is created using HPE 3PAR Virtual Copy software. When a snapshot is first created, all of its data is mapped indirectly to the parent volume's data. When a block is written to the parent, the original block is copied from the parent to the snapshot data space and the snapshot points to this data space instead. Similarly, when a block is written in the snapshot, the data is written in the snapshot data space and the snapshot points to this data space.

VVs have three types of space:

- The **user space** represents the user-visible size of the VV (i.e., the size of the SCSI LUN seen by a host) and contains the data of the base VV.
- The **snapshot data space** is used to store modified data associated with snapshots. The granularity of snapshot data mapping is 16 KB pages.
- The **snapshot admin space** is used to save the metadata (including the exception table) for snapshots.

Each of the three space types is mapped to LDs, with all of these LDs striped across all controller nodes; thus, VVs can be striped across multiple nodes for additional load balancing and performance.

The size limit for an individual virtual volume is 16 TiB. A VV is classified by its provisioning type, which can be one of the following types:

- **Fully provisioned VV (FPVV):** Has either no snapshot space or deprecated, statically allocated snapshot space.
- **Thinly provisioned VV (TPVV):** TPVV has space for the base volume allocated from the associated CPG and snapshot space allocated from the associated snapshot CPG (if any).

On creation, 256 MB per node is allocated to a TPVV. Storage is allocated on demand in the snapshot data area as required by the host operation being performed. The snapshot admin area contains the metadata indexes that point to the user data in the SD area. Because the SA metadata needs to be accessed to locate the user data, the indexes are cached in policy memory to reduce the performance impact of the lookups.

TPVVs associated with a common CPG share the same LDs and draw space from that pool as needed, allocating space on demand in small increments for each controller node. As the volumes that draw space from the CPG require additional storage, the HPE 3PAR Operating System automatically extends existing LDs or creates new LDs until the CPG reaches the user-defined growth limit, which restricts the CPG's maximum size.

- **Thinly deduped VV (TDVV):** TDVVs behave similarly to TPVV volumes with the fundamental difference that TDVVs within the same CPG will share common pages of data. The data shared is determined via the inline deduplication mechanism described later in this paper. TDVVs are supported only on CPGs that use SSDs as a tier of storage.
- **Commonly provisioned VV (CPVV):** The space for this VV is fully provisioned from the associated CPG, and the snapshot space is allocated from the associated snapshot CPG.

VLUNs and LUN masking

VVs are only visible to a host once the VVs are exported as VLUNs.

VVs can be exported in three ways:

- To specific hosts (set of Worldwide Names or WWNs)—the VV is visible to the specified WWNs, regardless of which port(s) those WWNs appear on. This is a convenient way to export VVs to known hosts.
- To any host on a specific port—this is useful when the hosts (or the WWNs) are not known prior to exporting, or in situations where the WWN of a host cannot be trusted (host WWNs can be spoofed).
- To specific hosts on a specific port.

Metadata handling

For metadata, HPE 3PAR StoreServ implements a mechanism of fast lookup tables that store location pointers to accelerate data access. This process relies on three layer addresses translation mechanism akin to Virtual Memory lookup tables. The VV metadata also known as Snapshot Admin space, contains bitmaps indicating which pages of the shared snapshot data areas are used along with the exception tables associated with each volume.

The exception tables provide the offset into the shared data area where the shared data is located. There are three levels (L1, L2, and L3) of exception tables for each Thin Volume or snapshot. The bits of the I/O request's logical block address (LBA) are used as indexes into the L1, L2, and L3 exception tables to determine the 16 KiB page being accessed.

Metadata management is shared across all cluster resources from CPU memory (control cache).

When HPE 3PAR Adaptive Flash Cache is configured, all Snapshot Admin space associated with any TPVV (or snapshot) on the array will always be also cached to Flash Cache.

For inline deduplication, where fast lookups are extremely important, HPE 3PAR StoreServ implements an innovative mechanism to detect duplicate pages called Express Indexing. The technology uses the computed hash signature as an index to determine whether a match already exists using the three level translation mechanisms described earlier.

When new a write I/O request comes in, the Logical Block Address (LBA) is used as an index into three different page tables as per a regular TPVV. However, instead of allocating a new page, the hash signature of the incoming data page is computed by the HPE 3PAR ASIC and compared to the signatures of the TDVV data already stored in the CPG.

If a match is found then the L3 page table entry will be set to point to the existing copy of the data page. Only if no match is found, a new page is allocated.

Please further details please refer to: [HPE 3PAR Thin Technologies white paper](#)

System-wide sparing

The HPE 3PAR Operating System has a logical volume manager that handles volume abstraction at the VV layer and also handles sparing. This logical volume manager reserves a certain number of chunklets as spare chunklets depending on the sparing algorithm and system configuration. Unlike many competitive arrays that reserve dedicated spare drives that then sit idle, system-wide sparing with HPE 3PAR

StoreServ Storage means that spare chunklets are distributed across all drives. This provides additional protection and enables a balanced load that extends the SSD life span by providing even wearing. It also protects against performance degradation by enabling a “many-to-many” rebuild in the event of a failure.

Flash-optimized innovations

HPE 3PAR StoreServ Mesh-Active Architecture, system-wide striping, fine grain virtualization, advanced metadata handling and system-wide sparing are just some of the pillars of HPE 3PAR StoreServ architecture that enhance flash-based media. Flash-based media can deliver many times the performance of conventional spinning HDDs and it can do so at very low, sub-millisecond latency. However, it is important to understand that these advantages can only be realized by an architecture that has optimized its entire I/O path to be performance-centric. If the storage controllers that sit between servers and back-end flash devices can't keep up with the performance of the flash drives, they become performance bottlenecks.

To work with flash-based media in the most performance-optimized manner, the HPE 3PAR StoreServ Architecture includes features designed to handle it in a substantially different way than spinning media. It also exploits every possible opportunity to extend flash-based media life span by reducing factors that contribute to media wear. On top of that all SSDs currently available for HPE 3PAR StoreServ Storage are enterprise-grade SAS SSDs, giving you a choice between the latest planar NAND and emerging 3D NAND technologies. HPE 3PAR StoreServ Storage currently offers several NAND flash options that leverage the latest enterprise Multi-level cell (MLC) technology.

This flash-optimized architecture relies on several new and unique HPE 3PAR StoreServ Storage innovations that accelerate performance and extend flash-based media life span:

- **Thin Deduplication with Express Indexing:** The system's Thin Deduplication software feature uses a hashing engine capability built into the HPE 3PAR ASICs in combination with a unique Express Indexing feature to deduplicate data inline and with a high degree of granularity. Hardware-accelerated Thin Deduplication delivers a level of capacity efficiency that is superior to other approaches without monopolizing CPU resources and degrading performance, thereby delivering the only primary storage deduplication solution in the industry that is truly enterprise-class. ASIC-assisted, block-level deduplication takes place inline, which provides multiple benefits, including increasing capacity efficiency, protecting system performance, and extending flash media life span.
- **Adaptive Flash Cache:** This performance acceleration feature can increase overall system IOPS and reduce latency by using SSD capacity to extend the system's cache for read data coming from spinning media, reducing the number of times the array needs to read from spinning disks and resulting in a higher number of read hits coming from cache.
- **Express Writes:** Protocol optimization within the HPE 3PAR OS enables HPE 3PAR StoreServ to deliver lower CPU interrupts per I/O transaction, which results in higher IOPS and reduced latency for writes. Express Writes does not require any changes on the initiator side and all supported HBAs and Hosts will benefit of this optimization.
- **Adaptive Read and Write:** This feature matches host I/O size reads and writes to flash media at a granular level to avoid excess writes that cause unnecessary wear to flash media. Adaptive reads and writes also significantly reduce latency and enhance back-end performance to enable more applications to be consolidated.
- **Express Layout:** This unique technology born from HPE 3PAR StoreServ's three-layer virtualization technology allows HPE 3PAR controller nodes to share access to SSDs in order to drive efficiency. Replacing traditional layouts for flash, Express Layout allows each physical drive (PD) to be “owned” by both controllers at the same time allowing both nodes to use chunklets from every drive to build Logical Disks (LDs). This technology allows smaller systems to realize reduced capacity overhead.
- **Autonomic Cache Offload:** Automatically adjusts the frequency at which data is offloaded from cache to flash media based on utilization rates. This reduces cache bottlenecks by automatically changing the frequency at which data is offloaded from cache to flash media—based on utilization rate and without requiring any user intervention. This helps achieve consistently high performance levels as you scale workloads to hundreds of thousands of IOPS.
- **Multi-tenant I/O processing:** Multi-tenant I/O processing enables performance improvement for mixed workloads or virtual desktop infrastructure (VDI) deployments by breaking large I/O into smaller chunks so that small read requests don't get held up or stuck behind larger I/O requests, which also helps ensure reduced latency expected of flash-based media.
- **Adaptive Sparing:** Using patented Adaptive Sparing technology, Hewlett Packard Enterprise has collaborated with SSD suppliers to extend usable capacity per drive by up to 20 percent. This is achieved by reducing capacity typically reserved by media suppliers for wear management and then using that space more efficiently. At a system level, increasing usable drive capacity also helps spread writes more broadly to extend SSD endurance.

For complete details of how the HPE 3PAR Architecture is flash optimized both at the hardware and software layers, refer to the white paper [HPE 3PAR StoreServ Storage: optimized for flash](#).

Workload-centric Storage Personas

HPE 3PAR StoreServ Storage is expressed by workload-centric Storage Personas. A persona is **the aspect of one’s character that is presented to or perceived by others**. Storage Personas are thus comprised of data access protocols and data services for the presentation of storage to hosts and clients. Specifically, HPE 3PAR StoreServ features a Block Persona and a File Persona that are engineered into the core of the HPE 3PAR OS and system architecture, and are managed seamlessly together via the HPE 3PAR StoreServ Management Console and scriptable HPE 3PAR CLI. Through these Storage Personas, HPE 3PAR StoreServ provides truly converged block, file, and object access to simultaneously support an expanse of workloads while allowing the best storage approach to be employed for a given workload.

HPE 3PAR Block Persona

HPE 3PAR StoreServ Storage is expressed by the Block Persona in the form of block volumes to server hosts via Fibre Channel, iSCSI, and FCoE. The block data services are as mentioned throughout this white paper.

HPE 3PAR File Persona

The HPE 3PAR File Persona can be enabled on a HPE 3PAR StoreServ storage system node pair with an optional license. It requires either a 2 port 10GbE or a 4 port 1GbE NIC to be installed in the system or the on-board 1GbE RCIP port to be enabled for File Persona.

The HPE 3PAR File Persona is designed for client workloads such as home directories and user shares; content management and collaboration; data preservation and governance; and custom cloud applications by presenting file shares via SMB (CIFS) and NFS as well as object shares via the Object Access API to client devices. File data services include: User Authentication Services; capacity and user/group Quota Management; File Store Snapshots with user-driven file restore; and Antivirus Scan Services for integration with third-party antivirus software.




Figure 6. Logical view of HPE 3PAR File Persona managed objects

As illustrated in figure 6, the HPE 3PAR File Persona has four types of managed objects. File Provisioning Groups, which are instances of the HPE intellectual property Adaptive File System, control how files are stored and retrieved. Each File Provisioning Group is transparently constructed from one or multiple Virtual Volumes. Virtual File Servers (which present virtual IP addresses), participate in User Authentication Services, and can have properties for such things as user/group Quota Management. File Stores are the slice of a Virtual File Server and File Provisioning Group at which File Store Snapshots are taken, capacity Quota Management can be performed, and Antivirus Scan Services policies customized. File Shares are presented to clients via SMB, NFS, and the Object Access API, and are where share permissions are set. File Stores and File Provisioning Groups are typically only explicitly managed for advanced operations.

Multi-tenant architecture benefits

With HPE 3PAR StoreServ Storage, you can securely partition resources within a shared infrastructure in order to pool physical storage resources for lower storage costs without compromising security or performance.

The HPE 3PAR StoreServ Storage platform was built from the ground up to deliver multi-tenant capacity that supports massive consolidation with ultra-high performance. The multi-controller scalability and extreme flexibility built into HPE 3PAR StoreServ Storage makes deploying and maintaining separate storage silos to deliver different QoS levels a thing of the past. Unlike application-centric approaches to storage, one-click autonomic rebalancing on HPE 3PAR StoreServ Storage enables you to enhance QoS levels without service disruption, pre-planning, or the need to purchase separate arrays to support different service levels. To support multiple tenants and workloads, HPE 3PAR StoreServ Storage provides secure administrative segregation of users, hosts, and application data. The following sections provide insight into the architectural elements that support each of these core capabilities.

Tier-1 resiliency

HPE 3PAR StoreServ Storage is designed to support massive consolidation by supporting mixed workloads and secure administrative segregation of users, hosts, and application data. Multi-tenancy allows IT organizations to deliver higher performance levels, greater availability, and next-generation functionality to multiple user groups and applications from a single storage system.

Today's IT realities—including virtualization, cloud computing, and ITaaS—demand the ability to deliver predictable service levels in an inherently unpredictable world, and make system resiliency the single most important requirement for multi-tenancy. Traditionally, Tier-1 storage has been characterized by hardware redundancy, advanced replication capabilities, and massive scalability in capacity and host connectivity. However, in order to enable the consolidation of multiple tenants onto a single system, hardware and software fault tolerance, as well the ability to predictably prevent downtime and handle failures in a way that is non-disruptive to users and applications, become critical. The HPE 3PAR StoreServ Architecture supports multi-tenancy by allowing you to consolidate with confidence and achieve higher service levels for more users and applications with less infrastructure.

Advanced iSCSI connectivity enhancements

HPE 3PAR StoreServ Storage offers a unified solution for IP-only environments for both block (over IP with iSCSI) and file (over IP with SMB/NFS). Although historically IP environments have not been considered suitable for flash, recent iSCSI connectivity enhancements to the platform including VLAN tagging give you the flexibility to deploy flash using Ethernet while preserving sub-millisecond latencies and Tier-1 resiliency and data services.

Latest iSCSI enhancements include:

- IPv6 support
- Support for more than 8,000 iSCSI initiators per system
- Support for VLAN tagging
- Enterprise iSCSI (iSCSI over DCB/lossless Ethernet)
- Send Targets discovery support
- Persistent Ports support

iSCSI VLAN tagging is an initiative that allows iSCSI ports to be configured with multiple IP addresses and VLAN tags, within HPE 3PAR StoreServ Storage. With support for up to 500 VLAN tags and 256 initiators per port, the traffic to these ports is well streamlined on VLAN membership. This is important in an iSCSI architecture which is governed by SCSI standards. Handling the network traffic and ensuring all the security protocols are intact becomes bit of a challenge. Especially when SCSI domain consists of more than one iSCSI initiators and iSCSI target; traditional iSCSI traffic is pretty static when frequent changes are to be made to underlying network, on the go. iSCSI VLAN tagging offers network simplification and makes HPE 3PAR StoreServ iSCSI traffic efficient by smoothing SCSI flows and giving that traffic higher priority. Support for IPv6 ensures that network activity of each device can potentially be tracked, thus offering the added security boost to the existing infrastructure.

Hardware and software fault tolerance

HPE 3PAR StoreServ Storage delivers Tier-1 resiliency with an architecture designed to eliminate any single point of failure (hardware or software) in the system. To mitigate single points of failure at the hardware layer, the system is designed with redundant components, including redundant power domains. In fact, to raise the bar with the fault tolerance mechanism, HPE 3PAR StoreServ 20000 Storage system is configured with two self-encrypting boot drives that work in redundant mode.

An independent copy of HPE 3PAR OS runs on every controller node, so even in the smallest configuration, with two controller nodes, the system offers resiliency also for the software stack.

HPE 3PAR StoreServ Storage components such as storage nodes, disk- and host-facing host bus adapters (HBAs), power supplies, batteries, and disks all feature N+1 and in some cases N+2 redundancy so that any of these components can fail without system interruption. The only non-redundant component in the system is a 100 percent completely passive controller node backplane/backplanes¹ that, given its passive nature, is virtually impervious to failure.

HPE 3PAR StoreServ Storage offers up to four current load-balanced power distribution units (PDUs) per rack, which provide a minimum of two separate power feeds. The system can support up to four separate data center power feeds, providing even more power resiliency and further protection against power loss as well as brownouts. Redundant power domains help ensure that as many as two disk chassis power supplies can fail without power being lost to back-end disk devices.

Controller nodes in an HPE 3PAR StoreServ Storage system includes a local physical drive that contains a separate instance of the HPE 3PAR Operating System as well as space to save cached write data in the event of a power failure.

The controller nodes are each powered by two (1+1 redundant) power supplies and backed up by two batteries. Each battery has sufficient capacity to power the controller nodes long enough to flush all dirty data in cache memory into the local physical drive in the event of a complete power failure to the node. Although many architectures use battery backed RAM as cache (to maintain the data in cache while waiting for power to be restored) these are not suitable for extended downtimes that are usually associated with natural disasters and unforeseen catastrophes.

Another common problem with many battery-powered backup systems is that it is often impossible to ensure that a battery is charged and working. To address this problem, the HPE 3PAR StoreServ Storage controller nodes are each backed by at least two batteries. Batteries are periodically tested by slightly discharging one battery while the other remains charged and ready in case a power failure occurs while the battery test is in progress. Following a power failure HPE 3PAR Operating System keeps track of battery charge levels and limits the amount of write data that can be cached based on the ability of the batteries to power the controller nodes while they are recharging following the power failure.

The HPE 3PAR StoreServ Storage power failure protection mechanisms eliminates the need for expensive batteries to power all of the system's drive chassis, while dirty cache data is destaged to disks on the back end of the array. Note that, because all cached write data is mirrored to another controller node, a system-wide power failure would result in saving cached write data onto the internal drives of two nodes. This offers further protection following a power failure in the event a node in the cluster is damaged by the power failure. The second node containing the data can be used for recovery of the saved data. Because each node's dual power supplies can be connected to separate AC power cords, providing redundant AC power to the system can further reduce the possibility of an outage due to an AC power failure.

Advanced fault isolation and RAID protection

Advanced fault isolation and high reliability are built into the HPE 3PAR StoreServ Storage system. The drive chassis, drive magazines, and physical drives themselves all report and isolate faults. A drive failure will not result in data being unavailable.

HPE 3PAR StoreServ Storage constantly monitors drives via the controller nodes and enclosures, and isolates faults to individual drives, then "offlines" only the failed component.

¹ HPE 3PAR StoreServ 20000 systems have a single backplane. 4-Node HPE 3PAR StoreServ 8000 systems have backplanes and cables.

HPE 3PAR StoreServ Storage is capable of RAID 1+0 (dual or triple mirror or striped), RAID 5+0 (RAID 5 distributed parity, striped in an X+1 configuration where $X \geq 2$ and $X \leq 8$), or RAID MP (multiple distributed parity, and currently striped with either a 4:2; 6:2; 8:2; 10:2; 14:2 configuration). In an appropriately configured HPE 3PAR StoreServ array, all available RAID options allow HPE 3PAR StoreServ Storage to create parity sets on different drives in different drive cages with separate power domains for greater integrity protection.

Each drive enclosure is divided into two redundant I/O modules cages that plug into the drive chassis midplane. The drive chassis components—power supplies, Fibre Channel or SAS Adapters, and drives—are serviceable online and are completely hot-pluggable. Redundant power supply/fan assemblies hot-plug into the rear of the midplane. Should the drive chassis midplane fail for any reason, partner cage or cages will continue to serve data for those volumes that were configured and managed as “High Availability (HA) Cage” volumes. If the “HA Cage” configuration setting is selected at volume creation, the controller node automatically manages the RAID 1+0, RAID 5+0, or RAID MP data placement to accommodate the failure of an entire cage without affecting data access.

Controller node redundancy

The HPE 3PAR Operating System instance running on each of the controller nodes is both statefully managed and self-healing, providing protection across all cache-coherent, Mesh-Active storage controller nodes should one or more processes fail and restart. Write cache is mirrored across controllers, and the system offers RAID 1+0 (mirroring + striping), RAID 5+0 (RAID 5 distributed parity + striping), and RAID MP (multiple distributed parity with striping) across the drives behind the nodes.

In addition, controller nodes are configured in logical pairs whereby each node has a partner. The partner nodes have redundant physical connections to the subset of physical drives owned by the node pair. Within the pair, each serves as the backup node for the LDs owned by the partner node.

If a controller node were to fail, data availability would be unaffected. In the event of a node failure, that node's partner takes over the LDs for the failed node. It then immediately flushes data in the write cache on other nodes in the array that belongs on the LDs it has taken over.

Furthermore, under certain circumstances, the system is capable of withstanding a second Node failure (however rare) without affecting data availability. After the Node failover recovery process for the initial Node failure is complete, a second Controller Node from the remaining Node pairs can fail without causing system downtime.

Data integrity checking

In addition to hardware fault tolerance, all HPE 3PAR StoreServ Storage systems offer automated end-to-end error checking during the data frames' journey through the HPE 3PAR StoreServ Storage array to the disk devices to help ensure data integrity in support of Tier-1 resilience. In addition to this HPE 3PAR Thin Express Gen5 ASIC comes with the Persistent Checksum feature known as T10-DIF that ensures end-to-end data protection, from host HBA to physical drives.

Embedded Cyclical Redundancy Checking (CRC) checking includes, but is not exclusive to, the following layers within all HPE 3PAR StoreServ Storage systems:

- CRC/parity checks on all internal CPU and serial buses
- Control cache ECC checks
- Data cache ECC checks
- PCIe I2C bus CRC/parity checks
- HPE 3PAR ASIC connection CRC/parity checks
- Protocol (Fibre Channel/iSCSI/FCoE) CRC checks at the frame level (hardware accelerated via the HPE 3PAR ASICs)
- Disk devices CRC checks at the block level, occurring once the data has landed and throughout the lifecycle of the data once it's stored to disk

CRC error checking is also extended to replicate data with HPE 3PAR Remote Copy software, which help ensure that potential cascaded data issues do not occur. HPE 3PAR StoreServ Storage replication includes a link pre-integration test to verify the stability of Remote Copy replication links in advance of adding these links within the HPE 3PAR StoreServ Storage system for use with HPE 3PAR Remote Copy over an IP network (RCIP).

All drives in the HPE 3PAR StoreServ 20000 and StoreServ 8000 storage systems are formatted with 520-byte blocks in order to provide space to store a CRC Logical Block Guard, as defined by the T10 Data Integrity Feature (T10-DIF) for each block. This value is computed by the HPE 3PAR HBA before writing each block and is checked when a block is read. NL SAS does not support 520-byte blocks, so on Enterprise NL SAS drives, data blocks are logically grouped with an extra block to store the CRC values. The CRC Logical Block Guard used by the T10-DIF is automatically calculated by the host HBAs to validate data stored on drives without additional CPU overhead.

HPE 3PAR StoreServ Storage continuously runs a background “PD scrubber” process to scan all blocks of the physical drives in the system. This is done to detect any potential issues at the device block layer and trigger RAID rebuilds down to 512-byte granularity if necessary. This is particularly important when it comes to flash media because it allows the system to proactively detect and correct any low-level CRC and bit errors.

Furthermore, Self-Monitoring, Analysis and Reporting Technology (SMART) predictive failures mean that any disk device crossing certain SMART thresholds would cause the storage controller nodes to mark a drive as “predictive failure,” identifying it for replacement before it actually fails.

When inline deduplication is used, the controller node ASICs will perform a bit-by-bit comparison before any new write is marked as a duplicate. This helps ensure data integrity by introducing this critical check into the deduplication process to support mission-critical environments.

HPE 3PAR StoreServ Storage systems also issue logical error status block (LESB) alerts if a frame arriving in the storage interface has CRC errors beyond a certain threshold. This indicates that a cable or component between the host and storage device needs replacing or cleaning.

Memory fencing

HPE 3PAR StoreServ Storage is able to correct single-bit (correctable) errors and detect double-bit (uncorrectable) errors.


It achieves this by using a thread (memory patrol) that continuously scans the memory and keeps track of correctable errors at a 16 KB page granularity. If, during the scan, the thread detects uncorrectable errors, areas of the memory are fenced and put onto a “do not use” list. The system will raise a service alert when the threshold of correctable errors is reached and/or memory is fenced such that replacing the DIMM is recommended.

Persistent technologies

No one has time for downtime, which is why modern Tier-1 resiliency requires that data access and service levels be maintained during failure recovery, maintenance, and software upgrades. Tier-1 resiliency demands that failures not only be prevented, but that the system can recover quickly in the event that something goes wrong. Not only is HPE 3PAR StoreServ Storage designed to be non-disruptively scalable and upgradable, but the system also has several advanced features to prevent unnecessary downtime and to maintain availability and performance levels during planned as well as unplanned outage events. These features are collectively known as persistent technologies.

HPE 3PAR Persistent Checksum

This feature, Persistent Checksum addresses media and transmission errors that can be caused by any component in the I/O stack all the way from the server HBA through the SAN switches and into the HPE 3PAR StoreServ HBAs making the data secure all the way from the hosts right to the drives and providing additional protection above CRC transmissions for Fibre Channel alone. Persistent Checksum is server and application independent (it does require server HBAs that support the feature) and offers elaborate host OS support. When using unsupported HBAs T10-DIF tags are added and verified on the array target ports, inter-node copies and back-end HBAs. When using supported HBAs, T10-DIF tags are added by the Host HBAs and verified throughout the HPE 3PAR StoreServ Storage system, making the data secure all the way from the hosts to the drives. Where Persistent Checksum detects media or transmission errors, graceful error recovery will take place, avoiding impact on the host application.



*For a list of supported T10-PI HBA see HP (now Hewlett Packard Enterprise) SPOCK

Figure 7. HPE 3PAR Persistent Checksum

HPE 3PAR Persistent Cache

HPE 3PAR Persistent Cache is a resiliency feature built into the HPE 3PAR Operating System that allows graceful handling of an unplanned controller failure or planned maintenance of a controller node. This feature eliminates the substantial performance penalties associated with traditional modular arrays and the cache “write-through” mode they have to enter under certain conditions. HPE 3PAR StoreServ Storage can maintain high and predictable service levels even in the event of a cache or controller node failure by avoiding cache write-through mode via HPE 3PAR Persistent Cache.

Under normal operation on an HPE 3PAR StoreServ Storage system, each controller has a partner controller in which the controller pair has ownership of certain logical disks. As mentioned earlier, LDs are the second layer of abstraction in the system’s approach to virtualization of physical resources. In the event of controller node failure, HPE 3PAR Persistent Cache preserves write caching by having the surviving node from a node pair, where one node has failed, mirror writes that would have gone to the failed node to other nodes in the cluster instead. By doing this the surviving node does not have to go into write-through mode for the LDs it owns to ensure data integrity in the unlikely event it should fail too.

For example, in a quad controller configuration (where Node 0 and Node 1 form a node pair and Node 2 and Node 3 form a second node pair), each node pair might own 100 LDs with each node within the pair fulfilling the role of the primary node for 50 of those LDs. If Node 2 fails, the system will transfer ownership of its 50 LDs to Node 3, and Node 0 and Node 1 will now be the backup (and thereby the cache mirroring partner) for the 100 LDs that Node 3 is now responsible for. The mirroring of write data coming into Node 3 for those 100 LDs will be distributed across Node 0 and Node 1.

HPE 3PAR Persistent Ports

Another persistent technology, HPE 3PAR Persistent Ports allows for a non-disruptive environment (from the host multipathing point of view) where host-based multipathing software is not required to maintain server connectivity in the event of a node or link outage on any SAN fabric. This applies to firmware upgrades, node failures, and node ports that are taken offline either administratively, or as the result of a hardware failure in the SAN fabric that results in the storage array losing physical connectivity to the fabric.

From a host standpoint, connections to HPE 3PAR StoreServ Storage systems continue uninterrupted with all I/O being routed through a different port on the HPE 3PAR StoreServ Storage array. This helps you achieve an uninterrupted service level for applications running on HPE 3PAR StoreServ Storage systems.

HPE 3PAR Persistent Port functionality works for Fibre Channel, FCoE, and iSCSI transport layers and provides transparent and uninterrupted failover in response to the following events:

- HPE 3PAR OS firmware upgrade
- HPE 3PAR node maintenance or failure
- HPE 3PAR array “loss sync” to the FC fabric or iSCSI network
- Array host ports being taken offline administratively
- Port laser loss for any reason (applies to FC only)

For more information, see the [HPE 3PAR Persistent Ports white paper](#).

HPE 3PAR Remote Copy

HPE 3PAR Remote Copy software brings a rich set of features that can be used to design disaster-tolerant solutions that cost-effectively address disaster recovery challenges. HPE 3PAR Remote Copy is a uniquely easy, efficient, and flexible replication technology that allows you to protect and share data from any application.

Implemented over a native IP network² (through the built-in Gigabit Ethernet interface available on all nodes) and native Fibre Channel, users may flexibly choose one of four different data replication modes—**Asynchronous Streaming** or Asynchronous Periodic (for asynchronous replication), **Synchronous**, or **Synchronous Long Distance**—to design a solution that meets their solution requirements for recovery point objectives (RPOs) and recovery time objectives (RTOs).

With all four of these modes, HPE 3PAR Remote Copy software allows you to mirror data between HPE 3PAR StoreServ Storage systems eliminating the incompatibilities and complexities associated with trying to mirror data between traditional vendors’ midrange and enterprise array technologies. Source and target volumes may also be flexibly and uniquely configured to meet users’ needs (e.g., different RAID levels, the use of FPVVs vs. TPVVs, or different drive types). HPE 3PAR Remote Copy is “thin aware” in that it is able to replicate both thin and thick volumes by using TPVV target volumes to provide the same cost savings associated with thin-provisioned source volumes created with HPE 3PAR Thin Provisioning software.

Synchronous mode solutions are suitable for data centers within metro distances and have a recovery point objective of zero seconds. When combined with HPE 3PAR Peer Persistence these solutions also enables recovery time objective that can be measured in seconds.

For asynchronous replication solutions, network bandwidth is efficiently utilized with **Asynchronous Periodic** mode. Changed data within an HPE 3PAR Remote Copy Volume Group is transferred only once—no matter how many times it may have changed—between synchronization intervals. Additionally, efficiencies in the initial copy creation of the target volumes that do not require replication of “zero” data across the replication network (regardless of target volume type, thick or thin) result in a faster initial synchronization and better network utilization. With support for replication over long distances, HPE 3PAR Remote Copy **Asynchronous Streaming** assures data protection against disasters by

² Asynchronous Streaming is not support on the built-in GiGe interface at the release of HPE 3PAR OS version 3.2.2.

providing an RPO in seconds, whereas HPE 3PAR Remote Copy Asynchronous Periodic mode ensures RPO in minutes over extended distances. Asynchronous Streaming is suitable for customers that want near zero RPO without the additional host latency associated with the synchronous mode of replication.

Synchronous Long Distance mode delivers a disaster recovery solution across long distances with a potential for zero data loss RPO and an RTO of minutes. This is achieved with a replication configuration that uses three sites (**3DC**) to simultaneously replicate a virtual volume from the primary array in synchronous mode to a target HPE 3PAR StoreServ Storage array located at a synchronous site (within a metropolitan area) and in asynchronous periodic mode to an HPE 3PAR StoreServ Storage array located at an asynchronous site (across a long distance). In addition to the HPE 3PAR Remote Copy connections from the primary array to the two backup arrays, a passive asynchronous periodic link is configured from the synchronous target array to the disaster recovery target array. Under the Synchronous Long Distance mode algorithm, the synchronous target array intelligently tracks the delta set of I/Os that have been acknowledged to the host but which have not yet been replicated to the asynchronous target array. In the event that a disaster takes the primary storage array down, the user has the flexibility to recover either from the synchronous target array or the asynchronous target array.

For more information, see the [HPE 3PAR Remote Copy white paper](#).

HPE 3PAR Virtual Domains

HPE 3PAR Virtual Domains software is an extension of HPE 3PAR virtualization technologies that delivers secure segregation of virtual private arrays (VPAs) for different user groups, departments, and applications while preserving the benefits delivered by the massive parallelism architected into the HPE 3PAR StoreServ platform. It supports HPE 3PAR's Multi-Tenancy Paradigm.

By providing secure administrative segregation of users and hosts within a consolidated, massively parallel HPE 3PAR StoreServ Storage system, HPE 3PAR Virtual Domains allows individual user groups and applications to affordably achieve greater storage service levels (performance, availability, and functionality) than previously possible.

HPE 3PAR Virtual Domains is completely virtual and represents no physical reservation of resources. To use HPE 3PAR Virtual Domains, a master administrator first creates a virtual domain, and then assigns logically defined entities to it. These include one or more host definitions based on Worldwide Name (WWN) groupings, one or more provisioning policies (RAID and disk type), and one or more system administrators (who are also granted role-based privileges by the master administrator).

Depending on the level of access, users can create, export, and copy standard volumes or thin-provisioned volumes. HPE 3PAR Virtual Domains is ideal for enterprises or service providers looking to leverage the benefits of consolidation and deploy a purpose-built infrastructure for their private or public cloud.

Data encryption

Data is perhaps the most important asset for organizations in today's digital age. Companies are looking to protect data against theft and misuse while meeting compliance requirements. The HPE 3PAR StoreServ Storage complies with the standards set forth by the National Institute of Standards and Technology (NIST) and FIPS 140-2 (Federal Information Processing Standard) and features Data-at-Rest (DAR) encryption that helps protect valuable data through self-encrypting drive (SED) technology. SED drives are HDDs and SSDs with a circuit (ASIC) built into the drive's controller chipset that automatically encrypts and decrypts all data being written to and read from the media.

HPE 3PAR StoreServ Storage supports Full Disk Encryption (FDE) based on the Advanced Encryption Standard (AES) 256 industry standard. The encryption is part of a hash code that is stored internally on physical media. All encryption and decryption is handled at the drive level and needs no other external mechanism.

Authentication keys are set by the user and can be changed at any time. The Local Key Manager (LKM) included with the HPE 3PAR StoreServ Storage encryption license is used to manage all drive encryption keys within the array and provides a simple management interface. In the event of a drive failure or the theft of a drive, a proper key sequence needs to be entered to gain access to the data stored within the drive. When an SED drive is no longer powered on, the drive goes into a locked state and requires an authentication key to unlock the drive when power is restored. Without the key, access to the data on the SED is not possible.

Hewlett Packard Enterprise also offers the enhanced encryption support on the HPE 3PAR StoreServ Storage systems by offering FIPS 140-2 compliant SED drives that provides the ability to use an external Enterprise Secure Key Manager (ESKM). ESKM is deployed whenever customers use encrypted storage or communication methods to protect their sensitive information. Herein, they store and serve keys to unlock the data stored on FIPS 140-2 compliant drives within the HPE 3PAR StoreServ Storage systems with strong access controls and security.

FIPS 140-2 compliance provides the customer the satisfaction of knowing their data is securely stored on the HPE 3PAR StoreServ array. Key Management on the array with either LKM or ESKM coupled with FIPS drives, offers customers a safe environment in which to securely store their data.

For more information, see the [HPE 3PAR StoreServ Data-At-Rest Encryption white paper](#).

Maintaining high and predictable performance levels

The ability of HPE 3PAR StoreServ Storage to maintain high and predictable performance in multi-tenant environments is made possible through architectural innovations that utilize all available array hardware resources at all times thereby eliminating resource contention, support mixed workloads, and enhance caching algorithms to accelerate performance and reduce latency.

Load balancing

Purpose-built for the Enterprise as well as virtual and cloud data centers, the HPE 3PAR StoreServ Architecture is unlike legacy controller architectures in that the Mesh-Active system design allows each volume to be active on any controller in the system via the high-speed, full-mesh interconnection that joins multiple controller nodes to form a cache-coherent Active/Active cluster. As a result, the system delivers symmetrical load balancing and utilization of all controllers with seamless performance scalability by adding more controllers and disk drives to the system.

Mixed-workload support

Unlike legacy architectures that process I/O commands and move data using the same processor complex, the HPE 3PAR StoreServ Storage controller node design separates the processing of SCSI control commands from data movement. This allows transaction-intensive and throughput-intensive workloads to run on the same storage resources without contention, thereby supporting massive consolidation and multi-tenancy. This means that, for example, the system can easily handle an OLTP application and an extremely bandwidth-consuming data warehousing application concurrently with ease.

This capability is made possible by the HPE 3PAR ASIC, which offloads data processing from the control processor, where metadata is processed. By pathing and processing data and metadata separately, transaction-intensive workloads are not held up behind throughput-intensive workloads. As a result, the HPE 3PAR StoreServ Storage platform, as compared to the architectures of traditional storage vendors—including many of today’s all-flash arrays—delivers excellent performance consistently, even in mixed-workload scenarios. Figure 8 illustrates the benefits of mixed-workload support.




Figure 8. HPE 3PAR StoreServ Storage with mixed-workload support

Control operations are handled as follows:

- With the HPE 3PAR StoreServ 20800 Storage system, control operations are processed by up to 16 high-performance Intel Six-Core processors.
- With the HPE 3PAR StoreServ 20840 Storage system, control operations are processed by up to 16 high-performance Intel Eight-Core processors.
- With the HPE 3PAR StoreServ 20850 Storage system, control operations are processed by up to 16 high-performance Intel Eight-Core processors.
- With the HPE 3PAR StoreServ 20450 Storage system, control operations are processed by up to 8 high-performance Intel Eight-Core processors.
- With the HPE 3PAR StoreServ 8450 and 8440 Storage systems, control operations are handled by up to four Intel 10-core processors.
- In the case of the HPE 3PAR StoreServ 8400 Storage system, control operations are handled by up to four Intel Hexa-core processors.
- For the HPE 3PAR StoreServ 8200 Storage system, control operations are handled by up to two Intel Hexa-core processors.

Data movement is handled as follows:

- For the HPE 3PAR StoreServ 20000 Storage system, all data movement is handled by the specially designed HPE 3PAR ASICs (two per controller node).
- For the HPE 3PAR StoreServ 8000 Storage system, all data movement is handled by the HPE 3PAR ASICs (one per controller node).

Storage quality of service

Quality of service (QoS) is an essential component for delivering modern, highly scalable multi-tenant storage architectures. The use of QoS moves advanced storage systems away from the legacy approach of delivering I/O requests with “best effort” in mind and tackles the problem of “noisy neighbors” by delivering predictable tiered service levels and managing “burst I/O” regardless of other users in a shared system. Mature QoS solutions meet the requirements of controlling service metrics such as throughput, bandwidth, and latency without requiring the system administrator to manually balance physical resources. These capabilities eliminate the last barrier to consolidation by delivering assured QoS levels without having to physically partition resources or maintain discreet storage silos.

HPE 3PAR Priority Optimization software enables service levels for applications and workloads as business requirements dictate, enabling administrators to provision storage performance in a manner similar to provisioning storage capacity. This allows the creation of differing service levels to protect mission-critical applications in enterprise environments by assigning a minimum goal for I/O per second and bandwidth, and by assigning a latency goal so that performance for a specific tenant or application is assured. It is also possible to assign maximum performance limits on workloads with lower service-level requirements to make sure that high-priority applications receive the resources they need to meet service levels.

Priority Optimization feature and industries leading “Latency goal” feature that enables the storage administrator to set SLAs as low as 500 μ s for volumes residing on SSD storage. It also makes it possible to configure service-level objectives in terms of KB/s and I/O bandwidth on a virtual volume set (VVset) or between different virtual domains.³ All host I/Os on the VVset are monitored and measured against the defined service-level objective. HPE 3PAR Priority Optimization control is implemented within the HPE 3PAR StoreServ Storage system and can be modified in real time. No host agents are required, and physical partitioning of resources within the storage array is not necessary.

For more information on HPE 3PAR Priority Optimization, please see the [HPE 3PAR Priority Optimization white paper](#).

Adaptive Flash Cache

HPE 3PAR Adaptive Flash Cache is built-in functionality of the HPE 3PAR StoreServ that allows SSDs to act as Level-2 read cache holding random read data for spinning media that has aged out of DRAM read cache. Adaptive flash cache reduces application response time for read intensive I/O workloads and can improve write throughput in mixed-workload environment. Adaptive Flash Cache effectively increases the

³ A VVset may contain a single volume or multiple volumes. A virtual volume may also belong to multiple virtual volume sets allowing users to create hierarchical rules.


amount of random read data cached on high-speed media on a node. By doing this, Adaptive Flash Cache can increase the overall random read IOPS an array can deliver by “unloading” the back end of the array potentially resulting in increased write throughput on the array. Using SSDs as Level-2 read cache to hold random read data that has aged out of DRAM cache is a cost-effective way of keeping more random read data on very fast media to improve overall system performance.

For more information on HPE 3PAR Adaptive Flash Cache, please see the [HPE 3PAR Adaptive Flash Cache white paper](#).

Express Writes

HPE 3PAR Express Writes represent a series of optimizations aimed to improve host write latency. This is achieved via the HPE 3PAR OS by optimizing SCSI transactions and thus improving the number of interrupts per transaction, which results in improved array CPU utilization and lower host write latency. Express Writes technology enables HPE 3PAR StoreServ array to achieve synchronous write latency below 200 microseconds. Depending on the workload, hosts may see an overall decrease in write latency of up to 20 percent, which may result in an increase in IOPS and throughput.

The following picture is from a system with express writes disabled and then enabled (after the line). With Express Writes disabled 90 percent of write I/O was committed in 0.250 microseconds or more and 10 percent was committed in 0.125 microseconds or less. With Express Writes Enabled 50 percent of writer I/O is committed in 0.125 microseconds or less.



Performance benefits of system-wide striping

In a traditional storage array, small volumes either suffer from poor performance by using few drives or waste expensive resources by using more drives than required for capacity in order to obtain sufficient performance. On HPE 3PAR StoreServ Storage systems, even modest-sized volumes will be widely striped using chunklets spread over multiple drives of the same type. Wide striping provides the full performance capabilities of the array (nodes, CPUs, buses, cache, disk drives) to small volumes without provisioning excess capacity and without creating hotspots on a subset of physical drives.

Physical drives can hold a mix of RAID levels because RAID groups with HPE 3PAR StoreServ Storage are constructed from chunklets rather than from whole drives. Different chunklets on a physical drive can be used for volumes with different RAID levels. On a traditional array, a storage administrator might be forced to use RAID 1 for an archival volume in order to use space that is available on a RAID 1 disk even though RAID 5 would deliver adequate performance and better space utilization. The chunklet-based approach employed by the HPE 3PAR Operating System allows all RAID levels to coexist on the same physical drives, using the better RAID level for each volume. Additional details about striping are provided in the [“Highly virtualized storage operating system”](#) section.

Bandwidth and communication

The ASICs within each HPE 3PAR StoreServ 20800, 20840, 20850, and 20450 Storage controller node serve as the high-performance engines that move data between three I/O buses, a four memory-bank data cache, and seven high-speed links to the other controller nodes over the full-mesh backplane. These ASICs perform RAID parity calculations and inline zero-detection to support the system's data compaction technologies. CRC Logical Block Guard used by T10-DIF is automatically calculated by the ASIC's to validate data stored on drives with no additional CPU overhead. An HPE 3PAR StoreServ 20800 Storage system with eight controller nodes has 16 ASICs totaling 224 GiB/s of peak interconnect bandwidth.

The single ASIC within each HPE 3PAR StoreServ 8450, 8400, and 8200 Storage controller node serves as the high-performance engine that moves data between three I/O buses, a dual memory-bank data cache, and three high-speed links to the other controller nodes over the full-mesh interconnect. As with the HPE 3PAR StoreServ 20000 Storage series, the ASICs for the HPE 3PAR StoreServ 8000 Storage series models perform parity RAID calculations and inline zero-detection; CRC Logical Block Guard used by the T10-DIF is automatically calculated by the ASIC's to validate data stored on drives with no additional CPU overhead. An HPE 3PAR StoreServ 8450 Storage system with four nodes has 4 ASICs totaling 48 GiB/s of peak interconnect bandwidth.

Data transfer paths

Figure 9 shows an overview of data transfers in an HPE 3PAR StoreServ Storage system with two simple examples: a write operation from a host system to a RAID 1 volume (lines labeled W1 through W4), and a read operation (Gray lines labeled R1 and R2). Only the data transfer operations are shown, not the control transfers.

The write operation consists of:

- **W1:** Host writes data to cache memory on a controller node
- **W2:** The write data is automatically mirrored to another node across the high-speed backplane link so that the write data is not lost even if the first node experiences a failure; only after this cache mirror operation is completed is the host's write operation acknowledged
- **W3** and **W4:** The write data is written to two separate drives (D1 and D1'), forming the RAID 1 set

In step W2, the write data is mirrored to one of the nodes that owns the drives to which data will be written (in this example, D1 and D1'). If the host's write (W1) is to one of these nodes, then the data will be mirrored to that node's partner.

HPE 3PAR Persistent Cache allows a node to mirror the write data to a node that does not have direct access to drives D1 and D1' in the event of a partner node failure.

The read operation consists of:

- **R1:** Data is read from drive D3 into cache memory
- **R2:** Data is transferred from cache memory to the host




Figure 9. Data transfer paths

I/O bus bandwidth is a valuable resource in the controller nodes, and is often a significant bottleneck in traditional arrays. As the data transfer example in figure 9 illustrates, I/O bus bandwidth is used only for data transfers between the host-to-controller node and controller node-to-drive transfers. Transfers between the controller nodes do not consume I/O bus bandwidth.

Processor memory bandwidth is another significant bottleneck in traditional architectures, and is also a valuable resource in the controller nodes. Unique to the HPE 3PAR StoreServ Storage system, controller node data transfers do not consume any of this bandwidth. This frees the processors to perform their control functions far more effectively. All RAID parity calculations are performed by the ASICs and do not consume processor or processor memory bandwidth.

Sharing and offloading of cached data

Because much of the underlying data associated with snapshot volumes is physically from the base VVs, data that is cached for the base VV can often be used to satisfy read accesses for a snapshot of that base VV.

In the event that two or more drives that underlay a RAID set become temporarily unavailable (or three or more drives for RAID MP volumes)—for example, if all cables to those drives are accidentally disconnected—the HPE 3PAR Operating System automatically moves any “pinned” writes in cache to dedicated Preserved Data LDs. This helps ensure that all host-acknowledged data in cache is preserved so that it can be properly restored once the destination drives come back online without compromising cache performance or capacity with respect to any other data by keeping cache tied up.

On flash-based systems, autonomic cache offload mitigates cache bottlenecks by automatically changing the frequency at which data is offloaded from cache to flash media. This helps ensure high performance levels consistently as workloads are scaled to hundreds of thousands of IOPS.

Pre-fetching

The HPE 3PAR Operating System keeps track of read streams for VVs so that it can improve performance by “pre-fetching” data from drives ahead of sequential read patterns. In fact, each VV can detect up to five interleaved sequential read streams and generate pre-fetches for each of them. Simpler pre-fetch algorithms that keep track of only a single read stream cannot recognize the access pattern consisting of multiple interleaved sequential streams.

Pre-fetching improves sequential read performance in two ways:

- The response time seen by the host is reduced.
- The drives can be accessed using larger block sizes than the host uses, resulting in more efficient operations.

Write caching

Writes to VVs are cached in a controller node, mirrored in the cache of another controller node, and then acknowledged to the host. The host, therefore, sees an effective response time that is much shorter than would be the case if a write were actually performed to the drives before being acknowledged. This is possible because the mirroring and power failure handling help ensure the integrity of cached write data.

In addition to dramatically reducing the host write response time, write caching can often benefit back-end drive performance by:

- Merging multiple writes to the same blocks so that many drive writes are eliminated
- Merging multiple small writes into single larger drive writes so that the operation is more efficient
- Merging multiple small writes to a RAID 5 or RAID MP LD into full-stripe writes so that it is not necessary to read the old data for the stripe from the drives
- Delaying the write operation so that it can be scheduled at a more suitable time

Additional benefit to write caching is provided via Adaptive Flash Cache; while the technology is aimed to accelerate reads, the fewer random read request to the HDD tier translate to more pages available for writes.

Fast RAID 5

The architectural design of the HPE 3PAR StoreServ Storage systems enables RAID 5 redundancy with performance levels that are on par with RAID 1 mirroring. This implementation combines the HPE 3PAR ASIC, a large, memory cache, and wide striping for reducing spindle contention to offer performance that approaches that of RAID 1, thus reducing the performance impact typical of RAID 5 on legacy storage architectures.

For certain workloads, Fast RAID 5 can provide higher performance than RAID 1. The write-back cache in HPE 3PAR StoreServ Storage systems allows sequential writes (as generated by transaction journals, logs, and similar performance-sensitive workloads) to be collected until a full parity group can be written, reducing disk I/O traffic and possible back-end bottlenecks. Given its layout algorithm, Fast RAID 5 is appropriate for volumes that are dominated by read activity. HPE 3PAR StoreServ Storage systems allow selection of the number of data blocks per parity block (N+1) to suit different needs. For RAID 5, 3+1 is the default, but any value from 2+1 to 8+1 can be selected. Higher values of N result in higher storage efficiency, but can reduce the chances for full-stripe writes. Hewlett Packard Enterprise customers using HPE 3PAR StoreServ Storage arrays typically choose HPE 3PAR Fast RAID 5 for most or all volumes, as Fast RAID 5 provides greater storage efficiency with minimal performance degradation compared to RAID 1.

HPE 3PAR RAID MP (Fast RAID 6)

Exponential growth in HDD capacity without commensurate improvements in reliability or performance results in greater risk of data loss. For example, consider the 600 GiB SAS disks and 6 TiB nearline SAS disks available on HPE 3PAR StoreServ Storage systems. The capacity difference alone implies that reconstruction of a failed disk upon replacement can be expected to take more than six times as long with the 6 TiB disk. The nearline disks are slower, too, which further increases the mean time to repair (MTTR) relative to smaller, faster SSD or SAS disks. A longer MTTR creates a larger window during which a second disk failure could cause data loss when using RAID 1 or RAID 5. Fast RAID 6 was created to address this problem. Like Fast RAID 5, Fast RAID 6 uses distributed parity, but it stores two different parity values in a manner that allows the data to be reconstructed, even in the event of two drive failures.

HPE 3PAR RAID MP (multiple, distributed parity) initially supports dual parity and is capable of supporting higher parity levels in the future.

Environments such as highly consolidated virtual host environments tend to have unusually high data protection requirements due to the large number of users that could be affected by data loss, and so demand the highest level of data protection.

High I/O loads make RAID 6 problematic on traditional arrays; the implementation of RAID MP on HPE 3PAR StoreServ Storage arrays is the only choice that provides the extra level of data protection without compromising I/O performance.

Data compaction

Reducing capacity requirements by mitigating overprovisioning; using enterprise-class data compaction technologies; and enabling fast, simple, automated space reclamation are essential to the industry-leading efficiency of HPE 3PAR StoreServ Storage systems. Compaction technologies such as thin provisioning, thin deduplication, and thin reclamation offer efficiency benefits for primary storage that can significantly reduce both capital and operational costs with spinning media and SSDs.

Thin technologies can vary widely in how they are implemented, and this can greatly impact the ability to reduce capacity requirements and extend SSD life span without forcing performance trade-offs. Not only is HPE 3PAR StoreServ Storage viewed as the industry's thin technology leader, but third-party testing and competitive analysis also confirm that HPE 3PAR StoreServ Storage offers the most comprehensive and efficient thin technologies among the major enterprise storage platforms.⁴

In addition to efficiency enhancements related to performance such as system-wide striping and Fast RAID implementations discussed in previous sections, HPE 3PAR StoreServ Storage offers the most comprehensive set of thin technologies available to drive up resource utilization while protecting array performance. The following sections describe various data compaction and underlying technologies. For an expanded discussion, refer to the brochure "[Thin Deduplication: HPE 3PAR StoreServ Storage with Thin Technologies for data compaction](#)" and the [HPE 3PAR Thin Technologies white paper](#).

HPE 3PAR ASIC for inline data reduction

At the heart of every HPE 3PAR StoreServ node there is an HPE 3PAR ASIC that features an efficient, silicon-based zero-detection and hashing mechanism. This unique hardware capability gives HPE 3PAR StoreServ Storage the power to perform inline deduplication and remove allocated but unused space inline and non-disruptively.

For thinly allocated volumes (and thin dedupe volumes) zero-detect capability can recognize an incoming write request of 16 KiB of zeros and either not allocate space for the zero block or free the space if it was already allocated for that block. All this happens in cache, and therefore no zeroes are written to the back end. When a read request comes in for a block that is unallocated, the HPE 3PAR StoreServ will immediately return zeros back to the host. Built-in zero detection can be controlled per Thin Volume, and it is enabled by default.

Many other storage arrays do not detect blocks of zeroes on write. Instead, the zeros are written to disk and a scrubbing process later detects these zeroed blocks and discards them. With this approach, the zeroed blocks consume space until they're scrubbed, and therefore they may not be available for use by other volumes when needed. Also, there is increased load placed on the storage as the scrubbing process examines the block contents on the physical storage.

The HPE 3PAR ASIC is also the only solution in the industry with a built-in, silicon-based hash calculation engine. With HPE 3PAR Thin Deduplication software, the process of calculating the hash signatures for incoming data and verifying reads are offloaded to the HPE 3PAR ASICs, freeing up processor cycles to deliver advanced data services and service I/O requests. This hardware-assisted approach, together with a unique Express Indexing feature, enables extremely efficient, extremely granular block-level inline deduplication that carries multiple benefits, including increased capacity efficiency, flash performance protection, and flash media life span extension. Unlike other approaches, HPE 3PAR Thin Deduplication software performs collision detection on all data before marking it as duplicated, which is essential to ensuring data integrity for mission-critical environments.

⁴ HPE Thin Technologies: A Competitive Comparison, Edison Group, September 2012, <http://www8.hp.com/h20195/v2/GetPDF.aspx/4AA5-3223ENW.pdf>.

HPE 3PAR Thin Technologies highlights

- HPE 3PAR Thin Technologies are completely automated.
- HPE 3PAR Operating System (OS) software uses a reservationless, dedicate-on-write approach to thin volumes and virtual copies that draws and configures capacity in fine-grained increments from a single free space reservoir without pre-dedication of any kind.
- HPE 3PAR Thin Technologies use an allocation unit size of just 16 KiB, so you don't have to worry about small writes consuming megabytes or even gigabytes of capacity.
- HPE 3PAR StoreServ is a storage platform built from the ground up to support thin technologies by reducing the diminished performance and functional limitations that plague bolt-on thin solutions.

HPE 3PAR Thin Deduplication and Thin Clones

HPE 3PAR Thin Deduplication software with Express Indexing relies on the HPE 3PAR ASICs to generate and assign a hash key to each unique incoming write request. Express Indexing, a mechanism that accelerates data signature comparison, is used for ultrafast detection of duplicate write requests in order to prevent duplicate data from being written.

When a new I/O request comes in, the Express Indexing feature performs instant lookups using metadata tables in order to compare the hash keys of the incoming request to signatures of data already stored in the array. When a match is found, Express Indexing flags the duplicate request and prevents it from being written to the back end. Instead, a pointer is added to the metadata table to reference the existing data blocks. To help ensure data integrity, HPE 3PAR Thin Deduplication software relies on the controller node ASICs to perform a bit-to-bit comparison checking for data collisions before any new write update is marked as a duplicate.

With this approach, the CPU-intensive jobs of calculating CRC's for incoming data and verifying reads are offloaded to the ASIC, freeing up processor cycles to deliver advanced data services and service I/O requests. Hardware-accelerated thin deduplication delivers a level of capacity efficiency that is superior to other approaches without monopolizing CPU resources and degrading performance, thereby delivering the only primary storage deduplication solution in the industry that is truly enterprise-class. ASIC-assisted, block-level deduplication takes place inline, which provides multiple benefits, including increasing capacity efficiency, protecting system performance, and extending flash media life span.

Without the purpose-built HPE 3PAR ASICs and HPE 3PAR Express Indexing, other storage architectures lack the processing power to simultaneously drive ultrafast inline deduplication and the high performance levels demanded by flash-based media. HPE 3PAR Thin Deduplication software enables these functions to take place without contention, without sacrificing performance, and while concurrently delivering advanced data services such as replication, federated data mobility, and quality of service-level enforcements.

An extension of HPE 3PAR Thin Deduplication for server virtualization environments, HPE 3PAR Thin Clones software enables the creation of non-duplicative VM clones with Microsoft® Hyper-V and VMware® ESXi. These VM clones are created instantly by leveraging copy offload for VMware vStorage API for Array Integration (VAAI) and Microsoft Offloaded Data Transfer (ODX) technology without increasing capacity consumption on the HPE 3PAR StoreServ Storage system. HPE 3PAR Thin Clones software leverages HPE 3PAR Thin Deduplication to update the metadata table without copying data, relying on inline deduplication technology to reduce capacity footprint as new write requests come in.

HPE 3PAR Thin Provisioning

Since its introduction in 2002, HPE 3PAR Thin Provisioning software has been widely considered the gold standard in thin provisioning. This thin provisioning solution leverages the system's dedicate-on-write capabilities to make storage more efficient and more compact, allowing customers to purchase only the disk capacity they actually need and only as they actually need it.

HPE 3PAR Thin Persistence

HPE 3PAR Thin Persistence software is an optional feature that keeps TPVVs and read/write snapshots of TPVVs small by detecting pages of zeros during data transfers and not allocating space for the zeros. This feature works in real time and analyzes the data before it is written to the source TPVV or read/write snapshot of the TPVV. Freed blocks of 16 KB of contiguous space are returned to the source volume, and freed blocks of 128 MB of contiguous space are returned to the CPG for use by other volumes.

HPE 3PAR Thin Conversion

With HPE 3PAR Thin Conversion software, a technology refresh does not require terabyte-for-terabyte replacement, but instead offers the opportunity to eliminate a significant amount of legacy capacity through fat-to-thin conversion made possible by the HPE 3PAR ASIC. In fact, the HPE Get Thinner Guarantee program stands behind the ability of new HPE 3PAR StoreServ Storage customers to reduce storage capacity requirements by 75 percent by deploying any model HPE 3PAR StoreServ Storage system and using HPE 3PAR Data Compaction Technologies to convert traditional volumes on legacy storage to TPVVs or TDVVs on the new system.⁵

HPE 3PAR Thin Conversion software makes this possible by leveraging the zero-detection and inline deduplication capabilities within the HPE 3PAR ASIC and a unique virtualization mapping engine for space reclamation that powers the simple and rapid migration of inefficient “fat” volumes on legacy arrays to more efficient, higher-utilization “thin” volumes on the HPE 3PAR StoreServ Storage array. Virtual volumes with large amounts of allocated but unused space are migrated to TPVVs or TDVVs that are much smaller than the original volumes. During the migration process, allocated but unused space is discarded and the result is a TPVV that uses less space than the original volume. If TDVVs volumes are chosen all duplicate data will be detected during the migration and only one copy of the data will be written to the HPE 3PAR StoreServ.

HPE 3PAR Thin Copy Reclamation

An industry first, HPE 3PAR Thin Copy Reclamation software keeps storage as lean and efficient as possible by reclaiming the unused space resulting from deleted virtual copy snapshots. This solution builds on a virtualization mapping engine for space reclamation called HPE 3PAR Thin Engine, which is included as part of the HPE 3PAR Operating System.

HPE 3PAR Thin Copy Reclamation software is an optional feature that reclaims space when snapshots are deleted from a system. As snapshots are deleted, the snapshot space is reclaimed from a TPVV or fully provisioned VV and returned to the CPG for reuse by other volumes. Deleted snapshot space can be reclaimed from virtual copies, physical copies, or remote copy volumes.

HPE 3PAR Virtual Copy

HPE 3PAR Virtual Copies are HPE 3PAR snapshot implementation used to provide a point-in-time virtual copy of a VV and ensure that the original data can always be obtained should a problem occur when updating the data on a VV. HPE 3PAR Virtual Copies implements both an efficient variant of copy-on-write (COW) and redirect-on-write (ROW) mechanisms. For COW, the 3PAR OS uses a delayed copy-on-write (DCOW) that eliminates any performance impact to host I/O, DCOW is used for snapshots of Fully Provisioned and Thinly Provisioned volumes. HPE 3PAR StoreServ delayed COW, relegates the reading of the original data, updating of the base volume with the new data, and the copy of the original data to a background process, after the write update has been acknowledged to the host.

For Thinly Deduped Virtual Volumes (TDVVs), the HPE 3PAR OS relies on the redirect-on-write to create snapshots and store the original data. This process allow the HPE 3PAR to use a single write I/O at most to store the host I/O to an existing block of data to the base TDVV. Furthermore, the in-line dedupe engine will check if the new write is deduplicable through the fast and efficient metadata lookup. In case of a match against any data block already in the dedupe store, then only metadata update will take place and I/O are sent to the backend.

HPE 3PAR Virtual Copies are always thin, non-duplicative (only one copy of the change blocks is kept and then referenced) and reservationless, they can also have variable quality of service as the user can specify a different CPG for snapshot data; for example user data (SD) on SSD and snapshots data (SD2) on a more cost efficient SAS tier. Snapshot of Thinly Deduped Virtual Volumes (TDVVs) that reside on an SSD tier will deduplicate snapshot data.

⁵ Subject to qualification and compliance with the Get Thinner Guarantee Program Terms and Conditions, which will be provided by your Hewlett Packard Enterprise Sales or Channel Partner representative. More information is available at [HPE 3PAR Get Thinner Guarantee Brochure](#).

Thanks to HPE 3PAR StoreServ efficient metadata handling, users can configure thousands of read-only and read-write snapshots and enables flexible management as promotion from any snapshot without destroying other snapshots. More information is available in the [HPE 3PAR Virtual Copy white paper](#).

Autonomic storage management

The HPE 3PAR Operating System helps simplify, automate, and expedite storage management by handling provisioning, tiering, and change management autonomically and intelligently, at a subsystem level, and without administrator intervention.

The system's user interfaces have been developed to offer autonomic administration, which means that the interfaces allow an administrator to create and manage physical and logical resources without requiring any overt action. Provisioning does not require any pre-planning, yet the system constructs volumes intelligently based on available resources, unlike manual provisioning approaches that require planning and the manual addition of capacity to intermediary pools.

Ease of use

The HPE 3PAR Operating System reduces training and administration efforts through the simple, point-and-click [HPE 3PAR StoreServ Management Console](#) and the scriptable HPE 3PAR Command Line Interface (CLI). Both management options provide uncommonly rich instrumentation of all physical and logical objects for one or more storage systems, thus eliminating the need for the extra tools and consulting often required for diagnosis and troubleshooting.

Open administration support is provided via SNMP, Storage Management Initiative Specification (SMI-S), and Web Service API.

Provisioning is managed intelligently and autonomically. Massively parallel and fine-grained striping of data across internal resources assures high and predictable service levels for all workload types. Service conditions remain high and predictable as system use grows or in the event of a component failure, while traditional storage planning, change management, and array-specific professional services are eliminated.

Managing performance is paramount to ensuring critical applications have the resources needed at the times of peak demands. These resources rely on a tool or tools to report metrics in a timely and accurate fashion. The monitoring of performance allows the storage administrator the flexibility to view real time data as well as performance trends by using historical data stored within a database. The tool which is used with HPE 3PAR StoreServ arrays is called System Reporter.

HPE 3PAR System Reporter software is a flexible, intuitive, performance- and capacity-management tool that aggregates fine-grained performance and capacity usage data for HPE 3PAR StoreServ Storage, regardless of location. HPE 3PAR System Reporter simplifies performance monitoring and helps create chargeback reports and resource planning.

[HPE 3PAR System Reporter](#) software is installed and runs by default on HPE 3PAR StoreServ systems, and automatically starts collecting performance data from the local system, there is no need to configure anything and HPE 3PAR StoreServ Management Console can be used to run historical capacity and performance reports.

HPE 3PAR Autonomic Groups takes autonomic storage management a step further by allowing both hosts and VVs to be combined into groups or sets that can then be managed as a single object. Adding an object to an autonomic group applies all previously performed provisioning actions to the new member. For example, when a new host is added to a group, all volumes that were previously exported to the group are autonomically exported to the new host with absolutely no administrative intervention required. Similarly, when a new volume is added to a group, this volume is also autonomically exported to all hosts the group has previously been exported to—intelligently and with no administrator action required.

In fact, management of the HPE 3PAR StoreServ Storage system requires only knowledge of a few simple, basic functions:

- Create (for VVs and LDs)
- Remove (for VVs and LDs)
- Show (for resources)
- Stat (to display statistics)
- Hist (to display histograms)

Although there are a few other functions, these commands represent 90 percent of the console actions necessary, returning simplicity to the storage environment. Both the CLI and the HPE 3PAR StoreServ Management Console communicate with the corresponding server process on the HPE 3PAR StoreServ Storage system over TCP/IP over the on-board Gigabit Ethernet port on one of the controller nodes.

Management of the HPE 3PAR StoreServ Storage system benefits from very granular instrumentation within the HPE 3PAR Operating System. This instrumentation effectively tracks every I/O through the system and provides statistical information, including service time, I/O size, KB/s, and IOPS for VVs, LDs, and physical drives (PDs). Performance statistics such as CPU utilization, total accesses, and cache hit rate for reads and writes are also available on the controller nodes that make up the system cluster.

These statistics can be reported through the HPE 3PAR StoreServ Management Console or through the CLI. Moreover, administrators at operation centers powered by the leading enterprise management platforms can monitor MIB-II information from the HPE 3PAR StoreServ Storage system. All alerts are converted into SNMP Version 2c and Version 3 traps and sent to any configured SNMP management station.

An even more powerful and flexible way to manage HPE 3PAR StoreServ Storage systems than through the CLI and the HPE 3PAR StoreServ Management Console is through the use of the HPE 3PAR Web Services API. This API enables programmatic management of HPE 3PAR Storage systems. Using the API, the management of volumes, CPGs, and VLUNs can be automated through a series of HTTPS requests. The API consists of a server that is part of the HPE 3PAR Operating System and runs on the HPE 3PAR StoreServ Storage system itself and a definition of the operations, inputs, and outputs of the API. The software development kit (SDK) of the API includes a sample client that can be referenced for the development of customer-defined clients.

HPE's commitment to the OpenStack® community brings the power of OpenStack to the enterprise with new and enhanced offerings that enable enterprises to increase agility, speed innovation, and lower costs. Over the past two years, HP (now Hewlett Packard Enterprise) has been a top contributor to the advancement of the OpenStack project. HPE's contributions have focused on continuous integration and quality assurance, which support the development of a reliable and scalable cloud platform that is equipped to handle production workloads. To support the need that many larger organizations and service providers have for enterprise-class storage, Hewlett Packard Enterprise has developed the HPE 3PAR StoreServ Block Storage Drivers, which support the OpenStack technology across both iSCSI and Fibre Channel (FC) protocols. This provides the flexibility and cost-effectiveness of a cloud-based open source platform to customers with mission-critical environments and high resiliency requirements.

For more information about HPE 3PAR StoreServ support for OpenStack, please see the [OpenStack HPE 3PAR StoreServ Block Storage Driver Configuration Best Practices](#) and [HPE 3PAR StoreServ Storage integration with OpenStack solution brief](#).

HPE Smart SAN

SAN plays a critical role in any data center by providing access and connectivity between storage arrays and servers via a dedicated network. Fibre Channel (FC) has been the dominant storage protocol that enjoys significant SAN market share. FC is popular for storage because of its enterprise-class performance, availability, and security. Fibre Channel zoning is a key feature that adds to security and better management of the SAN by providing necessary segregation and allowing controlled communication among selected devices within a large fabric. However, configuring zones still is a complex, tedious, and error-prone operation in a majority of SAN installations. Thus, signifying a need for automating these operations as much as possible to avoid human errors and reduce potential SAN downtime.

HPE Smart SAN for HPE 3PAR StoreServ comes with a set of innovative features, one of which is automated zoning to address the above issues. In addition, it also supports standards-based device registrations and diagnostic data collection for better configuration, visibility, and diagnostic purposes. Automated zoning as implemented on HPE 3PAR StoreServ as part of HPE Smart SAN 1.0, utilizes peer zoning as defined in FC standards, thus empowering HPE 3PAR StoreServ Array to configure zones automatically whenever hosts are provisioned on the target side. This drastically improves the customer experience with significant reduction in overall SAN configuration time as well as making the whole process less error prone. As an example, in a mid-sized SAN consisting of nine fabric switches, zone configuration for 128 host initiators and eight HPE 3PAR target ports, resulted in a saving of over 80 percent configuration time with Target Device Peer Zoning (TDPZ).⁶

HPE StoreFront Remote Management (SFRM)

HPE StoreFront Remote Management is a unique pro-active tool that takes management leadership to another level. It is a Web portal dash board that maximizes asset utilization and optimizes data center environment by analyzing the performance delivered and capacity used by a HPE 3PAR StoreServ array over a period of time; this helps to forecast and plan the growth efficiently; as also by identifying capacity and performance issues early remedial measures can be taken on time; and all of these unique insights are presented at one place, at one time.

What's further unique about SFRM is its ability to anywhere and anytime access a HPE 3PAR StoreServ array over a Web portal with no restrictions on the number of HPE 3PAR StoreServ arrays accessed. With a user interface that is consistent with HPE 3PAR StoreServ Management Console (SSMC) and HPE OneView, this provides customers a unique ability to back data up into the cloud if they do sense any problem anytime.

Data optimization

HPE 3PAR StoreServ storage offers several products that can be used for service-level optimization. These solutions match data to the most cost-efficient resource capable of delivering the needed service level at any given time.

HPE 3PAR Dynamic Optimization is a powerful software feature that enables storage administrators to perform several online volume optimizations:

- Online data movement of existing volumes to different RAID levels and different tiers of storage. For example, application that requires high performance only during certain windows can be tuned to RAID 1 on SSDs and during lower activity periods moved to more cost-effective RAID 6 storage on nearline disks.
- Convert existing volumes to a different volume type. For example, Thinly Provisioned volumes on a HDD CPG to a Thinly Deduped volume on an SSD tier. Or the conversion of a full volume to a thinly provisioned volume. This conversion happens transparently and non-disruptively.

HPE 3PAR Adaptive Optimization software is another autonomic storage tiering tool that takes a fine-grained, policy-driven approach to service-level optimization. HPE 3PAR Adaptive Optimization works by analyzing performance (access rates) for sub-volume regions, then selects the most active regions (those with the highest I/O rates) and uses the proven sub-volume data movement engine built into the HPE 3PAR Operating System to relocate those regions to the fastest storage tier available. It also moves less active regions to slower tiers to help ensure space availability for newly active regions.

Traditional storage arrays require storage administrators to choose between slow, inexpensive storage and fast, expensive storage for each volume—a process that depends on the knowledge of the application's storage access patterns. Moreover, volumes tend to have hotspots rather than evenly distributed accesses, and these hotspots can move over time.

Using HPE 3PAR Adaptive Optimization, an HPE 3PAR StoreServ Storage system configured with nearline disks and Fast Class disks plus a small number of SSDs can approach flash performance at little more than the cost per megabyte of NL SAS-based storage, adapting autonomically as access patterns change. HPE 3PAR Dynamic Optimization and HPE 3PAR Adaptive Optimization are shown in figure 10.

⁶ HPE internal testing performed in one of the HPSD SQA labs, 2015

For more information about HPE 3PAR Adaptive Optimization Software, please see the [HPE 3PAR Adaptive Optimization white paper](#).




Figure 10. HPE 3PAR Dynamic Optimization and HPE 3PAR Adaptive Optimization

Autonomic rebalancing

The HPE 3PAR Operating System automatically creates a balanced system layout by mapping VVs to many LDs, which are composed of chunklets drawn from many physical disks. As new hardware, drives, drive chassis, and nodes are added to a system, the existing data can be laid out to use the new components and benefits from system-wide striping.

The HPE 3PAR Autonomic Rebalance feature provides the ability to analyze how volumes on the HPE 3PAR StoreServ Storage system are using physical disk space and makes intelligent, autonomic adjustments to help ensure better volume distribution when new hardware is added to the system. This rebalancing is achieved via the “tunesys” command.

HPE 3PAR Storage Federation

HPE 3PAR Storage Federation is a suite of technologies that solves one of the major pain-point when it comes to tech-refresh, the act of migrating the data itself, that enables customers to move off a 3rd party array to 3PAR and then by streamlining the entire 3PAR asset lifecycle management and technology refresh. HPE 3PAR Federation technologies elevate data to newer, current hardware technologies—while avoiding the pain associated with “fork-lift” upgrades that require extensive planning and preparation (and may include such difficulties as application downtimes or other impactful events.) And unlike other technologies, 3PAR federation gives the option to either repurpose legacy arrays, extend their value through further federation or replication, or simply retire them. Most importantly HPE 3PAR Storage Federation is an Enterprise solution, at no moment in time it will present a SPOF (like running on a single controller) and supports on-line migration of volumes that are exported to clusters and/or replicated.

Storage federation is a native functionality built into the HPE 3PAR Operating System that enables users to move data and workloads between arrays without impacting applications, users, or services. Simply and non-disruptively shift data between any model of HPE 3PAR StoreServ Storage system without additional overhead impact to hosts management layers or appliances. With a very broad support matrix of Host Operating Systems, seamless bidirectional data mobility on HPE 3PAR StoreServ Storage systems can also improve availability in not just physical hosted environments but also Hypervisor based environments like VMware, Microsoft Hyper-V etc.

Storage federation is the delivery of distributed volume management across self-governing, homogenous peer storage arrays, all managed by a single management interface call HPE 3PAR StoreServ Management Console. Federated data mobility allows live data to be easily and non-disruptively moved between HPE 3PAR StoreServ Storage systems without any impact to host performance of any complexity of an added appliance in the data path. This is very similar to the virtual machine mobility enabled by products like VMware vMotion, Microsoft Hyper-V Migration etc. but in the case of HPE 3PAR Storage Federation, data volume mobility is enabled between storage systems. Storage federation on HPE 3PAR StoreServ Storage systems is peer-to-peer relationship, something that is native to the system itself.




Figure 11. HPE 3PAR Storage Federation

Note

That Storage Federation is different from hierarchical virtualization, which is the delivery of consolidated or distributed volume management through appliances that hierarchically control a set of heterogeneous storage arrays. Hierarchical virtualization, also sometimes referred to as external storage virtualization, adds a new layer that has to be purchased and managed. A new layer not only introduces additional fault domains that need to be handled but also compromises the overall functionality of the entire system to be at the lowest common denominator.

In contrast, storage federation on HPE 3PAR StoreServ Storage systems delivers the following benefits:

- Keeps costs low (no redundant layer of intelligent controllers)
- Reduces failure domains (no additional layers)
- Maintains functionality of each of the peers
- Simplifies administration
- Reduced complexity of Interoperability between storage systems

Storage federation has emerged as a way to address and improve storage agility and efficiency at the data center and even metropolitan area level. The primary building blocks of HPE 3PAR Storage Federation are:

- HPE 3PAR Peer Motion
- HPE 3PAR Online Import
- HPE 3PAR Peer Persistence

HPE 3PAR Peer Motion

HPE 3PAR Peer Motion software is the first non-disruptive, do-it-yourself data mobility tool for enterprise block storage that does not require any external appliance to be included in the data path, nor does it introduce any additional overhead on the host resources. Unlike traditional block migration approaches, HPE 3PAR Peer Motion enables non-disruptive data mobility of storage volume(s) between any HPE 3PAR StoreServ Storage system that are a part of the Storage Federation, without complex planning or dependency on extra tools. HPE 3PAR Peer Motion offers bidirectional data mobility which helps setup your federation upfront and create a setup with back and forth data movement between two HPE 3PAR StoreServ Storage system within the same data center.

Federated multi-array bidirectional data mobility between HPE 3PAR StoreServ Storage systems is easy to implement and manage via the HPE 3PAR StoreServ Management Console (SSMC). Simple predefined workflows have been implemented within the SSMC that give you the

ability to move data—for example, moving all volumes associated with a host from one HPE 3PAR StoreServ Storage system to another or moving individual volumes for workload balancing purposes—all with just one click.

The HPE 3PAR Peer Motion software leverages the same built-in technology that powers the simple and rapid inline thin conversion of inefficient fat volumes on source arrays to more efficient, higher-utilization thin volumes on the destination HPE 3PAR StoreServ Storage system.

HPE 3PAR Peer Motion software allows all HPE 3PAR StoreServ Storage systems to participate in peering relationships with each other in order to provide the following flexibility benefits:

- **Federated workload balancing**—Moves workloads from over-utilized assets to underutilized ones
- **Federated asset management**—Non-disruptively adds new storage to the infrastructure or migrates data from older systems to newer ones
- **Federated thin provisioning**—Manages storage utilization and efficiency at the data center level, not the individual system level

For more information about HPE 3PAR Peer Motion, please see the [Storage Federation for non-disruptive data mobility white paper](#).

HPE 3PAR Online Import

Based on HPE 3PAR Peer Motion technology, HPE 3PAR Online Import software leverages federated data mobility on the HPE 3PAR StoreServ Storage array to simplify and expedite data migration from HPE EVA Storage, EMC VMAX, EMC VNX, and EMC CLARiiON CX4 arrays, Hitachi Data Systems (HDS), Hitachi TagmaStore Network Storage Controller (NSC), Hitachi Universal Storage Platforms (USP), Hitachi Virtual Storage Platforms (VSP), IBM XiV Gen2 & IBM XiV Gen3 Storage systems. With HPE 3PAR Online Import software, migration from these platforms can be performed in only five stages:

1. Set up the online import environment
2. Zone the host to the new system
3. Configure host multipathing
4. Unzone from the source, and start the migration
5. Validate the application

HPE 3PAR Online Import will virtualize your HPE EVA Storage, EMC Storage, HDS Storage or IBM Storage systems, so you can migrate from your legacy Storage platforms to HPE 3PAR StoreServ Storage system, without adding any external appliance in the data path or impacting your host performance. With an easy, Do-It-Yourself migration, HPE 3PAR Online Import software offers the ability to move the host data with minimal to no-disruption (depending on the Host OS) from EMC VMAX, EMC VNX, and EMC CLARiiON CX4, HDS NSC, HDS USP, HDS VSP, IBM XiV Gen2 or IBM XiV Gen3 Storage arrays to any model of HPE 3PAR StoreServ Storage system. The orchestration of data mobility provided for EMC HDS and IBM Storage systems is via an easy to use command driver tool called HPE 3PAR Online Import Utility.

HPE 3PAR Online Import for EVA uses HPE EVA Command View as the orchestration platform to enable direct migration of data from a source HPE EVA Storage system to a destination HPE 3PAR StoreServ Storage array without requiring host resources for data migration. The entire process can be completed with only minimal to no disruption (depending on host OS), sending EVA virtual disk and host configuration information to the HPE 3PAR StoreServ Storage array without the need to change host configurations or interrupt data access in most cases.

For more information on HPE 3PAR Online Import for HPE EVA, see the HPE EVA P6000 to [HPE 3PAR StoreServ Online Import Best Practices white paper](#). For more information on HPE 3PAR Online Import for EMC, HDS and IBM Storage systems, see the [HPE 3PAR Online Import Software for EMC Storage Solution brief](#), and the [HPE 3PAR Online Import Data sheet](#).

HPE 3PAR Peer Persistence

HPE 3PAR Peer Persistence software enables HPE 3PAR StoreServ Storage systems located within a metropolitan distance to act as peers to each other for delivering a high-availability, transparent failover solution for the connected VMware vSphere, Microsoft Hyper-V and Microsoft Windows® clusters. HPE 3PAR Peer Persistence allows an array-level, high-availability solution between two sites or data centers where failover and failback remains completely transparent to the hosts and applications running on those hosts. Unlike traditional disaster recovery models where the hosts (and applications) must be restarted upon failover, HPE 3PAR Peer Persistence allows hosts to remain online serving their

business applications, even when the serving of the I/O workload migrates transparently from the primary array to the secondary array, resulting in zero downtime.

In an HPE 3PAR Peer Persistence configuration, a host cluster would be deployed across two sites and an HPE 3PAR StoreServ Storage system would be deployed at each site. All hosts in the cluster would be connected to both of the HPE 3PAR StoreServ Storage systems. These HPE 3PAR StoreServ systems present the same set of VVs and VLUNs with same volume WWN to the hosts in that cluster. The VVs are synchronously replicated at the block level so that each HPE 3PAR StoreServ Storage system has a synchronous copy of the volume. A given volume would be primary on a given HPE 3PAR StoreServ Storage system at any one time. Using Asymmetric Logical Unit Access (ALUA), HPE 3PAR Peer Persistence presents the paths from the primary array (HPE 3PAR StoreServ Storage system on which the VV is primary) as “active/optimized” and the paths from the secondary array as “standby” paths. Issuing a switchover command on the array results in the relationship of the arrays to swap, and this is reflected back to the host by swapping the state of the paths from active to standby and vice versa. Under this configuration, both HPE 3PAR StoreServ Storage systems can be actively serving I/O under normal operation (albeit on separate volumes).

With HPE 3PAR Peer Persistence, Oracle RAC stretch cluster support is now available on Linux® (RHEL 6) over FC, FCoE, and iSCSI protocols.

For more information, see the [HPE 3PAR Peer Persistence white paper](#).




Figure 12. Transparent failover with HPE 3PAR Peer Persistence software

Proactive support

HPE Support for HPE 3PAR StoreServ Storage provides a global support infrastructure that leverages advanced system and support architectures for fast, predictive response and remediation. The HPE 3PAR Secure Service Architecture provides secure service communication between the **HPE 3PAR StoreServ Storage systems** at the customer site and **HPE 3PAR Central**, enabling secure diagnostic data transmission and remote service connections. Key diagnostic information such as system health statistics, configuration data, performance data, and system events can be transferred frequently and maintained centrally on a historical basis. As a result, proactive fault detection and analysis are improved and manual intervention is kept to a bare minimum.

This implementation provides automated analysis and reporting that delivers accuracy and consistency, full system information in hand that reduces onsite dependencies, and fully scripted and tested automated point-and-click service actions that reduce human error.

HPE 3PAR StoreServ Storage systems include a dedicated service processor, a server that monitors one or more.

HPE 3PAR StoreServ systems and enables remote monitoring and remote servicing of the array. The service processor is a physical server that is external to the HPE 3PAR StoreServ Storage system and communicates to it via TCP/IP. A virtual service processor is available for the HPE 3PAR StoreServ 8000 Storage series. Service processor with a redundant power supply is also available for the HPE 3PAR StoreServ 20000 series for higher availability.

The requirement for a dedicated server outside of the storage sub-systems separates the fault domains and ensures greater reliability for the storage array without compromising serviceability. The service processor is a physical server with the necessary OS already pre-installed from the factory. It is ready-to-go racked in a factory integrated storage setup that minimizes complexity in setup, installation and usage for the customer.

The service processor functions as the communication interface between a customer's IP network and HPE 3PAR Central by managing all service-related communications. It leverages the industry-standard Hypertext Transfer Protocol Secure (HTTPS) to secure and encrypt data for all inbound and outbound communications. The information collected via the service processor is sent to HPE 3PAR Central. This information includes system status, configuration, performance metrics, environmental information, alerts, and notification debug logs. No customer data is sent.

The data sent is used by Hewlett Packard Enterprise support teams to proactively monitor the array and contact the customer if potential issues are discovered. Customers are warned proactively about potential problems before they occur. In the case of switch issues, the customer is advised of an issue and replacement parts are dispatched. Trained Hewlett Packard Enterprise service personnel can service the system at the customer's convenience. If the service processor cannot dial Hewlett Packard Enterprise for any reason, both the HPE 3PAR StoreServ Storage system and Hewlett Packard Enterprise support centers will receive alerts.

The service processor is also used to download new patches, maintenance updates, new firmware revisions and diagnostics; it will store them and push them to the HPE 3PAR StoreServ Storage system for software upgrades. If remote access is needed for any reason, the customer can configure inbound secure access for OS upgrades, patches, and engineering access. If the customer's data center does not permit "phone home" devices, then all alerts and notifications will be sent to the customer's internal support team. The customer can then notify Hewlett Packard Enterprise support of an issue or suspected issue, either over the phone or via the Web.

In addition, all HPE 3PAR StoreServ Storage systems support a complimentary support service known as Over-Subscribed System Alerts (OSSA) in addition to and concurrent with automated remote monitoring, alerting, and notification. This automated monitoring tool performs proactive utilization checks on key system elements utilizing data that resides at HPE. This data is collected periodically from the system and sent to Hewlett Packard Enterprise. The intent is to provide valuable information such as storage node CPU utilization, disk IOPS, the number of host initiators per port, and other metrics to keep the HPE 3PAR StoreServ Storage system running optimally.

For more information about HPE 3PAR Secure Service Architecture, please see the [HPE 3PAR Secure Service Architecture white paper](#).

Summary

Virtualization, cloud computing, and ITaaS are driving new requirements around storage agility and efficiency that are pushing legacy architectures to their brink. With a modern, scale-out architecture that is designed to meet these new demands, HPE 3PAR StoreServ Storage delivers the flexibility to respond flexibly and efficiently to change, allowing you to:

- Consolidate with confidence onto a multi-tenant platform with six nines of availability
- Deliver uncompromising QoS for even the most demanding workloads
- Accelerate performance with a flash-optimized architecture featuring inline deduplication
- Cut capacity requirements by 50 percent and double virtual machine density
- Respond 8X faster with autonomic management
- Seamlessly refresh storage and maintain load balancing across arrays
- Painlessly migrate to Tier-1 storage built for cloud computing and ITaaS

HPE 3PAR StoreServ Storage does this all while driving up efficiency and resource utilization with hardware acceleration that enhances performance and lowers total cost of ownership for storage. With a range of models that all leverage the same scale-out architecture and a single operating system to bring Tier-1 data services to the midrange, deliver all-flash array performance, and provide mission-critical resiliency and QoS, HPE 3PAR StoreServ Storage is the only primary storage platform you will ever need for block, file, and object access.

For more information

For detailed and up-to-date specifications on each of these products, please refer to the product QuickSpecs:

- [HPE 3PAR StoreServ 20000 Storage QuickSpecs](#)
- [HPE 3PAR StoreServ 8000 Storage QuickSpecs](#)
- [HPE 3PAR StoreServ Software QuickSpecs](#)

Learn more at

hpe.com/storage/3PARStoreServ



Sign up for updates

★ Rate this document

© Copyright 2011, 2014–2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel is a trademark of Intel Corporation in the U.S. and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle is a registered trademark of Oracle and/or its affiliates. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries.



Hardware and Software Interoperability Matrix

This document provides interoperability matrices for Cisco Unified Computing System components and configurations that have been tested and validated by Cisco, by Cisco partners, or both. Use this document as a reference for supported hardware and software.

This document contains the following sections:

- [Audience](#)
- [Operating System Interoperability Matrix](#)
- [VM-FEX Software Interoperability Matrix](#)
- [Converged Network Adapter Interoperability Matrix](#)
- [Network Interface Card Interoperability Matrix](#)
- [RAID Controller on Motherboard Interoperability Matrix](#)
- [UCS Storage Accelerator Interoperability Matrix](#)
- [Storage Array Interoperability Matrix](#)
- [Enterprise Backup Interoperability Matrix](#)
- [Switch Interoperability Matrix](#)
- [Related Documentation](#)
- [Obtaining Documentation and Submitting a Service Request](#)

For C-Series servers managed by UCSM, *unless otherwise indicated in this document*, the supported firmware and drivers are those listed in the 1.4(7) and 1.4(8) C-Series Compatibility guide, located at http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html.

For information about compatibility with different infrastructure bundles, please refer to the Release Note for the desired infrastructure bundle release.

Audience

This document is designed for use by Cisco TAC, sales, support engineers, professional service partners, and systems administrators responsible for the design and deployment of the Unified Computing System in the data center environment.

This document applies to the B Series Servers in Cisco Unified Computing System Release 2.1(1). The content of this document is updated on a regular basis. The current version was created on 2014-05-14.

Operating System Interoperability Matrix

Table 1 lists the operating systems that the Unified Computing System B Series Servers support.

Table 1: Operating System

UCS Blade Server	Vendor	Operating System Version	Notes
B200, B250, B230, B440 M2 and B200, B22, B420 M3	Citrix	XenServer 6.0.2	21
B200, B250, B230, B440 M2 and B200, B22, B420 M3	Citrix	XenServer 6.1	21
B200, B250, B230, B440 M1 & M2 and B22, B200 M3	Microsoft	Windows Server 2008 SP2 x64	18, 21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	Microsoft	Windows Server 2008 R2 SP1 x64	18, 21
B200, B250, B230, B440 M2 and B200, B22, B420 M3	Microsoft	Windows Server 2012	18, 21
B230, B440 M2 and B200, B22, B420 M3	Oracle	OL 6.3 64bit with UEK 2.6.39-200	22
B230, B440 M2 and B200, B22, B420 M3	Oracle	OVM 3.1.1 64bit with UEK 2.6.39-200	21
B200, B250, B230, B440 M2	Oracle	Solaris 10 09/10 (U9)	21
B200, B250, B230, B440 M2	Oracle	Solaris 10 08/11 (U10)	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	Red Hat	Red Hat Enterprise Linux 5.7 64bit	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	Red Hat	Red Hat Enterprise Linux 5.8 64bit	21

UCS Blade Server	Vendor	Operating System Version	Notes
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	Red Hat	Red Hat Enterprise Linux 5.9 64bit	21
B200, B250 M2 and B230, B440 M1 & M2 and B200, B22, B420 M3	Red Hat	Red Hat Enterprise Linux 6.2 64bit	21
B200, B250 M2 and B230, B440 M1 & M2 and B200, B22, B420 M3	Red Hat	Red Hat Enterprise Linux 6.3 64bit	21
B200, B250, B230, B440 M2 and B200, B22, B420 M3	SuSE	SUSE Linux Enterprise Server 11.2 64bit	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	VMware	vSphere 4.1 U2	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	VMware	vSphere 4.1i U2	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	VMware	vSphere 4.1 U3	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	VMware	vSphere 4.1i U3	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	VMware	vSphere 5.0 U1	21
B200, B250, B230, B440 M1 & M2 and B200, B22, B420 M3	VMware	vSphere 5.0 U2	21
B200, B250, B230, B440 M2 and B200, B22, B420 M3	VMware	vSphere 5.1	21
B200, B250, B230, B440 M2 and B200, B22, B420 M3	VMware	vSphere 5.1 U1	21

UCS Blade Server	Vendor	Operating System Version	Notes
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Citrix	XenServer 6.0.2	20
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Citrix	XenServer 6.1	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Microsoft	Windows Server 2008 SP2 x64	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24M3(SFF),C420 M3	Microsoft	Windows Server 2008 R2 SP1 x64	20
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Microsoft	Windows Server 2012	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Novell	SUSE Linux Enterprise Server 11.2 64bit	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Oracle	OL 6.3 64bit with UEK 2.6.39-200	20

UCS Blade Server	Vendor	Operating System Version	Notes
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Oracle	OVM 3.1.1 64bit with UEK 2.6.39-200	20
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Oracle	Solaris 10 08/11 (U10)	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Red Hat	Red Hat Enterprise Linux 5.7 64bit	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Red Hat	Red Hat Enterprise Linux 5.8 64bit	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	Red Hat	Red Hat Enterprise Linux 6.2 64bit	20
C260 M2, C460 M2 C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	Red Hat	Red Hat Enterprise Linux 6.3 64bit	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	VMware	vSphere 4.1 U2	20

UCS Blade Server	Vendor	Operating System Version	Notes
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2,	VMware	vSphere 4.1i U2	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	VMware	vSphere 4.1 U3	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2,	VMware	vSphere 4.1i U3	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	VMware	vSphere 5.0 U1	20
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	VMware	vSphere 5.1	20

VM-FEX Software Interoperability Matrix

Table 2 lists the VM-FEX software that the Unified Computing System B Series Servers support.

Table 2: VM-FEX Software

Software Type	Software Version	Operating System	Driver
VM-FEX	VEM410-201108406	vSphere 4.1 U2	v132-4.2.1.1.4.1.0-2.0.4
VM-FEX	VEM410-201108406	vSphere 4.1i U2	v132-4.2.1.1.4.1.0-2.0.4
VM-FEX	VEM410-201108406	vSphere 4.1 U3	v132-4.2.1.1.4.1.0-2.0.4
VM-FEX	VEM410-201108406	vSphere 4.1i U3	v132-4.2.1.1.4.1.0-2.0.4

Software Type	Software Version	Operating System	Driver
VM-FEX	VEM500-20110825132140-BG-release	vSphere 5.0 U1	v132-4.2.1.1.4.1.0-3.0.4
VM-FEX	VEM500-20110825132140-BG-release	vSphere 5.0 U2	v132-4.2.1.1.4.1.0-3.0.4
VM-FEX	cisco-vem-v151-5.1-1.1.1.1	vSphere 5.1	cisco-vem-v151-5.1-1.1.1.1
VM-FEX	cisco-vem-v151-5.1-1.1.1.1	vSphere 5.1 U1	cisco-vem-v151-5.1-1.1.1.1
VM-FEX		Red Hat Enterprise Linux 6.2 64bit	2.1.1.41 (enic)
VM-FEX		Red Hat Enterprise Linux 6.3 64bit	2.1.1.41 (enic)
VM-FEX		Windows Server 2012	2.2.0.13 (PF, VF enic) / 2.2.0.11 (vmfex_ext)

Converged Network Adapter Interoperability Matrix

Table 3 lists the converged network adapters that the Unified Computing System B Series Servers support.

Table 3: Converged Network Adapters

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	OL 6.3 64bit with UEK 2.6.39-200	1.5.0.18 (FNIC) / 2.1.1.24 (NIC)	2.1(1)	14, 22
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	OVM 3.1.1 64bit with UEK 2.6.39-200	1.5.0.18 (FNIC) / 2.1.1.24 (NIC)	2.1(1)	14, 23, 28
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Red Hat Enterprise Linux 5.7 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Red Hat Enterprise Linux 5.8 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Red Hat Enterprise Linux 5.9 64bit	1.5.0.20 (FNIC) / 2.1.1.33 (ENIC)	2.1(1)	12, 14, 27
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Red Hat Enterprise Linux 6.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Red Hat Enterprise Linux 6.3 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	SUSE Linux Enterprise Server 11.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	14
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 4.1 U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 14, 31
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 4.1i U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	10, 12, 14, 19, 31
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 4.1 U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 14, 31
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 4.1i U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	10, 12, 14, 19, 31
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 5.0 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 5.0 U2	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 5.1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	vSphere 5.1 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B22 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Windows Server 2008 SP2 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	12, 14, 17, 25
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Windows Server 2008 R2 SP1 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	12, 14, 25
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	Windows Server 2012	2.2.0.17 (FNIC) / 2.2.0.13 (ENIC)	2.1(1)	12, 14, 30
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	XenServer 6.0.2	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	14
B200, B22, B420 M3	UCS 1240 Virtual Interface Card and Port Expander Card	XenServer 6.1	1.5.0.20 (FNIC) / 2.1.1.39-1 (ENIC)	2.1(1)	4, 14, 29
B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	OL 6.3 64bit with UEK 2.6.39-200	1.5.0.18 / 2.1.1.24 (NIC)	2.1(1)	14, 22
B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	OVM 3.1.1 64bit with UEK 2.6.39-200	1.5.0.18 / 2.1.1.24 (NIC)	2.1(1)	14, 23, 28

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Red Hat Enterprise Linux 5.7 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Red Hat Enterprise Linux 5.8 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Red Hat Enterprise Linux 5.9 64bit	1.5.0.20 (FNIC) / 2.1.1.33 (ENIC)	2.1(1)	12, 14, 27
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Red Hat Enterprise Linux 6.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Red Hat Enterprise Linux 6.3 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12, 14
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	SUSE Linux Enterprise Server 11.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	14
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 4.1 U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 14, 31
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 4.1i U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	10, 12, 14, 19, 31
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 4.1 U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 14, 31

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 4.1i U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	10, 12, 14, 19, 31
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 5.0 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 5.0 U2	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 5.1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	vSphere 5.1 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 14, 19, 31
B200, B230, B440 M2 and B200, B22 M3	UCS 1280 Virtual Interface Card	Windows Server 2008 SP2 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	12, 14, 17, 25
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Windows Server 2008 R2 SP1 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	12, 14, 25
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	Windows Server 2012	2.2.0.17 (FNIC) / 2.2.0.13 (ENIC)	2.1(1)	12, 14, 30
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	XenServer 6.0.2	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	14

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B230, B440 M2 and B200, B22, B420 M3	UCS 1280 Virtual Interface Card	XenServer 6.1	1.5.0.20 (FNIC) / 2.1.1.39-1 (ENIC)	2.1(1)	4, 14, 29
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	Red Hat Enterprise Linux 5.7 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	Red Hat Enterprise Linux 5.8 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	Red Hat Enterprise Linux 5.9 64bit	1.5.0.20 (FNIC) / 2.1.1.33 (ENIC)	2.1(1)	12, 27
B200, B250 M2 and B230, B440 M1 & M2	UCS M81KR Virtual Interface Card	Red Hat Enterprise Linux 6.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12
B200, B250 M2 and B230, B440 M1 & M2	UCS M81KR Virtual Interface Card	Red Hat Enterprise Linux 6.3 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	12
B200, B250, B230, B440 M2	UCS M81KR Virtual Interface Card	SUSE Linux Enterprise Server 11.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	vSphere 4.1 U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 31
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	vSphere 4.1i U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	10, 12, 19, 31
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	vSphere 4.1 U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 31
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	vSphere 4.1i U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	10, 12, 19, 31

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	vSphere 5.0 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 19, 31
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	vSphere 5.0 U2	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 19, 31
B200, B250, B230, B440 M2	UCS M81KR Virtual Interface Card	vSphere 5.1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 19, 31
B200, B250, B230, B440 M2	UCS M81KR Virtual Interface Card	vSphere 5.1 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 12, 19, 31
B200, B250, B230, B440 M1 & M2	UCS M81KR Virtual Interface Card	Windows Server 2008 SP2 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	12, 17, 25
B200, B250, B230, B440 M1 and M2	UCS M81KR Virtual Interface Card	Windows Server 2008 R2 SP1 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	12, 25
B200, B250, B230, B440 M2	UCS M81KR Virtual Interface Card	Windows Server 2012	2.2.0.17 (FNIC) / 2.2.0.13 (ENIC)	2.1(1)	12, 30
B200, B250, B230, B440 M2	UCS M81KR Virtual Interface Card	XenServer 6.0.2	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	
B200, B250, B230, B440 M2	UCS M81KR Virtual Interface Card	XenServer 6.1	1.5.0.20 (FNIC) / 2.1.1.39-1 (ENIC)	2.1(1)	4, 29
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF), C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Red Hat Enterprise Linux 5.7 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Red Hat Enterprise Linux 5.8 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Red Hat Enterprise Linux 6.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Red Hat Enterprise Linux 6.3 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	SUSE Linux Enterprise Server 11.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	vSphere 4.1 U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	vSphere 4.1i U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	vSphere 4.1 U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	vSphere 4.1i U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	vSphere 5.0 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	vSphere 5.1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Windows Server 2008 SP2 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	20, 25
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Windows Server 2008 R2 SP1 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	20, 25
C260 M2, C460 M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF),C420 M3	UCS VIC 1225 10-Gbps 2 port CNA SFP+	Windows Server 2012	2.2.0.17 (FNIC) / 2.2.0.13 (ENIC)	2.1(1)	20, 30

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Red Hat Enterprise Linux 5.7 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Red Hat Enterprise Linux 5.8 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Red Hat Enterprise Linux 6.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Red Hat Enterprise Linux 6.3 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	SUSE Linux Enterprise Server 11.2 64bit	1.5.0.20 (FNIC) / 2.1.1.41 (ENIC)	2.1(1)	20
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	vSphere 4.1 U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	vSphere 4.1i U2	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	vSphere 4.1 U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	vSphere 4.1i U3	1.4.0.213 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	vSphere 5.0 U1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	vSphere 5.1	1.5.0.20 (FNIC) / 2.1.2.38 (ENIC)	2.1(1)	11, 19, 20, 31
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Windows Server 2008 SP2 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	20, 25

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
C200/C210 M1 & M2, C200 M2(SFF), C250 M1 & M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Windows Server 2008 R2 SP1 x64	2.1.0.20 (FNIC) / 2.2.0.13 (ENIC) / 2.1.0.4 (NIC Teaming Driver)	2.1(1)	20, 25
C200/C210 M2, C200 M2(SFF), C250 M2, C260 M2, C460 M1 & M2, C220/C240 M3(SFF/LFF), C22/C24 M3(SFF)	P81E Dual Port 10-Gbps Ethernet to PCIe Virtual Interface Card	Windows Server 2012	2.2.0.17 (FNIC) / 2.2.0.13 (ENIC)	2.1(1)	20, 30
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.7 64bit	8.2.0.106 (FC) / 4.0.462.0 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.8 64bit	8.2.0.108.4p (FC) / 4.0.100r (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.9 64bit	8.2.0.128.3p (FC) / 4.2.116r (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Red Hat Enterprise Linux 6.2 64bit	8.3.5.45.4p (FC) / 4.0.160r (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Red Hat Enterprise Linux 6.3 64bit	8.3.6.68.5p (FC) / 4.1.307r (ETH)	4.0.467.106	8

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 4.1 U2	8.2.1.105.34 (FC) / 4.0.306.0 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 4.1i U2	8.2.1.105.34 (FC) / 4.0.306.0 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 4.1 U3	8.2.1.105.34 (FC) / 4.0.306.0 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 4.1i U3	8.2.1.105.34 (FC) / 4.0.306.0 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 5.0 U1	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 5.1	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	vSphere 5.1 U1	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Windows Server 2008 R2 SP1 x64	7.2.50.007 (FC) / 4.0.449.0 (ETH)	4.0.467.106	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-E Converged Network Adapter	Windows Server 2012	2.72.012.001 (FC) / 4.2.313.0 (ETH)	4.0.467.106	8

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.7 64bit	8.2.0.106 (FC) / 4.0.462.0 (ETH)	4.0.467.0	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.8 64bit	8.2.0.108.4p (FC) / 4.0.100r (ETH)	4.0.467.0	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.9 64bit	8.2.0.128.3p (FC) / 4.2.116r (ETH)	4.0.467.0	8
B200, B250 M2 and B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Red Hat Enterprise Linux 6.2 64bit	8.3.5.45.4p (FC) / 4.0.160r (ETH)	4.0.467.0	8
B200, B250 M2 and B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Red Hat Enterprise Linux 6.3 64bit	8.3.5.68.5p (FC) / 4.1.307r (ETH)	4.0.467.0	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	vSphere 4.1 U3	8.2.1.105.34 (FC) / 4.0.306.0 (ETH)	4.0.467.0	8, 11
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	vSphere 4.1i U3	8.2.1.105.34 (FC) / 4.0.306.0 (ETH)	4.0.467.0	8, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	vSphere 5.0 U1	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.0	8, 11, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	vSphere 5.0 U2	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.0	8, 11, 19

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M2	UCS M72KR-E Converged Network Adapter	vSphere 5.1	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.0	8, 11, 19
B200, B250, B230, B440 M2	UCS M72KR-E Converged Network Adapter	vSphere 5.1 U1	8.2.2.105.36 (FC) / 4.0.355.1 (ETH)	4.0.467.0	8, 11, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Windows Server 2008 SP2 x64	7.2.50.007 (FC) / 4.0.449.0 (ETH)	4.0.467.0	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-E Converged Network Adapter	Windows Server 2008 R2 SP1 x64	7.2.50.007 (FC) / 4.0.449.0 (ETH)	4.0.467.0	8
B200, B250, B230, B440 M2	UCS M72KR-E Converged Network Adapter	Windows Server 2012	2.72.012.001 (FC) / 4.2.313.0 (ETH)	4.0.467.0	8
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.7 64bit	8.2.0.96.2p (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.8 64bit	8.2.0.108.4p (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	Red Hat Enterprise Linux 5.9 64bit	8.2.0.128.3p (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	vSphere 4.1 U2	8.2.1.30.1-58vmw (FC)	2.1(1)	8, 13

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	vSphere 4.1i U2	8.2.1.30.1-58vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	vSphere 4.1 U3	8.2.1.30.1-58vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	vSphere 4.1i U3	8.2.1.30.1-58vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	Windows Server 2008 SP2 x64	7.2.41.002 (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-E Converged Network Adapter	Windows Server 2008 R2 SP1 x64	7.2.41.002 (FC)	2.1(1)	8, 13
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.7 64bit	8.03.07.12.5.6-k (FC) / 5.0.25.1 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.8 64bit	8.03.07.09.05.08-k (FC) / 5.0.24 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.9 64bit	8.03.07.15.05.09-k (FC) / 5.0.29 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Red Hat Enterprise Linux 6.2 64bit	8.03.07.12.06.0-k (FC) / 5.0.25.1 (ETH)	1.10.84	8

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Red Hat Enterprise Linux 6.3 64bit	8.04.00.04.06.3-k (FC) / 1.00.00.30 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 4.1 U2	841.k1.34.1-1vmw (FC) / 4.0.735 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 4.1i U2	841.k1.34.1-1vmw (FC) / 4.0.735 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 4.1 U3	841.k1.34.1-1vmw (FC) / 4.0.735 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 4.1i U3	841.k1.34.1-1vmw (FC) / 4.0.735 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 5.0 U1	911.k1.1-19vmw (FC) / 5.0.736 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 5.1	911.k1.1-19vmw (FC) / 5.0.736 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	vSphere 5.1 U1	911.k1.1-19vmw (FC) / 5.0.736 (ETH)	1.10.84	8
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Windows Server 2008 R2 SP1 x64	9.1.9.38 (FC) / 4.6.6.1213 (ETH)	1.10.84	8

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B230, B440 M2 and B200, B22, B420 M3	UCS M73KR-Q Converged Network Adapter	Windows Server 2012	9.1.10.15 (FC) / 4.6.6.1213 (ETH)	1.10.84	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.7 64bit	8.03.07.03.05.07-k (FC) / 1.0.00.25 (ETH)	1.02.13	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.8 64bit	8.03.07.09.05.08-k (FC) / 1.0.00.29 (ETH)	1.02.13	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.9 64bit	8.0.3.07.15.05.09-k (FC) / 1.0.00.30 (ETH)	1.02.13	8
B200, B250 M2 and B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Red Hat Enterprise Linux 6.2 64bit	8.03.07.05.06.2-k (FC) / 1.00.00.29 (ETH)	1.02.13	8
B200, B250 M2 and B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Red Hat Enterprise Linux 6.3 64bit	8.04.00.04.06.3-k (FC) / 1.00.00.30 (ETH)	1.02.13	8
B200, B250, B230, B440 M2	UCS M72KR-Q Converged Network Adapter	Solaris 10 09/10 (U9)	20100301-3.00 (FC) / 100429-v1.05 (ETH)	1.02.13	8
B200, B250, B230, B440 M2	UCS M72KR-Q Converged Network Adapter	Solaris 10 08/11 (U10)	3.06 (FC) / 1.10 (ETH)	1.02.13	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	vSphere 4.1 U2	832.k1.27vmw (FC) / 1.0.0.43 (ETH)	1.02.13	8, 11

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	vSphere 4.1i U2	832.k1.27vmw (FC) / 1.0.0.43 (ETH)	1.02.13	8, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	vSphere 4.1 U3	832.k1.27vmw (FC) / 1.0.0.43 (ETH)	1.02.13	8, 11
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	vSphere 4.1i U3	832.k1.27vmw (FC) / 1.0.0.43 (ETH)	1.02.13	8, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	vSphere 5.0 U1	911.k1.1-19vmw (FC) / 1.0.0.47 (ETH)	1.02.13	8, 11, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	vSphere 5.0 U2	911.k1.1-19vmw (FC) / 1.0.0.47 (ETH)	1.02.13	8, 11, 19
B200, B250, B230, B440 M2	UCS M72KR-Q Converged Network Adapter	vSphere 5.1	911.k1.1-19vmw (FC) / 1.0.0.47 (ETH)	1.02.13	8, 11, 19
B200, B250, B230, B440 M2	UCS M72KR-Q Converged Network Adapter	vSphere 5.1 U1	911.k1.1-19vmw (FC) / 1.0.0.47 (ETH)	1.02.13	8, 11, 19
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Windows Server 2008 SP2 x64	9.1.8.27 (FC) / 1.0.1.12 (ETH)	1.02.13	8
B200, B250, B230, B440 M1 & M2	UCS M72KR-Q Converged Network Adapter	Windows Server 2008 R2 SP1 x64	9.1.8.27 (FC) / 1.0.1.12 (ETH)	1.02.13	8

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M2	UCS M72KR-Q Converged Network Adapter	Windows Server 2012	9.1.10.15 (FC) / 1.0.1.15 (ETH)	1.02.13	8
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.7 64bit	8.03.07.03.05.07-k (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.8 64bit	8.03.07.09.05.08-k (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	Red Hat Enterprise Linux 5.9 64bit	8.03.07.15.05.09-k (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	vSphere 4.1 U2	831.k1.28.1-1vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	vSphere 4.1i U2	831.k1.28.1-1vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	vSphere 4.1 U3	831.k1.28.1-1vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	vSphere 4.1i U3	831.k1.28.1-1vmw (FC)	2.1(1)	8, 13
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	Windows Server 2008 SP2 x64	9.1.8.27 (FC)	2.1(1)	8, 13

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250 M1 & M2 and B440 M1	UCS M71KR-Q Converged Network Adapter	Windows Server 2008 R2 SP1 x64	9.1.8.27 (FC)	2.1(1)	8, 13

Network Interface Card Interoperability Matrix

Table 4 lists the network interface cards that the Unified Computing System B Series Servers support.

Table 4: Network Interface Cards

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.7 64bit	1.62.00-6 (ETH) / 2.1.3b (iSCSI)	6.2.15.23.7.1	
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.8 64bit	1.70.00-0 (ETH)	6.2.15.23.7.1	
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.9 64bit	7.2.51-0 (ETH)	6.2.15.23.7.1	
B200, B250 M2 and B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 6.2 64bit	1.62.00-6 (ETH)	6.2.15.23.7.1	
B200, B250 M2 and B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 6.3 64bit	1.70.00-0 (ETH)	6.2.15.23.7.1	
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	vSphere 4.1 U2	1.62.15.v41 (ETH) / 1.9.1t.v41.2 (iSCSI)	6.2.15.23.7.1	11

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	vSphere 4.1i U2	1.62.15.v41 (ETH)	6.2.15.23.7.1	12, 19
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	vSphere 5.0 U1	1.61.15.v50.1 (ETH)	6.2.15.23.7.1	
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Windows Server 2008 SP2 x64	6.2.9.0 (ETH)	6.2.15.23.7.1	
B200, B250, B230, B440 M1 & M2	UCS M51KR-B 10-Gigabit Ethernet Adapter	Windows Server 2008 R2 SP1 x64	6.2.9.0 (ETH) / 6.2.7.0 (iSCSI)	6.2.15.23.7.1	12, 17
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.7 64bit	3.2.9 (ETH)	2.1.60.1.1	
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.8 64bit	3.4.8-k (ETH)	2.1.60.1.1	
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.9 64bit	3.4.8-k (ETH)	2.1.60.1.1	
B200, B250 M2 and B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 6.2 64bit	3.3.9 (ETH)	2.1.60.1.1	
B200, B250 M2 and B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 6.3 64bit	3.4.8-k (ETH)	2.1.60.1.1	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	vSphere 4.1 U2	3.1.17.1 (ETH)	2.1.60.1.1	11
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	vSphere 4.1i U2	3.1.17.1 (ETH)	2.1.60.1.1	19
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	vSphere 4.1 U3	3.1.17.1 (ETH)	2.1.60.1.1	11
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	vSphere 4.1i U3	3.1.17.1 (ETH)	2.1.60.1.1	19
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	vSphere 5.0 U1	3.7.13 (ETH)	2.1.60.1.1	11, 19
B200, B250, B230, B440 M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	vSphere 5.1	3.7.13 (ETH)	2.1.60.1.1	11, 19
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Windows Server 2008 SP2 x64	2.7.28.0 (ETH)	2.1.60.1.1	
B200, B250, B230, B440 M1 & M2	UCS M61KR-I 10-Gigabit Ethernet Adapter	Windows Server 2008 R2 SP1 x64	2.7.28.0 (ETH)	2.1.60.1.1	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.7 64bit	3.2.9 (ETH)	2.1.11.0.70	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.8 64bit	3.4.8-k (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	Red Hat Enterprise Linux 5.9 64bit	3.4.8-k (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	vSphere 4.1 U2	2.0.38.2.5.1-1vmw (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	vSphere 4.1i U2	2.0.38.2.5.1-1vmw (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	vSphere 4.1 U3	2.0.38.2.5.1-1vmw (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	vSphere 4.1i U3	2.0.38.2.5.1-1vmw (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	Windows Server 2008 SP2 x64	2.4.29.1 (ETH)	2.1.11.0.70	
B200, B250 M1 & M2	UCS 82598KR-CI 10-Gigabit Ethernet Adapter	Windows Server 2008 R2 SP1 x64	2.4.29.1 (ETH)	2.1.11.0.70	

RAID Controller on Motherboard Interoperability Matrix

Table 5 lists the RAID controllers on motherboard that the Unified Computing System B Series Servers support.

Table 5: RAID Controller on Motherboard

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Red Hat Enterprise Linux 5.7 64bit	3.0.4.18rh	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Red Hat Enterprise Linux 5.8 64bit	3.0.4.20rh	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Red Hat Enterprise Linux 5.9 64bit	3.04.20rh1	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Red Hat Enterprise Linux 6.2 64bit	3.04.19	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Red Hat Enterprise Linux 6.3 64bit	3.0.4.20	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Solaris 10 09/10 (U9)	1.100	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Solaris 10 08/11 (U10)	1.109	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	SUSE Linux Enterprise Server 11.2 64bit	4.28.00.00	1.32.04.00	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 4.1 U2	4.21.00.01.1-5vmw	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 4.1i U2	4.21.00.01.1-5vmw	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 4.1 U3	4.21.00.01.1-5vmw	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 4.1i U3	4.21.00.01.1-5vmw	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 5.0 U1	4.23.01.00-5	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 5.0 U2	4.23.01.00-5	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 5.1	4.23.01.00-5	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	vSphere 5.1 U1	4.23.01.00-5	1.32.04.00	
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Windows Server 2008 SP2 x64	1.34.02.00	1.32.04.00	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B250 M1 & M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Windows Server 2008 R2 SP1 x64	1.34.02.00	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	Windows Server 2012	1.34.2.6	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	XenServer 6.0.2	4.22	1.32.04.00	
B200, B250 M2	LSI Logic SAS1064E PCI-Express Fusion-MPT	XenServer 6.1	4.28.00.00	1.32.04.00	
B230 M2	LSI Logic MegaSAS 2008 (9240)	OL 6.3 64bit with UEK 2.6.39-200	00.00.05.40	20.10.1-0100	22
B230 M2	LSI Logic MegaSAS 2008 (9240)	OVM 3.1.1 64bit with UEK 2.6.39-200	00.00.05.40	20.10.1-0100	
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Red Hat Enterprise Linux 5.7 64bit	00.00.05.38-rh1	20.10.1-0100	
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Red Hat Enterprise Linux 5.8 64bit	00.00.05.40-rh2	20.10.1-0100	
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Red Hat Enterprise Linux 5.9 64bit	00.00.06.15-rh	20.10.1-0100	
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Red Hat Enterprise Linux 6.2 64bit	00.00.05.40-rh2	20.10.1-0100	
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Red Hat Enterprise Linux 6.3 64bit	00.00.06.14-rh1	20.10.1-0100	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B230 M2	LSI Logic MegaSAS 2008 (9240)	Solaris 10 09/10 (U9)	03.03.00	20.10.1-0100	3
B230 M2	LSI Logic MegaSAS 2008 (9240)	Solaris 10 08/11 (U10)	03.03.00	20.10.1-0100	3
B230 M2	LSI Logic MegaSAS 2008 (9240)	SUSE Linux Enterprise Server 11.2 64bit	00.00.05.38-rc1	20.10.1-0100	
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	vSphere 4.1 U2	5.33	20.10.1-0100	11
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	vSphere 4.1i U2	5.33	20.10.1-0100	19
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	vSphere 4.1 U3	5.33	20.10.1-0100	11
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	vSphere 4.1i U3	5.33	20.10.1-0100	19
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	vSphere 5.0 U1	5.34	20.10.1-0100	11, 19
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	vSphere 5.0 U2	5.34	20.10.1-0100	11, 19
B230 M2	LSI Logic MegaSAS 2008 (9240)	vSphere 5.1	5.34	20.10.1-0100	11, 19
B230 M2	LSI Logic MegaSAS 2008 (9240)	vSphere 5.1 U1	5.34	20.10.1-0100	11, 19
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Windows Server 2008 SP2 x64	4.32.0.64	20.10.1-0100	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B230 M1 & M2	LSI Logic MegaSAS 2008 (9240)	Windows Server 2008 R2 SP1 x64	4.32.0.64	20.10.1-0100	
B230 M2	LSI Logic MegaSAS 2008 (9240)	Windows Server 2012	5.2.122.0	20.10.1-0100	
B230 M2	LSI Logic MegaSAS 2008 (9240)	XenServer 6.0.2	5.41	20.10.1-0100	
B230 M2	LSI Logic MegaSAS 2008 (9240)	XenServer 6.1	6.15	20.10.1-0100	
B440 M2	LSI Logic MegaSAS 2108 (9260)	OL 6.3 64bit with UEK 2.6.39-200	00.00.05.40	12.12.0-0050	22
B440 M2	LSI Logic MegaSAS 2108 (9260)	OVM 3.1.1 64bit with UEK 2.6.39-200	00.00.05.40	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Red Hat Enterprise Linux 5.7 64bit	00.00.05.38-rh1	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Red Hat Enterprise Linux 5.8 64bit	00.00.05.40-rh2	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Red Hat Enterprise Linux 5.9 64bit	00.00.06.15-rh	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Red Hat Enterprise Linux 6.2 64bit	00.00.05.40-rh2	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Red Hat Enterprise Linux 6.3 64bit	00.00.06.14-rh1	12.12.0-0050	
B440 M2	LSI Logic MegaSAS 2108 (9260)	Solaris 10 09/10 (U9)	2.4	12.12.0-0050	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B440 M2	LSI Logic MegaSAS 2108 (9260)	Solaris 10 08/11 (U10)	2.9	12.12.0-0050	
B440 M2	LSI Logic MegaSAS 2108 (9260)	SUSE Linux Enterprise Server 11.2 64bit	00.00.05.38-rc1	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	vSphere 4.1 U2	5.33	12.12.0-0050	11
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	vSphere 4.1i U2	5.33	12.12.0-0050	19
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	vSphere 4.1 U3	5.33	12.12.0-0050	11
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	vSphere 4.1i U3	5.33	12.12.0-0050	19
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	vSphere 5.0 U1	5.34	12.12.0-0050	11, 19
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	vSphere 5.0 U2	5.34	12.12.0-0050	11, 19
B440 M2	LSI Logic MegaSAS 2108 (9260)	vSphere 5.1	5.34	12.12.0-0050	11, 19
B440 M2	LSI Logic MegaSAS 2108 (9260)	vSphere 5.1 U1	5.34	12.12.0-0050	11, 19
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Windows Server 2008 SP2 x64	4.32.0.64	12.12.0-0050	
B440 M1 & M2	LSI Logic MegaSAS 2108 (9260)	Windows Server 2008 R2 SP1 x64	4.32.0.64	12.12.0-0050	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B440 M2	LSI Logic MegaSAS 2108 (9260)	Windows Server 2012	5.2.122.0	12.12.0-0050	
B440 M2	LSI Logic MegaSAS 2108 (9260)	XenServer 6.0.2	5.41	12.12.0-0050	
B440 M2	LSI Logic MegaSAS 2108 (9260)	XenServer 6.1	6.15	12.12.0-0050	
B200 M3	LSI Logic MegaRAID SAS 2004	OL 6.3 64bit with UEK 2.6.39-200	00.00.05.40	20.10.1-0100	22
B200 M3	LSI Logic MegaRAID SAS 2004	OVM 3.1.1 64bit with UEK 2.6.39-200	00.00.05.40	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Red Hat Enterprise Linux 5.7 64bit	00.00.05.38-rh1	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Red Hat Enterprise Linux 5.8 64bit	00.00.05.40-rh2	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Red Hat Enterprise Linux 5.9 64bit	00.00.06.15-rh	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Red Hat Enterprise Linux 6.2 64bit	00.00.05.40-rh2	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Red Hat Enterprise Linux 6.3 64bit	00.00.06.14-rh1	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	SUSE Linux Enterprise Server 11.2 64bit	00.00.05.38-rc1	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 4.1 U2	5.33	20.10.1-0100	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 4.1i U2	5.33	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 4.1 U3	5.33	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 4.1i U3	5.33	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 5.0 U1	5.34	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 5.0 U2	5.34	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 5.1	5.34	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	vSphere 5.1 U1	5.34	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Windows Server 2008 SP2 x64	5.1.112.64	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Windows Server 2008 R2 SP1 x64	5.1.112.64	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	Windows Server 2012	5.2.122.0	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	XenServer 6.0.2	5.41	20.10.1-0100	
B200 M3	LSI Logic MegaRAID SAS 2004	XenServer 6.1	6.15	20.10.1-0100	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B22 M3	LSI Logic MegaRAID SAS 2002	OL 6.3 64bit with UEK 2.6.39-200	00.00.05.40	20.10.1-0100	22
B22 M3	LSI Logic MegaRAID SAS 2002	OVM 3.1.1 64bit with UEK 2.6.39-200	00.00.05.40	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Red Hat Enterprise Linux 5.7 64bit	00.00.05.38-rh1	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Red Hat Enterprise Linux 5.8 64bit	00.00.05.40-rh2	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Red Hat Enterprise Linux 5.9 64bit	00.00.06.15-rh	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Red Hat Enterprise Linux 6.2 64bit	00.00.05.40-rh2	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Red Hat Enterprise Linux 6.3 64bit	00.00.06.14-rh1	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	SUSE Linux Enterprise Server 11.2 64bit	00.00.05.38-rc1	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 4.1 U2	5.33	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 4.1i U2	5.33	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 4.1 U3	5.33	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 4.1i U3	5.33	20.10.1-0100	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 5.0 U1	5.34	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 5.0 U2	5.34	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 5.1	5.34	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	vSphere 5.1 U1	5.34	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Windows Server 2008 SP2 x64	5.1.112.64	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Windows Server 2008 R2 SP1 x64	5.1.112.64	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	Windows Server 2012	5.2.122.0	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	XenServer 6.0.2	5.41	20.10.1-0100	
B22 M3	LSI Logic MegaRAID SAS 2002	XenServer 6.1	6.15	20.10.1-0100	
B420 M3	LSI Logic MegaRAID SAS 2208	OL 6.3 64bit with UEK 2.6.39-200	00.00.05.40	23.2.1-0056	22
B420 M3	LSI Logic MegaRAID SAS 2208	OVM 3.1.1 64bit with UEK 2.6.39-200	00.00.05.40	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	Red Hat Enterprise Linux 5.7 64bit	00.00.05.38-rh1	23.2.1-0056	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B420 M3	LSI Logic MegaRAID SAS 2208	Red Hat Enterprise Linux 5.8 64bit	00.00.05.40-rh2	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	Red Hat Enterprise Linux 5.9 64bit	00.00.06.15-rh	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	Red Hat Enterprise Linux 6.2 64bit	00.00.05.40-rh2	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	Red Hat Enterprise Linux 6.3 64bit	00.00.05.40-rh2	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	SUSE Linux Enterprise Server 11.2 64bit	00.00.05.38-rc1	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 4.1 U2	5.33	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 4.1i U2	5.33	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 4.1 U3	5.33	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 4.1i U3	5.33	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 5.0 U1	5.34	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 5.0 U2	5.34	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 5.1	5.34	23.2.1-0056	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B420 M3	LSI Logic MegaRAID SAS 2208	vSphere 5.1 U1	5.34	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	Windows Server 2008 R2 SP1 x64	5.2.116.64	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	Windows Server 2012	5.2.122.0	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	XenServer 6.0.2	5.41	23.2.1-0056	
B420 M3	LSI Logic MegaRAID SAS 2208	XenServer 6.1	6.15	23.2.1-0056	

UCS Storage Accelerator Interoperability Matrix

Table 6 lists the UCS storage accelerators that are supported with this version of the Unified Computing System.

Table 6: UCS Storage Accelerators

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	Windows Server 2008 R2 SP1 x64	2.10.58.00	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	Windows Server 2012	2.10.58.00	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	Red Hat Enterprise Linux 5.7 64bit	14.00.00.00	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	Red Hat Enterprise Linux 5.8 64bit	14.00.00.00	08.65.02.00	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	Red Hat Enterprise Linux 6.2 64bit	14.00.00.00	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	Red Hat Enterprise Linux 6.3 64bit	14.00.00.00	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	vSphere 4.1i U3	14.00.00.00.1	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	vSphere 5.0 U1	14.00.00.00.1	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	vSphere 5.1	14.00.00.00.1	08.65.02.00	
B200, B22, B420 M3	LSI 400GB SLC WarpDrive	vSphere 5.1 U1	14.00.00.00.1	08.65.02.00	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	Windows Server 2008 R2 SP1 x64	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	Windows Server 2012	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 5.7 64bit	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 5.8 64bit	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 6.2 64bit	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 6.3 64bit	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	vSphere 4.1i U3	3.2.2.869	7.1.13	

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	vSphere 5.0 U1	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	vSphere 5.1	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 365GB MLC Fusion-io ioDrive2	vSphere 5.1 U1	3.2.2.869	7.1.13	
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	Windows Server 2008 R2 SP1 x64	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	Windows Server 2012	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 5.7 64bit	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 5.8 64bit	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 6.2 64bit	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	Red Hat Enterprise Linux 6.3 64bit	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	vSphere 4.1i U3	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	vSphere 5.0 U1	3.2.2.869	7.1.13	34
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	vSphere 5.1	3.2.2.869	7.1.13	34

UCS Blade Server	Adapter Model	Operating System	Adapter Driver	Adapter Firmware	Notes
B200, B22, B420 M3	UCS 785GB MLC Fusion-io ioDrive2	vSphere 5.1 U1	3.2.2.869	7.1.13	34

Storage Array Interoperability Matrix

Table 7 lists the storage array models that are supported with this version of the Unified Computing System.

Cisco's interoperability guidelines in this table apply to the following topologies and protocols:

- FC/FCoE-attached storage arrays (Fabric Interconnects in NPV mode or FC Switching Mode); unless otherwise indicated FCoE NPV connections are supported to the devices listed in the FC switch table
- UCSM-configured iSCSI boot
- iSCSI data LUN access if the storage array is directly connected to UCS Fabric Interconnects using Appliance Ports
- NFS or CIFS data access if the storage array is directly connected to UCS Fabric Interconnects using Appliance Ports



Note

Cisco UCS has no interoperability requirements for storage arrays configured to use iSCSI (Data Only), NFS, or CIFS when connected through UCS Fabric Interconnect Uplink Port.

Table 7: Storage Array

Vendor	Storage Array Model	Array Firmware	Notes
Dell Compellent	SC040, SC8000	Refer to storage vendor's support list, http://www.compellent.com/Support.aspx	5, 6, 7, 32
Dell Equallogic	PS6010, PS6110	Refer to storage vendor's support list, https://support.equallogic.com/secure/login.aspx	6, 7, 12
EMC	EMC Symmetrix family (VMAX & DMX), EMC CLARiiON (AX & CX) EMC Unified Storage VNX (5300, 5500, 5700, 7500) EMC Celerra (Native Block) EMC Isilon (NFS/CIFS) EMC VPLEX	Refer to storage vendor's support list, https://elabnavigator.emc.com	5, 6, 7, 12, 26, 32

Vendor	Storage Array Model	Array Firmware	Notes
HP	EVA 4000/6000/8000/4100/6100/8100, EVA 4400/6400/8400 XP 10000/12000, XP 20000/24000, P9500, MSA 2000, P2000	Refer to storage vendor's support list, http://h20272.www2.hp.com/	6, 7
HP	HP 3PAR F200, F400, T400, T800 HP 3PAR StoreServ 10000 Storage HP 3PAR StoreServ 7000 Storage	Refer to storage vendor's support list, http://h20272.www2.hp.com/	5, 6, 7
HDS	USP-V, USP-VM, VSP, AMS 1000, 2000 Series: AMS 2100, AMS 2300, AMS 2500, HUS100 Series, HUS-VM	Refer to storage vendor's support list, http://www.hds.com/products/interoperability/	5, 6, 7, 12
IBM	DS 3500, DS 4800, DS 5300, DS 8100, DS 8800, SVC, Storwize V7000, IBM N-Series, XIV(Gen 2)	Refer to storage vendor's support list, http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss	5, 6, 7, 33
Netapp	FAS2000 series, FAS3000 series, FAS6000 series, FAS3100 series, FAS3200 series, FAS6200 series, V3000 series, V6000 series	Refer to storage vendor's support list, http://support.netapp.com/	5, 6, 7, 12
Nimble	CS-Series Family	Refer to storage vendors support list, http://www.nimblestorage.com/support/	6, 7, 12
Nexsan	Nexsan E60, E60X, E48 , E48X, E18, E18X	Refer to storage vendor's support list, https://partners.nexsan.com/Nexsan_PHP/webFileLocation/compatibility_guide/Nexsan_Compatibility_Guide.pdf	2
SolidFire	SF3010, SF6010, SF9010	Refer to storage vendor's support list, http://www.solidfire.com/storage-system/	6, 7, 12
Tegile	Zebi Storage Array Family	Refer to storage vendors support list, http://www.tegile.com/support	5, 6
Tintri	Tintri VMstore T540	Refer to storage vendors support list, http://www.tintri.com/support	26
Violin	3000, 6000 Series Flash Memory Arrays	Refer to storage vendors support list, http://www.vmem.com/support	5, 6, 7

Vendor	Storage Array Model	Array Firmware	Notes
Whiptail	ACCELA Storage arrays	Refer to storage vendors support list, http://whiptail.com/hcl-certifications	6, 7
X-IO	X-IO: Emprise 5000, ISE-2, ISE-1, Hyper-ISE	Refer to storage vendor's support list, http://xiostorage.com/products/product-certifications/	6, 7

Enterprise Backup Interoperability Matrix

Table 8 lists the tape library models that are supported with this version of the Unified Computing System.



Note Direct connect of tape libraries is not supported.

Table 8: Tape Libraries

Vendor	Tape Library Model	Backup Application	Firmware	Notes
Oracle	SL500: HP LTO6	Symantec Netbackup 7.5	Refer to storage vendor's support list, http://www.quantum.com/swcompguide.aspx	24
Oracle	SL500: HP LTO3, HP LTO4, HP LTO5, IBM LTO3, IBM LTO4, IBM LTO5 SL3000: HP LTO3, HP LTO4, HP LTO5, IBM LTO3, IBM LTO4, IBM LTO5, T10000C	EMC Networker 7.6.3		24
Quantum	i500: HP LTO4, HP LTO5, IBM LTO4, IBM LTO5 DXi6702: All supported emulated tape drives	Symantec Netbackup 7.5	Refer to storage vendor's support list, http://www.quantum.com/swcompguide.aspx	24
IBM	TS3584: IBM LTO4	TSM 6.3		24

Switch Interoperability Matrix

Table 9 lists the switch models that are supported with this version of the Unified Computing System.

Table 9: Switch

Vendor	Switch Model	Notes
Brocade	6510: 7.0.2, 7.0.1b, 7.0.1a	2, 16
Brocade	300, 5100, 5300: 7.0.2, 7.0.1b, 7.0.1a, 6.4.3b, 6.3.2e, 6.3.2d, 6.4.2b, 6.4.1b	2, 16
Brocade	4100, 48000, 4900, 5000: 5.2.1, 6.4.1b, 6.3.2a, 6.4.2a	2, 16
Brocade	7500: 5.2.1, 6.4.1b	2, 16
Brocade	DCX-4S: 7.0.2, 7.0.1b, 7.0.1a, 6.4.3b, 6.3.2e, 6.3.2d, 6.4.2b, 6.4.1b	2, 16
Brocade	DCX: 7.0.2, 7.0.1b, 7.0.1a, 6.4.3b, 6.3.2e, 6.3.2d, 6.4.2b, 6.4.1b	2, 16
Cisco	MDS: 4.2.1b	1, 2, 9
Cisco	Nexus 7000: 5.2(4)	1, 2, 9
Cisco	Nexus 5000: 5.1(3)N1(1)	1, 2, 9

Notes

- 1 Code levels greater than or equal to these listed versions are deemed by Cisco to be interoperable
- 2 The storage vendor's interoperability matrix should be consulted for switch software revisions that are supported with their arrays
- 3 SAN Boot requires removal of all local disks
- 4 Ethernet or FC driver driver required for installation not included in OS image or not supported, download driver from cisco.com
- 5 Direct Connect FC/FCoE topologies are supported with this array
- 6 Non-default adapter settings are typically required by storage vendors and found on their support list
- 7 Multiple array and switch models supported with UCS by the storage vendor are on their support list
- 8 Fiber Channel Storage that supports equivalent adapters and drivers is supported by UCS. UCS M71KR-E is equivalent to Emulex LP2100x, UCS M71KR-Q is equivalent to QLogic QLE804x, UCS M72KR-Q is equivalent to QLogic QLE8142/QLE8152, UCS M72KR-E is equivalent to Emulex OCe10102-F, UCS M73KR-Q is equivalent to QLogic QLE8242, UCS M73KR-E is equivalent to Emulex Oce11102-F
- 9 Specific switch features must be supported on all switches used, please consult release notes for each switch product to determine support for the desired feature
- 10 To upgrade to ESXi 4.0 U3 or U4 and ESXi 4.1 U1 or U2, 1st install the Cisco custom image (refer to note 23)

- 11 The listed driver version is required for support. The updated Driver CD is available from VMware.com > All Downloads > VMware vSphere > Drivers & Tools > VMware ESX/ESXi Driver CD for Cisco FCoE CNA
- 12 iSCSI Boot is supported. Both Appliance and Uplink ports are supported unless otherwise noted
- 13 Refer to the UCS 82598KR-CI 10-Gigabit Ethernet Adapter for the supported Ethernet driver on this Converged Network Adapter
- 14 The 2104XP fabric extender is not compatible when the 1240 and 1280 or 1240 and VIC Port Expander are combined on the same blade. Applies to B22 M3 and B200 M3 as well as B420 with 1240 and adapter slot 2 populated
- 15 Listed support is for the UEK Kernel only. For the RHEL compatible kernel, please refer to the RHEL versions
- 16 UCS Fabric Interconnect must be in FC End Host mode and upstream switch NPIV enabled
- 17 iSCSI Boot is not supported on Windows Server Core
- 18 All variations of OS are supported for this version, including Windows Server Core, Full and Hyper-V Standalone
- 19 The listed driver version is required for support. They are contained in the Cisco OEM Customized ESX ISO available from VMware.com > downloads > VMware vSphere Hypervisor (ESXi)
- 20 For C-Series servers managed by UCSM, unless stated in this document, the supported firmware and drivers will be in the 1.4(7) and 1.4(8) C-Series Compatibility guide
- 21 2.0(4) and 2.0(5) B-Series Product bundle is supported with the 2.1(1) Infrastructure bundle. When running this combination, the 2.0(4) or 2.0(5) Interoperability guide would be applicable
- 22 Listed support is for the UEK Kernel only. For the Red Hat Enterprise Linux compatible kernel, please refer to the Red Hat Enterprise Linux versions
- 23 Updating drivers during OS installation is currently not supported by Oracle. Update driver after the OS has completed installation.
- 24 Enable FCP error recovery in the service profile adapter policy
- 25 NIC Teaming driver is not currently supported with Hyper-V
- 26 Supported for NFS/CIFS access via appliance ports on the Fabric Interconnect
- 27 Requires Errata RHSA-2013:0168, reference bugzilla 884740
- 28 Requires Oracle OVM patch 14775391
- 29 For SAN installations, requires Citrix Hotfix CTX137403 during the installation, after installation is complete, apply CTX138731
- 30 Windows 2012 NPIV Host Feature is not supported
- 31 ESX NPIV Host Feature is not supported
- 32 Legacy Mode support only
- 33 Virtual Storage appliances are NPV Mode only. These require target to target zoning in direct connect topologies which is not supported with UCSM zoning.
- 34 Only UCS 785GB MLC Fusion-io ioDrive2, serial numbers FIO1717xxxx and newer, are supported on the B420 M3

Related Documentation

Links to the latest versions of related Cisco documentation are available in the *Cisco UCS Documentation Roadmap*, located at:

<http://www.cisco.com/go/unifiedcomputing/b-series-doc>

Cisco UCS B-Series Blade Server Software is available for download here:

<http://www.cisco.com/cisco/software/type.html?mdfid=283853163&flowid=25821>

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Configuration Summary

STORAGE TIERS

DRIVES	MODEL	RAID	HA	WORKLOAD	R/W	RAW	USABLE	MB/S	IOPS	RT (ms)	
1x SSD	16	1.02TB SSD SFF	Raid 5 (7+1)	Random	8k	50% Read	27.92 TiB	21.67 TiB	997	127,657	0.7

USABLE CAPACITY WITH COMPACTION



21.7 TiB

RATIO
4 to 1

86.7 TiB



RAW CAPACITY

27.9 TiB


USABLE CAPACITY

21.7 TiB

MAX FRONT-END IOPS *

127,657

USABLE VS. OVERHEAD



SUMMARY

Model:	8200	SFF Enclosures:	1
Nodes:	2	Free Slots:	6
1Gb Eth RCIP Ports:	2	LFF Enclosures:	-
FC 8Gb Ports:	-	Free Slots:	-
FC 16Gb Ports:	4	Total Disks:	16

PHYSICAL SPECIFICATIONS

Total Rack Space:	2U
Total Weight (without rack):	26.7 kg
Total BTU (balanced):	1419
Total Power (balanced):	0.416 kW

* Estimate Only

Performance Summary


STORAGE TIERS

DRIVES	MODEL	RAID	HA	WORKLOAD	R/W	MB/S	IOPS	RT (ms)
SSD 16	1.92TB SSD SFF	Raid 5 (7+1)	<input type="checkbox"/>	Random 8k	50% Read	997	127,657	0.7

SSD: 16 x 1.92TB SSD SFF
RAID 5 - Random 8k 50% read/50% write



Configuration Layout



■ 1.92TB SSD SFF

Bill of Materials

Part Number	Description	Qty
K2Q35A	HP 3PAR StoreServ 8200 2N Storage Base	1
K2P89A	HP 3PAR 8000 1.92TB SAS cMLC SFF SSD	16
L7B45A	HP 3PAR 8200 OS Suite Base LTU	1
L7B46A	HP 3PAR 8200 OS Suite Drive LTU	16



Hewlett Packard
Enterprise

HPE 3PAR StoreServ Storage: designed for mission-critical high availability

Modern Tier-1 storage for the new style of Business

Contents

Executive summary.....	3
Achieving native HA	3
HPE 3PAR Architecture	4
Native end-to-end error checking.....	4
Native End-to-End Error Checking.....	5
No single point of failure	5
Zero downtime for maintenance and upgrade services.....	6
Persistent Ports feature	6
Rapid recovery from failed disk devices.....	6
Traditional HA vs. HPE 3PAR StoreServ resiliency features.....	7
Flexible and resilient enterprise RAID levels.....	8
Automated bypass control and hot-plugging mode.....	8
Drive enclosure availability.....	8
No cache data loss in any scenario.....	8
Persistent Cache feature	9
Persistent Checksum.....	9
Agile QoS management on HPE 3PAR StoreServ Storage	9
High availability across sites using Peer Persistence.....	10
Disaster tolerant solutions with HPE 3PAR Remote Copy software.....	10
Flat backup for end-to-end availability and data protection	11
Federation and non-disruptive data migration with HPE 3PAR Peer Motion software.....	11
Customer-configurable priority for snapshot promotion	12
Universal management interface and software functionality	12
Proactive support services and remote operations.....	12
Automated proactive utilization monitoring.....	13

Executive summary

In the modern data center, the quest for predictability and order has given way to the need for flexibility and the ability to address the unexpected. New service delivery models like IT as a service (ITaaS) are driving new requirements around storage agility, efficiency, and resiliency.

The marketplace is filled with “me too” traditional SAN storage suppliers, but only HPE offers the industry’s broadest experience delivering solutions and services to enable business continuity while maintaining flexibility of IT infrastructure. For example, HPE 3PAR StoreServ Storage features a modern, purpose-built architecture that spans from low to high and encompasses both traditional spinning media and all-flash capacity, delivering new innovations to storage that are far beyond traditional SAN storage environments. HPE 3PAR StoreServ Storage removes the last barrier to consolidation by delivering assured quality of service (QoS) levels without having to physically partition resources or maintain discreet storage silos.

HPE 3PAR StoreServ Storage is designed to help minimize operational impact by delivering planned resilience. Full hardware redundancy is built into the system, mitigating all single points of failure. When appropriately configured, HPE 3PAR StoreServ Storage is capable of greater than “five nines” or 99.999 percent availability from a hardware perspective. This is reinforced by autonomic and proactive error checking, and software features that deliver seamless failover/failback to help ensure complete system resilience, even when the unexpected happens.

A unique set of products power HPE 3PAR StoreServ Storage in delivering high availability (HA) and mission-critical resilience to multi-tenant cloud and ITaaS environments. HPE 3PAR Peer Persistence software keeps mission-critical applications running seamlessly, even in the event of a disaster, by enabling automated, transparent failover and failback, between both systems and sites. HPE 3PAR Persistent Cache removes performance impacts that can result from unplanned component failures, making it a must-have for maintaining service levels in the virtual data center. HPE 3PAR Persistent Ports allows non-disruptive upgrades on HPE 3PAR StoreServ Storage without relying on multi-pathing software and without initiating failover. HPE 3PAR Persistent Checksum ensures end-to-end data integrity, protecting against silent corruption from the host to the storage array. HPE 3PAR Priority Optimization software provides assurance of predictable and consistent service levels by providing storage controls that allow applications to get the right level of performance.

Achieving native HA


All HPE 3PAR StoreServ Storage models offer built-in, native hardware and software availability that combine to create massively available Tier-1 systems. Many enterprise, financial, and hosted service providers have standardized on HPE 3PAR StoreServ Storage at a global level to offer mission-critical ITaaS services. This technical white paper describes all of the HPE 3PAR StoreServ Storage hardware components and software features that natively support seamless, Tier-1 availability.

HPE 3PAR Architecture

The foundation for all HPE 3PAR StoreServ Storage systems, a single HPE 3PAR StoreServ Architecture combines superior open technologies with extensive innovations in hardware and software design. This architecture reduces class distinctions across storage tiers by offering a range of models to meet the needs of small to large enterprise data centers running key business applications, up through the enterprise-wide deployment of mission-critical applications and beyond. With a single architecture that meets mid-range to enterprise as well as all-flash storage demands, HPE 3PAR StoreServ Storage offers a common operating system, management experience, feature set, and Tier-1 data services across the entire product family.

The HPE 3PAR StoreServ Architecture features a high-speed, full-mesh backplane that joins multiple controller nodes (the high-performance data movement engines of the HPE 3PAR StoreServ Architecture) to form a cache-coherent, “mesh-active” cluster. This low-latency interconnect allows for tight coordination among the system’s controller nodes. Within this architecture, controller nodes are paired via either Fibre Channel (FC) or SAS connections (depending on model) from each node in the pair to the dual-ported drive chassis (or drive enclosures) owned by each pair. In addition, each controller node may have one or more paths to hosts (either directly or over a network). The clustering of storage controller nodes enables HPE 3PAR StoreServ Storage to present to hosts a single-instance, highly available, self-optimizing, and high-performance storage system.

HPE 3PAR Mesh-Active design allows each volume to be active on every controller in the system



Native end-to-end error checking

HPE 3PAR StoreServ Storage offers automated end-to-end error checking during the journey of data frames through the system to the disk devices. These technologies have been put in place to always ensure that customer data frames does not get corrupted at any point in its journey through the array and where there may be corruption, detection is done promptly and handled appropriately. For example, the HPE 3PAR Thin Express ASIC comes with the Persistent Checksum feature that ensures end-to-end data protection, from host HBA to physical drives. In addition to this, Self-Monitoring Analysis and Reporting Technology (SMART) is used to provide predictive failure detection such that any disk device crossing certain SMART thresholds can be designated as a “predictive failure” by the storage node and identifying that device as in need of replacement before it actually fails.

Error checking includes, but is not exclusive to, the following layers within all HPE 3PAR StoreServ Storage systems:


- Cyclic redundancy checking (CRC) and parity checks on all internal CPU and serial buses
- Control cache and data cache error-correcting code (ECC) checks
- PCIe bus CRC checks
- HPE 3PAR Gen5 Thin Express ASIC Connection CRC checks
- Protocol FC/iSCSI/FCoE CRC checks at the frame level (hardware-accelerated via the systems’ HPE 3PAR Gen5 Thin Express ASICs)
- Disk devices CRC at the block level, both when the data lands and during the lifecycle of the data on disk
- End-to-end T10-DIF error checking at host HBA, storage adapter, data cache and disk drives. This feature is also called HPE 3PAR Persistence Checksum

HPE 3PAR Persistence Checksum works on all customer data while all other above checks work on customer data, metadata and program code.

HPE 3PAR StoreServ Storage also issues logical error status block (LESB) alerts if a frame arriving in the storage interface has CRC errors beyond a certain threshold. This indicates that a cable or component between the host and storage device needs replacing or cleaning.

With HPE 3PAR Remote Copy software, proactive error detection is also extended to replicated data to help ensure that potential cascaded data issues are not experienced. Additionally, HPE 3PAR StoreServ replication with Remote Copy includes a link pre-integration test to verify the stability of Remote Copy replication links in advance of adding these links for use with Remote Copy over an IP network.

Native End-to-End Error Checking



No single point of failure

HPE 3PAR StoreServ Storage components such as storage nodes, cache cards, disk- and host-facing host bus adapters (HBAs), power supplies, batteries, and disks feature 2N or better redundancy. As a result, system interruption does not occur as a result of failures with any of these components. The system's power domains are designed with redundancy to prevent power loss to back-end disk devices due to a drive enclosure power supply failure. HPE 3PAR StoreServ Storage systems offer up to four current load-balanced power distribution units (PDUs) per rack, providing a minimum of one and support for up to two separate data center power feeds to offer even more power redundancy and protection against events such as brownouts.

All storage nodes run a separate copy of HPE 3PAR Operating System (OS) software. The software is both statefully managed and self-healing in the case of a process's or processes' failure and restart, across all cache-coherent active/active storage controller nodes.

Zero downtime for maintenance and upgrade services

With HPE 3PAR StoreServ Storage, all maintenance and upgrade activities are performed online and completely non-disruptively. This includes hardware replacements and upgrades as well as HPE 3PAR Operating System software updates and patches. Component firmware at all levels can be upgraded dynamically (en masse or individually), including the disk firmware, as all disk devices in the HPE 3PAR StoreServ Storage system are dual-ported.

Intuitive functionality such as **checkupgrade** and **checkhealth** functionality helps confirm that an upgrade can be completed successfully and without risk when performed prior to an HPE 3PAR OS upgrade. These functions check firmware levels, system services, and system status in terms of any degraded or failed components that may need replacing. In addition, the health check can be executed on an ongoing basis to obtain comprehensive information on the health of the system.

Proactive system monitoring and support services can also be used to remotely apply software upgrades and dynamically examine a system at the customer's request, helping ensure successful upgrade experiences when change control windows are limited.

Persistent Ports feature

When upgrading controller node software, the HPE 3PAR StoreServ system gives you the ability to upgrade one controller node at a time, pausing between nodes to confirm that multi-pathing is active. HPE 3PAR Persistent Ports is an OS feature that provides greater resiliency and decreases the need for host multi-pathing software.

By using HPE 3PAR Persistent Ports in SAN connect environments where the FC SAN switch to which the HPE 3PAR StoreServ Storage system is connected supports N_Port ID virtualization (NPIV) capability, you can perform software upgrades on a node-by-node basis without any reliance on host multi-pathing software, thereby avoiding any path failovers. The HPE 3PAR Persistent Ports feature also works with iSCSI and FCoE ports and allows non-disruptive and online upgrades to take place completely transparent to hosts.

In addition to enabling non-disruptive and online software upgrades, HPE 3PAR Persistent Ports also protects applications against any possible outage due to the loss of a controller node. Regardless of the cause for this loss—planned upgrade/maintenance, power failure, and others—the Persistent Ports feature allows seamless failover of all the ports on the impacted node to designated partner ports on a different node. This is seamless to hosts, which continue to use the original ports assigned while in reality IO is transparently re-routed through the partner port. This helps organizations maintain uninterrupted service levels for their applications running on HPE 3PAR StoreServ Storage and protects against any service level penalties that might apply. For Fibre Channel attached hosts, HPE 3PAR Persistent Port functionality also provides transparent and uninterrupted failover in response to “loss_sync” events, which may be triggered as the result of a failure between an array HBA and the switch.

Rapid recovery from failed disk devices

The most common hardware failure for storage arrays in general occurs in the disk devices themselves. In the event of a disk failure, it is imperative that rebuild takes place quickly and with little or no impact to system performance. HPE 3PAR StoreServ Storage systems utilize unique capacity virtualization technology to help ensure that all potentially impacted logical unit numbers (LUNs) associated with a physical disk device failure are rebuilt from a “degraded” mode to full parity protection in the fastest possible time.


On the HPE 3PAR StoreServ Storage system, all disk devices are virtualized into “chunklets.” The system tracks utilization and writes allocation for each chunklet at all times. All disk devices have a number of “spare” chunklets that are reserved for sparing purposes.

The system uses a global spare chunklet pool that spans all physical drives in the system instead of more limited and traditional hot-sparing concepts. This global pool is comprised of a specific amount of spare capacity, regardless of the size and type of disk devices in the system. If the system is upgraded, an embedded algorithm helps ensure that the global chunklet sparing pool grows in proportion to the capacity installed. If a disk device fails, an automated RAID rebuild occurs to spare chunklets distributed across all disk devices, resulting in an expedited and extremely rapid rebuild window. Traditional storage rebuild models involve a few disks in a disk RAID group rebuilding to a hot spare, which can take significantly longer for RAID rebuild and result in I/O hotspots.

The massive parallelism resulting from global sparing helps ensure limited impact on application performance. In addition, HPE 3PAR StoreServ Storage disk rebuilds are “thin-provisioning aware”, so that only chunklets that have been written to are rebuilt. This also accelerates the rebuild time to full parity. For example, if a 4 TB disk fails but only 300 GB of written data resides on it, only 300 GB needs to be rebuilt to return any affected LUNs back to full RAID parity.

The system checks the Worldwide Name (WWN) of each replaced disk device to help ensure that the failed disk has not accidentally been re-inserted, thus help minimizing extended service times.

Traditional HA vs. HPE 3PAR StoreServ resiliency features



* Hotspots occur during rebuild and/or after disk replacement.

Flexible and resilient enterprise RAID levels

HPE 3PAR StoreServ Storage supports the following RAID levels: RAID 0, 1, 5, and 6 (multi-parity).

For RAID 5, the following RAID set sizes are supported concurrently (2 + 1 through 8 + 1, inclusively).

For RAID 6, the following RAID set sizes are supported concurrently (4 + 2, 6 + 2, 8 + 2, 10 + 2, and 14 + 2 inclusively), regardless of the disk device types installed.

For RAID 1, multiple mirror depths are supported on a LUN by LUN basis. From traditional RAID 1 (data + mirror) to mirror depth 4 (data + mirror + mirror + mirror), represents an unprecedented availability. With this RAID level, 75 percent of the disk devices in the system can be lost, yet any LUN configured with mirror depth 4 will still be online and accessible by hosts and applications.

HPE 3PAR StoreServ Storage supports the ability of administrators to change RAID level, RAID stripe size, device type, and availability level on the fly. In addition, based on user-specified policies, the system allows new storage resources to be dynamically shared as necessary—for example, additional storage controller nodes and disk devices. This helps ensure that the right applications receive the right levels of protection, HA, and performance, which is especially important when consolidating existing or rolling out new applications in an always-on infrastructure.

Automated bypass control and hot-plugging mode

HPE 3PAR StoreServ Storage supports switching intelligence to automate isolation of disk devices and drive enclosure SAS adapters from the SAS backend to which they are connected once a specific error threshold is reached. This process isolates the faulty component and protects other devices on that backend from other potential failures.

In addition, when a faulty disk or component is to be replaced, it is automatically placed into “hot-plug” mode. There is a visible LED on all devices to indicate this mode and reduces the risk of replacing the wrong device. Hot-plug mode also fences and isolates power to the component that needs to be replaced, which keeps power spikes from happening as a result of the component being removed or replaced.

Once the replaced or upgraded component is swapped out or added, the replacement routine automatically performs a series of checks to validate the new component. If required, the firmware on the new component can also be automatically and non-disruptively upgraded prior to activation.

Drive enclosure availability

The default provisioning for all HPE 3PAR StoreServ Storage systems includes high availability across drive enclosures for every LUN provisioned (This is the HA-cage or HA-enclosure mode.). This default protects from and helps ensure data availability in the event of the loss of an entire disk chassis or enclosure. Even in the event of an entire disk enclosure loss, the hosts do not lose access to the provisioned LUNs and applications still stay online. Since not all data may be of the same importance, this functionality is dynamically configurable per LUN. The minimum disk availability is magazine level. As the system scales, the availability level per LUN can be non-disruptively changed to disk-shelf availability as desired.

No cache data loss in any scenario

The protection of customer data is the number one priority on HPE 3PAR StoreServ Storage. HPE 3PAR StoreServ Storage supports redundant batteries to help ensure that the contents of write cache are written to specific and secure disks if a site power failure occurs. There is no pinning of data in cache, which puts data at risk when power is not restored in a specified period of time. Because cache data is written to disk, there is no limit with regard to the time to data recovery for cases where power is down for an extended period. Once power is restored, the cached data that was written to the secure disks is restored to its original location as part of the power-fail recovery process.

Persistent Cache feature

All storage arrays grow (no matter how efficient they are) as a result of increasing amounts of data and the need to keep that data secure and active for compliance purposes.

Unlike many competitive arrays that can only scale to two storage controller nodes or storage processors, HPE 3PAR StoreServ Storage can scale non-disruptively from two to eight cache-coherent active/active storage controller nodes (depending on model). As soon as the system detects that four or more storage controller nodes are installed, a resilience feature called Persistent Cache is autonomically enabled.

The Persistent Cache feature helps ensure that no storage controller node is placed into performance-limiting “cache write thru” mode as a result of a losing its partner in the node pair. Any node that loses its adjacent node can dynamically form a mirrored cache relationship with another storage controller node. This helps ensure that most efficient performance requirements are met and service levels remain unaffected during the failure or maintenance of a storage controller node.

Note

Other vendors that only scale to two storage controller nodes support a feature that enables the customer to keep mirrored cache mode enabled on a single storage controller in case of the loss or maintenance of the other storage controller node. For an enterprise storage customer, this represents an extreme risk to data as the system will now have a single point of failure. This creates the potential for data loss should the surviving controller node fail. Other vendors will quote mean time between failures (MTBF) of the second storage controller node as being such that the risk of a second error on the surviving storage controller node is minimal. However, if the risk value is not zero, there is usually a risk of dual failure and hence data loss. The biggest risk when running with a single point of failure is not component failure, but user error.

Persistent Checksum

HPE 3PAR Persistent Checksum addresses media and transmission errors that can be caused by any component in the I/O stack and provides additional protection above CRC transmissions for Fibre Channel protocol. Persistent Checksum is application independent and offers elaborate host OS support. When using supported HBAs, T10-DIF2 tags are added and verified on all customer data packets which are also checked and verified in the HPE 3PAR StoreServ Storage sub-systems like data cache and disk drives. When using unsupported HBAs T10-DIF tags are added and verified on the back-end HBAs and/or array target ports, storage data cache and disk drives, making the data secure all the way from the hosts right to the drives.

Agile QoS management on HPE 3PAR StoreServ Storage

HPE 3PAR Priority Optimization software for HPE 3PAR StoreServ Storage systems implements and manages a priority policy per Virtual Volume set (VVset) proxy for an application, or per virtual domain that serves as a proxy for a tenant. HPE 3PAR Priority Optimization software enables end users to take full control of performances via min goals and max limits for IOPS and bandwidth, along with the ability to set latency goals for volume sets as low as 500 microseconds for the most important applications. If these goals are not met, the system automatically adjusts the service levels of lower-priority applications and workloads in order to help assure that necessary QoS levels for your highest priority applications are frequently maintained.

With the combination of Priority Optimization software and HPE 3PAR Virtual Domains, you can specify these thresholds to protect individual tenants as well—for example, to prevent a single tenant from monopolizing resources. These capabilities removes the last barrier to consolidation by allowing you to deliver assured QoS levels without having to physically partition resources or maintain discreet storage silos.

HPE 3PAR Priority Optimization software is flexible and easy to configure and monitor, and it requires minimal supervision from storage system administrators. In contrast, some competitive QoS implementations require that you assign a workload to a predefined priority level or that the array be logically partitioned to reserve part of the system resources to a particular workload. These solutions are inflexible and don't allow real-time enforcement. HPE 3PAR Priority Optimization software is resident on the storage system and runs completely on the HPE 3PAR StoreServ Storage system. There are no host-based components to install.

High availability across sites using Peer Persistence

HPE 3PAR Peer Persistence software enables HPE 3PAR StoreServ Storage systems located within a metropolitan distance to act as peers to each other for delivering a high-availability, transparent failover solution for the connected VMware® vSphere, Microsoft® Hyper-V, Microsoft Windows® clusters, Red Hat® Enterprise Linux® (RHEL), and HPE-UX. HPE 3PAR Peer Persistence software allows an array-level HA solution between two sites or data centers where failover and failback remains completely transparent to the hosts and applications running on those hosts.

Unlike traditional disaster recovery models where the hosts (and applications) must be restarted upon failover, HPE 3PAR Peer Persistence software allows hosts to remain online serving their business applications, even when the serving of the I/O workload migrates transparently from the primary array to the secondary array, resulting in zero downtime.

In an HPE 3PAR Peer Persistence configuration, a host cluster would be deployed across two sites and an HPE 3PAR StoreServ Storage system would be deployed at each site. All hosts in the cluster would be connected to both of the HPE 3PAR StoreServ Storage systems. These HPE 3PAR StoreServ systems present the same set of VVs and VLUNs with same volume WWN to the hosts in that cluster. The VVs are synchronously replicated at the block level so that each HPE 3PAR StoreServ Storage system has a synchronous copy of the volume. A given volume would be primary on a given HPE 3PAR StoreServ Storage system at any one time. Using Asymmetric Logical Unit Access (ALUA), HPE 3PAR Peer Persistence presents the paths from the primary array (HPE 3PAR StoreServ Storage system on which the VV is primary) as “active/optimized” and the paths from the secondary array as “standby” paths. Issuing a switchover command on the array results in the relationship of the arrays to swap, and this is reflected back to the host by swapping the state of the paths from active to standby and vice versa. Under this configuration, both HPE 3PAR StoreServ Storage systems can be actively serving I/O under normal operation (albeit on separate volumes).

Disaster tolerant solutions with HPE 3PAR Remote Copy software

HPE 3PAR StoreServ Storage offers the most complete set of Disaster Recovery (DR) solutions in the industry. With a rich set of features, HPE 3PAR Remote Copy software is a uniquely easy, efficient, and flexible replication technology that allows you to protect and share data from any application. Choose between four different replication options from the same license—and across mid-range, enterprise-flash, and all-flash arrays of any model.

Implemented over a native IP network (through the built-in GigE interface available on all nodes) and native Fibre Channel, users may flexibly choose one of four different data replication modes: **Asynchronous Streaming** or **Asynchronous Periodic** (for asynchronous replication), **Synchronous, Synchronous Long Distance**—to design a solution that meets their solution requirements from the perspective of recovery-point objectives (RPOs) and recovery-time objectives (RTOs).

With all four of these modes, HPE 3PAR Remote Copy software allows data to be mirrored between any two HPE 3PAR StoreServ Storage systems of any model, reducing the incompatibilities and complexities associated with trying to mirror between traditional vendors' mid-range, enterprise, or all-flash arrays. Source and target volumes may also be flexibly and uniquely configured to meet user's needs (for example, different RAID levels, thick or thin volumes or different drive types, and others). HPE 3PAR Remote Copy is “thin aware” in that it is able to replicate both thin and thick volumes by using TPVV target volumes to provide the same cost savings associated with thin-provisioned source volumes created with HPE 3PAR Thin Provisioning software.

Synchronous mode solutions are suitable for data centers within metro distances and have a recovery point objective of zero seconds. When combined with HPE 3PAR Peer Persistence these solutions also enables recovery time objective of zero seconds.

For asynchronous replication solutions, network bandwidth is efficiently utilized with **Asynchronous Periodic** mode. Changed data within an HPE 3PAR Remote Copy Volume Group is transferred only once—no matter how many times it may have changed—between synchronization intervals. Additionally, efficiencies in the initial copy creation of the target volumes that do not require replication of “zero” data across the replication network (regardless of target volume type, thick or thin) result in a faster initial synchronization and better network utilization. With support for replication over long distances, HPE 3PAR Remote Copy **Asynchronous Streaming** assures data protection against disasters by providing an RPO in seconds, whereas HPE 3PAR Remote Copy Asynchronous Periodic mode ensures RPO in minutes over extended distances. Asynchronous Streaming is suitable for customers that want near zero RPO without the additional host latency associated with the synchronous mode of replication.

Synchronous Long Distance mode delivers a disaster recovery solution across long distances with a potential for zero data loss RPO and an RTO of minutes. This is achieved with a replication configuration that uses three sites (**3DC**) to simultaneously replicate a virtual volume from the primary array in synchronous mode to an HPE 3PAR StoreServ Storage array located at a synchronous site (within a metropolitan area) and in asynchronous periodic mode to an HPE 3PAR StoreServ Storage array located at an asynchronous site (across a long distance). In addition to the HPE 3PAR Remote Copy connections from the primary array to the two backup arrays, a passive asynchronous periodic link is configured from the synchronous array to the disaster recovery array. Under the Synchronous Long Distance mode algorithm, the synchronous site intelligently tracks the delta set of I/Os that have been acknowledged to the host but which have not yet been replicated to the asynchronous site. In the event that a disaster takes the primary storage array down, the user has the flexibility to recover either from the synchronous site or the asynchronous site.

Flat backup for end-to-end availability and data protection

HPE 3PAR Peer Motion software enables the seamless movement of data between HPE 3PAR StoreServ Storage systems of any model, enabling one-click workload balancing. Up to four systems can be pooled into the same federation to support up to 60 PB of usable capacity and 10 Million IOPS. Multi-array, bi-directional, non-disruptive data mobility across systems enables you to load balance at will, refresh technology seamlessly, reduce costs of asset lifecycle management, and lower technology refresh CAPEX.

HPE 3PAR Peer Motion Software is a non-disruptive, do-it-yourself (DIY) data mobility tool for enterprise block storage that does not require any external appliance to be included in the data path, nor does it introduce any additional overhead on the host resources. Unlike traditional block migration approaches, the HPE 3PAR Peer Motion Software enables non-disruptive data mobility of storage volumes between any HPE 3PAR StoreServ system that is part of a storage federation, and without complex planning or dependency on extra tools.

Federated multi-array, bi-directional data mobility between HPE 3PAR StoreServ Storage systems is easy to implement and manage via the HPE 3PAR StoreServ Management Console (SSMC). Simple predefined workflows have been implemented within SSMC that give you the ability to move data—for example, moving all volumes associated with a host from one HPE 3PAR StoreServ Storage system to another or moving individual volumes for workload balancing purposes—all with just a single click.

Federation and non-disruptive data migration with HPE 3PAR Peer Motion software

HPE 3PAR Peer Motion software enables the seamless movement of data between HPE 3PAR StoreServ Storage systems of any model, enabling one-click workload balancing. Up to four systems can be pooled into the same federation to support up to 60 PB of usable capacity and 10 Million IOPS. Multi-array, bi-directional, non-disruptive data mobility across systems enables you to load balance at will, refresh technology seamlessly, reduce costs of asset lifecycle management, and lower technology refresh CAPEX.

HPE 3PAR Peer Motion Software is a non-disruptive, do-it-yourself (DIY) data mobility tool for enterprise block storage that does not require any external appliance to be included in the data path, nor does it introduce any additional overhead on the host resources. Unlike traditional block migration approaches, the HPE 3PAR Peer Motion Software enables non-disruptive data mobility of storage volumes between any HPE 3PAR StoreServ system that is part of a storage federation, and without complex planning or dependency on extra tools.

Federated multi-array, bi-directional data mobility between HPE 3PAR StoreServ Storage systems is easy to implement and manage via the HPE 3PAR StoreServ Management Console (SSMC). Simple predefined workflows have been implemented within SSMC that give you the ability to move data—for example, moving all volumes associated with a host from one HPE 3PAR StoreServ Storage system to another or moving individual volumes for workload balancing purposes—all with just a single click.

Customer-configurable priority for snapshot promotion

HPE 3PAR StoreServ Storage supports thousands of reservation less and QoS-aware snapshots (up to 500 per parent LUN). These snapshots are stateful, non-fragmenting, and “thin-aware”—meaning they are non-duplicative with respect to writes. In addition, any snapshot in the tree can be restored and promoted at any time while newer and earlier point-in-time (PIT) snapshots still remain valid. These snapshot features provide continued and active data consistency at all times, which is a must for all enterprise customers implementing a Tier-1 business solution.

When a LUN must be restored for any reason (corrupt database, file system resizing, and others), support for customizable prioritization order for snapshot promotes to the parent LUN that enable a restore to the desired recovery point as quickly as possible. This helps ensure that the customer can generally restore crash-consistent or application-consistent snapshots obtained by HPE 3PAR Recovery Manager products to a precise recovery point at all times.

Universal management interface and software functionality

HPE 3PAR StoreServ Storage supports a converged operating system (the HPE 3PAR OS) across all models, meaning that generations of HPE 3PAR StoreServ Storage systems—regardless of size, model, or deployment timeframe—can all be managed from the same interface. All software functionality is identical across all HPE 3PAR StoreServ models, which means that any administrator only needs one set of storage skills when managing HPE 3PAR StoreServ Storage.

In addition to a simple management interface, the use of HPE 3PAR OS across the entire family drastically reduces the risk of provisioning or management errors when utilizing different HPE 3PAR StoreServ models.

Proactive support services and remote operations

HPE 3PAR StoreServ Storage has a dedicated service processor server that is used mainly for the outbound transfer of specific array-based information required by HPE Services. This information includes system status, configuration, performance metrics, environmental information, alerts, and notification debug logs. No customer data is ever sent over these links.

The data sent is used by HPE Services teams to proactively monitor the array and contact the customer if potential issues are discovered. This results in the customer being proactively warned about a potential problem before the issue can occur. Once the customer is advised of the issue, parts are dispatched and non-disruptively replaced by trained HPE service personnel at the customer's convenience.

If the service processor stops transferring information for any reason, HPE customer support centers will receive an alert. The service processor is also the firmware database repository for all current and previous firmware levels (including HPE 3PAR OS, disk device, and disk chassis [AKA “enclosure”] firmware).

If remote access is needed for any reason, the customer can configure inbound access for OS upgrades, patches, and engineering access.

If the customer's data center does not permit “phone home” devices, then all alerts and notifications, are sent to the customer's internal support team. The customer can then notify HPE Services of any suspected issue over the phone or via the Web.

Automated proactive utilization monitoring

In addition to and concurrent with automated remote monitoring, alerting, and notification, HPE 3PAR StoreServ Storage also supports a complimentary support service for customers known as Over-Subscribed System Alerts (OSSA). This automated monitoring tool performs proactive utilization checks on key system elements using data that resides at HPE. This data is collected periodically from HPE 3PAR StoreServ Storage systems and sent to HPE. The intent is to provide clients with valuable information to keep their HPE 3PAR StoreServ Storage systems running more effectively at all times.

In addition, optional HPE 3PAR System Reporter software enables customers to set thresholds and alerts for performance on system components and even configure a custom set of alerts. New metrics are added dynamically as needed.

Current metrics include:

1. Storage node controller CPU utilization
2. Disk IOPS (per device type, to highlight potential spindle-bound applications and overloaded systems)
3. Number of host initiators per port (guards against over-subscription and non-tested configurations)
4. Total host initiators per system
5. PCIe bus bandwidth
6. Port bandwidth (per host- and disk-facing system ports)

Learn more at

hp.com/go/3PAR



Sign up for updates

★ Rate this document



© Copyright 2011–2015 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

4AA3-8316ENW, December 2015, Rev. 6



Hewlett Packard
Enterprise

Disaster-tolerant solutions with HPE 3PAR Remote Copy

Contents

Overview	3
Data replication and disaster tolerant solution challenges.....	3
Disaster-tolerant solution metrics	3
HPE 3PAR Remote Copy software	4
What's new for HPE 3PAR Remote Copy with HPE 3PAR OS 3.2.2?	4
HPE 3PAR Remote Copy replication modes.....	5
Remote Copy Link Layer support.....	7
Remote Copy over IP	7
Remote Copy over Fibre Channel	9
Remote Copy over Fibre Channel bridged over an IP network.....	9
Remote Copy connections	10
Latency support with Remote Copy	10
Remote Copy functionality.....	11
Remote Copy coordinated snapshots	11
Auto recovery policy	11
New creatercopygroup and setrcopygroup options for asynchronous streaming mode	11
Remote Copy volume group VLUN limitations.....	12
System sizing considerations when using Remote Copy	13
Data replication link sizing.....	13
Remote Copy topologies.....	15
Remote Copy performance design.....	18
Solution benefits.....	19
Simplified implementation.....	19
Simplified management	19
Failover management	20
Transparent failover with HPE 3PAR Peer Persistence.....	20
Cross-product interoperability	21
Simple verification of replicated data	22
Access to production data without impacting replication.....	22
Recovery managers	22
Integration with cluster products	22
Summary.....	23

Overview

Today's IT organizations are faced with the difficult task of satisfying the diverse disaster-tolerant needs of the entire enterprise. Business-, government-, and industry-driven requirements compel the need to store more data and make it continuously available.

With the increased need for data availability comes an increased demand for fast disaster recovery, and an even stricter demand for very small recovery point objective (RPO) in the face of a data center disaster. Most types of replication do expose the user to some data loss in the event of a disaster.

An enterprise that is unable to recover its data assets quickly with minimal data loss following a disaster may be at risk for regulatory action, or worse, an inability to continue business. How do IT organizations protect more applications and data than ever while at the same time limiting the potential for data loss without adverse effects on server performance that can result from synchronous replication?

HPE 3PAR StoreServ Storage addresses all of these challenges by offering Remote Copy, a powerful yet simple solution for remote data replication that is the foundation of a properly designed and deployed disaster tolerant solution. New to the HPE 3PAR Remote Copy family with the release of HPE 3PAR OS 3.2.2 is Asynchronous Streaming Remote Copy software, which provides the capability to deploy replication solutions that can provide an RPO measured in seconds.

Data replication and disaster tolerant solution challenges

Even after natural or human-induced disasters that drastically affect day-to-day operations, businesses must continue to function. Compliance with business standards, industry trends, or federal regulations may place additional requirements on an enterprise looking to create or expand disaster recovery. Maximum RPO is one such important requirement.

For some organizations, adequate funding for disaster recovery is difficult to obtain because it may be perceived as an added expense for a very limited subset of corporate data. Clearly articulating how and why disaster recovery is necessary to meet the requirements, put in place by management or the Federal Government, is paramount. This helps to deploy a workable solution that meets expectations for defined RPO and recovery time objective (RTO).

Many storage administrators believe that by simply replicating data from the primary data center to a backup data center, they have fulfilled the enterprise's requirement for disaster tolerance. This couldn't be farther from the truth. A proper disaster-tolerant solution is a combination of technologies, software, and processes. These that are combined into a solution designed to meet a defined goal for RPO and RTO, and not just a technology that replicates data from one location to another.

Most probably, you'll agree that planning and implementing a disaster recovery solution is one of the most complex, time-consuming, and expensive projects that any enterprise will undertake. Designing one that meets a very small RPO can make the objective even more daunting.

Disaster-tolerant solution metrics

The primary metrics on which a disaster-tolerant solution is designed and measured are RPO and RTO. Even though for many enterprises RTO is not a very important metric (for some it is but generally not), the maximum amount of data that can be lost in the event of a disaster, RPO is very important. At the very core of any disaster-tolerant solution is the ability to ensure that data loss will not exceed the defined RPO for the solution following recovery from a disaster.

RPO is a definition of the maximum amount of data that can be lost in the event a disaster occurs. RPO is generally defined as an amount of time and not a given quantity of data. For example, an RPO of two hours would guarantee that following a disaster that occurs at 2 pm, once recovery completes at the disaster recovery site, all transactions that were committed up to and including 12 pm would be present in the replicated copy of the database. The RPO could, in fact, contain transactions committed after 12 pm but the defined two-hour RPO guarantees transaction committed up to 12 pm will always be present. With HPE 3PAR Asynchronous Streaming Remote Copy, solutions can be designed to meet requirements with RPO definitions measured in seconds without imposing negative performance implications on the servers whose data is being replicated.

RTO is a definition of the amount of downtime that may elapse following a disaster before the database needs to be up and running, and consuming new transactions. In the vast majority of disaster-tolerant designs, this is secondary to the amount of data loss that can be tolerated. A good disaster-tolerant solution will have some type of automatic recovery mechanism, such as clustering, which is integrated with the data replication software to both automate and limit the amount of downtime suffered following a disaster. In most cases, RTO is driven by the amount of time it takes to initiate a failover of the database to the disaster recovery data center, restart the database and have it complete crash recovery.

HPE 3PAR Remote Copy software

Traditionally, enterprise and midrange array-based replication solutions have not been interoperable, which leads to complex high-cost replication solutions. It also increases the knowledge base that must be available to the local staff or the hired consultants responsible for maintaining the replication solution. Not providing interoperability between enterprise and midrange storage often leads administrators to implement storage arrays at both sites for every array protected at the primary site in a 1:1 ratio of array type. The problem with this approach is that the cost of the storage is tied closely to the number of storage arrays being protected instead of the amount of data being protected.

HPE 3PAR Remote Copy supports replication of data between all array models with one exception. Support for Asynchronous Streaming mode on HPE 3PAR OS version 3.2.2 is provided between all HPE 3PAR StoreServ models except for the 7200 and 7200c models. Support of Asynchronous Streaming mode on 7200 and 7200c models will be supported on a following release. Contact your HPE representative for details on the support of Asynchronous Streaming Remote Copy with HPE 3PAR StoreServ array models 7200 and 7200c.

Note

Contact your HPE representative for details on support of Remote Copy Asynchronous Streaming mode on 7200 and 7200c models of the HPE 3PAR StoreServ arrays.

Implemented using either the native IP network interface on the HPE 3PAR arrays or native Fibre Channel SAN fabric (or via Fibre Channel extended over IP via FC-IP routers), users may flexibly choose from three different data replication modes to design a solution that meets their requirements for RPO, RTO, and cost.

HPE 3PAR Remote Copy allows you to replicate data between any two models HPE 3PAR StoreServ Storage systems. Eliminating the incompatibilities and complexities associated with trying to replicate between other vendors' midrange and enterprise array technologies. Source and target volumes may also be flexibly and uniquely configured to meet user's needs (for example, different RAID levels; thick, thin, or deduped volumes; or different drive types).

What's new for HPE 3PAR Remote Copy with HPE 3PAR OS 3.2.2?

There is a lot of new and enhanced functionality as well as configuration support with HPE 3PAR Remote Copy in the HPE 3PAR OS 3.2.2 release. The new functionality and enhanced configuration support include:

- Support for asynchronous streaming replication—a new replication mode that can be used to design solutions that deliver an RPO measured in seconds compared to periodic asynchronous mode's minutes (see the section titled "[Asynchronous Streaming mode](#)" for details)
- Support for greater network latency for solutions using Remote Copy synchronous replication mode
- Support for synchronous replication mode for solutions that bridge Fibre Channel to IP to extend a SAN fabric (FCIP)
- "Certified" support for Remote Copy over Fibre Channel solutions (RCFC) that leverage dark fiber solutions. These certified solutions provide for increased latencies via HPE or Alcatel-Lucent-tested optical networks over dense wavelength division multiplexing (DWDM). These certified optical network solutions are supported for all Remote Copy replication modes

HPE 3PAR Remote Copy replication modes

HPE 3PAR Remote Copy software now offers a full set of features that can be used to design disaster-tolerant solutions requiring an RPO as small as zero to an RPO measured in seconds to an RPO measured in hours or days or even longer. Users can choose between the three different data replication modes offered, synchronous, Asynchronous Streaming, or Periodic Asynchronous replication to design the most cost-effective solution that meets their solution requirement for RPO and RTO.

Synchronous mode

In synchronous mode, host-initiated write data is mirrored to write cache on both the primary and the secondary StoreServ arrays before the write completion is acknowledged to the host. On the primary StoreServ array, data is mirrored across the cache of two nodes. The write request is then sent to the backup StoreServ array via a communication link.

The backup StoreServ array also mirrors the data in its cache (again, on two nodes) and then sends an acknowledgement to the primary system. The host write is acknowledged to the host server after the remote array's acknowledgement is received by the primary array (figure 1). As with all synchronous replication solutions, Synchronous Remote Copy provides an RPO of zero or no data loss.




Figure 1. Synchronous mode

When used with an RCFC transport (or Fibre Channel extended with IP via FC-IP routers) HPE 3PAR Synchronous replication utilizes a patented protocol that only requires a single round trip across the FC network to replicate data vs. the standard double round trip SCSI protocol found in most implementations. Synchronous Remote Copy does this by having the secondary array post a group of SCSI read requests to the primary array. As host I/Os are received by the primary array it simply responds to the previously posted read requests from the secondary array, hence saving one of the round trips found in a normal SCSI write request.

Synchronous Remote Copy is now supported on all network transports provided by Remote Copy, RCIP, RCFC, and FCIP. The maximum latency supported by Remote Copy when configured for synchronous replication mode is 5 ms of round-trip latency. When paired with the HPE and Alcatel-Lucent optical network replicating over DWDM replication distances of up to 500 km are now possible.

Periodic asynchronous mode

Periodic asynchronous mode is asynchronous replication that is ideal for environments that need to replicate data where the write I/O service times would be too large if the data were being replicated synchronously, and where an RPO of 10 minutes or greater can be tolerated. Periodic asynchronous mode insulates the host write I/Os from any replication latency (and target array write latency) resulting from the link latency or speed.

The network throughput requirements when using periodic asynchronous mode are not as stringent as they would be with synchronous or asynchronous streaming mode. This is because with periodic asynchronous mode the replication link speed can be sized close to the average data generation rate that occurs during the largest delta resync—providing cost saving on replication.




Figure 2. Periodic asynchronous mode

In periodic asynchronous mode, host writes are completed on the primary array and the host write is acknowledged as soon as the data is mirrored across two nodes. The volumes on the primary array and volumes on the secondary array are then resynchronized “periodically”—when scheduled or when resynchronization is manually initiated via the `syncdcopy` command.

In this example (figure 2), what the user gets on the secondary array would be an I/O consistent copy of the source data as it looked at 12:10 initially and then again at 12:20. The process would then repeat itself 10 minutes later (in this example, the resync interval has been set to 10 minutes) and the user will get an I/O consistent copy of the source data as it looks at 12:30, 12:40, and so on.

When using Periodic Asynchronous Remote Copy choose the largest delta-resync interval practical, which will allow the solution to meet your RPO. The larger the delta resync interval, the closer the replication network can be in Megabits to the “average” data generation rate of the replicated data. This saves overall network cost.

Asynchronous streaming mode

For synchronous replication solutions as link latencies increase so does the latency between a host write and array acknowledgement. In most cases, this puts a practical limitation on the link latency that can be used for synchronous replication. With the release of HPE 3PAR OS v3.2.2, support for the new HPE 3PAR Asynchronous Streaming Remote Copy is available. It is an excellent option when an RPO of zero is desired but where the write I/O latency associated with a synchronous replication solution cannot be tolerated.

Asynchronous Streaming Remote copy is, as the name implies, asynchronous replication of data between two HPE 3PAR arrays but it differs from Periodic Asynchronous Remote Copy. The difference is instead of replicating data periodically at a defined resynchronization interval (for example, every 10 minutes) Asynchronous Streaming Remote Copy constantly places I/O data from the primary array onto the replication link for delivery at the secondary array. Asynchronous Streaming Remote Copy does this by holding data in cache on the primary array while it is being replicated to the secondary array. Once receipt of the data at the secondary array has been confirmed to the primary, the cache space holding the data is freed up.

This continuous “streaming” of the data between arrays will allow solutions based on Asynchronous Streaming Remote Copy to provide RPOs measured in seconds or even less. Meanwhile, journal-based asynchronous replication solutions meet RPO requirements defined in multiple minutes or hours.

Asynchronous Streaming Remote Copy is perfect for environments, where very small RPOs are required. Or for environments, where synchronous replication is desired but replicating the link latencies are so large it will result in unacceptable write I/O response times for the data to be replicated synchronously. Deploying an Asynchronous Streaming Remote Copy solution in such an environment solves the write I/O response time problem while providing the smallest RPO possible short of zero.

Asynchronous Streaming Remote Copy solutions generally requires replication link speed to be sized within 95–99 percent of the maximum data generation rate. This ensures cache on the primary array is not saturated and very small RPO is delivered.




Figure 3. Asynchronous streaming mode

Remote Copy Link Layer support

In order to satisfy the varying business or technical needs of users, Remote Copy offers two native methods of connectivity: native Gigabit and 10GbE (RCIP) and native Fibre Channel over optical fiber (RCFC). With suitable FC-IP routers, FCIP is also supported (with FCIP Remote Copy over Fibre Channel can utilize a SAN fabric that is extended over an IP network to provide long-distance Remote Copy support over IP).

Remote Copy over IP

Remote Copy over IP (RCIP) is a native IP implementation of Remote Copy over Ethernet. Every HPE 3PAR StoreServ array comes with a standard Gigabit or 10GbE port on every node that is dedicated for RCIP. RCIP is only supported for synchronous and periodic asynchronous replication solutions. It is ideal for short, medium, and long-haul replication. It is most often used by organizations that require disaster recovery for small to medium environments with only moderate data generation rates and a limited number of volumes to be replicated.

Synchronous mode is supported on IP networks with up to 5 ms of round-trip latency. Periodic asynchronous mode is supported on IP networks with up to 150 ms round-trip latency. Check with your HPE representative for Asynchronous Streaming Remote Copy support via RCIP.

Note

Check with your HPE representative in regards to Asynchronous Streaming Remote Copy support via RCIP.

RCIP uses multiple links between arrays (minimum of two and up to eight) to maximize bandwidth and ensure availability. By offering native IP connectivity, along with the site's existing IP infrastructure, RCIP is quicker to implement than solutions that don't offer native IP replication.

When using RCIP, it is recommended that the network allocated to Remote Copy is a dedicated network (or dedicated VLAN carved out of a site network) and that it not share network bandwidth with other applications. If the Remote Copy solution is deployed on a shared network, there is no way to ensure the solution can meet the defined RPO when using asynchronous replication. Poor performance on a shared network is a possibility even in synchronous replication.

For synchronous solutions, it is possible that non-Remote Copy traffic using bandwidth on the network will result in higher than expected write I/O latencies for the host I/Os being replicated synchronously across the shared network. For solutions based on Periodic Asynchronous Remote Copy, sharing the network with other applications may prevent data replication within the chosen resync interval. This results in missing the RPO target of the solution. The system log reporting will also send out messages that the delta resync interval has exceeded.

The use of VLANs to dedicate a given bitrate to RCIP is supported and is, in fact, recommended over having Remote Copy share bandwidth on the network with other applications.

While RCIP is compatible with partner solutions that optimize bandwidth by compressing data on the network (making a 100 Mbps network look like a 400 Mbps network, for example, by compressing 400 MB of data so it fits on a 100 Mbps link) the nature of the data being replicated can affect the effectiveness of these tools. For example, if the data being replicated is poorly compressible or not compressible at all, then the WAN optimizer will not provide the network bandwidth the solution was sized for resulting in the solution not achieving its designed for RPO.

Remote Copy does not cause this. It is a result of the data being generated and whether the WAN optimizer can do a good job of compressing that data or not. It can also result in situations that prevent Remote Copy from replicating data within the chosen resync interval, eventually missing the RPO target of the solution. The system log reporting will send out messages that the delta resync interval has exceeded.

It is also possible where dropped Remote Copy heartbeat messages and/or TCP retry messages show up in the system log. This happens because the WAN accelerator drops IP packets when its buffers overflow and it has to ask Remote Copy to resend the IP packets.

In situations where the anticipated amount of data to be replicated is high, FCIP (RCFC extended over a long distance WAN using FC-to-IP routers) can be a better choice than RCIP. HPE recommends the use of FCIP.

Note

For environments with high data replication rates, where replication over an IP network is necessary, HPE recommends the use of FCIP over RCIP.

In summary, RCIP supports:

- Round-robin load balancing across all available RCIP links between a pair of HPE 3PAR StoreServ arrays
- GiGe and/or 10-GiGe connection between HPE 3PAR StoreServ systems through LAN/MAN/WAN switches
- Synchronous and periodic asynchronous modes
- Solutions requiring the replication of a moderate amount of data
- One RCIP connection per node (up to eight RCIP connections on an 8-node array)
- Can be used in conjunction with RCFC and FCIP
- All Remote Copy topologies
- Maximum network latencies of 5 ms round trip for synchronous replication and 150 ms round trip for periodic asynchronous modes

Remote Copy over Fibre Channel

For customers that choose Fibre Channel connectivity between arrays HPE 3PAR StoreServ Storage offers Remote Copy over Fibre Channel connectivity (RCFC). RCFC is most often used for shorter distance solutions such as a campus or a metropolitan. It is a solution that can support distances of up to 500 km if desired (5 ms maximum of round-trip network latency over the SAN fabric).

RCFC uses Fibre Channel connections (at least two) between arrays for availability as well as for increasing total available bandwidth for replication. Customers have the flexibility to use connections across any approved Fibre Channel fabric to create multiple hops between arrays. These hops can include any HPE fabric vendor-approved connectivity such as FC ISLs between buildings, fabric extension via long-haul ISLs, and more. The ISLs can be provided through long-wavelength GBICs between switches and wave division multiplexing solutions such as DWDM that provide extended ISLs.

Note

Refer to the HPE SAN Design Guide for details on supported SAN fabric switches, GBICs, and distances. The HPE SAN Design Guide can be accessed via SPOCK: h20272.www2.hp.com/spock/

In summary, RCFC supports:

- Load balancing across available links between a source target Remote Copy pair for links configured for the same replication mode (synchronous, asynchronous streaming, or periodic asynchronous)
- All HPE 3PAR supported SAN infrastructure as defined in the [HPE SAN design guide](#)
- Synchronous, asynchronous streaming, and periodic asynchronous replication modes
- Solutions requiring the replication of large amounts of data
- Up to four RCFC links per node (maximum of 32 per array)
- All Remote Copy topologies
- HPE and Alcatel-Lucent—certified DWDM extended optical networks
- Maximum network latencies of 5 ms round-trip latency for synchronous, asynchronous streaming, and periodic asynchronous modes

Remote Copy over Fibre Channel bridged over an IP network

In addition to support over native optical Fibre Channel networks, Remote Copy is also supported over Fibre Channel fabrics bridged via IP for extended distance (i.e., latency) support for high throughput solutions where RCIP is not recommended. In situations where latency and bandwidth are concerns, or where the anticipated amount of data to be replicated is high, FCIP (RCFC extended over a long distance WAN using FC-to-IP routers) can be a better choice than RCIP.

In summary, FCIP supports:

- Fibre Channel SAN networks that are extended over IP using HPE approved FC-IP routers
- Load balancing across available links between a source target pair and Remote Copy pair for links configured for the same replication mode (synchronous, asynchronous streaming, or periodic asynchronous)
- Solutions requiring high replication throughput
- Up to four RCFC links per node (maximum of 32 per array)
- All Remote Copy topologies
- Maximum round-trip network latencies of 5 ms for synchronous, 10 ms for asynchronous streaming, and 120 ms for periodic asynchronous modes

Remote Copy connections

3.2.2 Remote Copy supports:

- Four RCFC connections per node (for nodes configured with enough FC ports)
- One RCIP connection per node

With these increases it is possible for a single eight-node HPE 3PAR StoreServ array to have up to 40 Remote Copy links:

- 32 combined RCFC and FCIP connections (four per node) + 8 RCIP connections (one per node) = 40 total Remote Copy connections

The Remote Copy connections on an array can all be used between a single pair of HPE 3PAR StoreServ arrays in a one-to-one topology or they can be used to connect the array to multiple HPE 3PAR StoreServ arrays together in an “M-to-N” or a synchronous long distance (SLD) topology (more on these topologies in the section titled “[Remote Copy topologies](#)”). For a single pair of eight-node HPE 3PAR array, if both arrays have enough ports to support 40 connections then all 40 connections may be used. See table 1 for details on the number of RCIP and RCFC links supported on HPE 3PAR OS 3.1.3.

Table 1. RCIP and RCFC link count support for HPE 3PAR OS 3.1.3 and later

	PRE 3.1.3	3.1.3 AND LATER (7000, 8000, 10000, AND 20000 ONLY)
Maximum RCFC and FCIP ports/node	1	4
Maximum RCFC and FCIP ports/storage system	4	32
Maximum RCIP ports/node	1	1
Maximum RCIP ports/storage system	4	8

Latency support with Remote Copy

The maximum latency supported on the network used by Remote Copy has increased significantly with the release of HPE 3PAR OS 3.2.2. Table 2 shows the maximum supported latency values starting with HPE 3PAR OS 3.2.2.

Table 2. Remote Copy latency support starting with HPE 3PAR OS 3.2.2

Table 2. Remote Copy latency support starting with HPE 3PAR OS 3.2.2

REPLICATION MODE	TRANSPORT LAYER	LATENCY
Synchronous	RCIP	5 ms
	RCFC	5 ms
	FCIP	5 ms
Periodic asynchronous	RCIP	150 ms
	RCFC	5 ms
	FCIP	120 ms
Asynchronous streaming	RCIP	Not supported
	RCFC	5 ms
	FCIP	10 ms

Remote Copy functionality

Remote Copy coordinated snapshots

Starting with HPE 3PAR OS 3.2.2, Remote Copy now supports coordinated snapshots for VLUNs being replicated synchronously, with asynchronous streaming mode and with periodic asynchronous mode. Prior versions of the HPE 3PAR OS only supported coordinated snapshots for Remote Copy groups that were being replicated in synchronous mode.

Coordinated snapshots allow the storage administrator to, with a single “createsv –rcopy” command, create a snapshot of the VVs in a Remote Copy group on the source array and at the exact point in time a duplicate of those snapshots will also be created on the target array. By quiescing a database before creating coordinated snapshots the storage admin can create transactionally consistent snapshots on the primary and secondary arrays that are suitable for use for database backup or other work.

Auto recovery policy

For Remote Copy groups, HPE 3PAR Remote Copy software provides for an “auto recovery” policy to be set on each Remote Copy volume group. This policy **does not** define whether replication groups should be reversed and replication resume automatically following an RC Group failover. Rather, it defines whether or not replication should automatically resume when the replication links return following a complete link failure. The default behavior is **no_auto_recover**, which prevents automatic restart in case of recovery from a complete link failure. The default behavior allows the administrator to ensure that the network issue is adequately resolved before resuming replication. In the case of a “flapping” WAN link, this allows Remote Copy to remain quiescent rather than attempt multiple restarts across an unreliable network.

The administrator can override the default setting for “auto recovery” by choosing `auto_recover` via the `setrcopygroup` CLI command or via the IMC or SSMC, which allows setting the automatic restart of a Remote Copy group once the link between the local and remote site is recovered. This mode is useful when the administrator does not wish to manually restart the copy operations after a link failure.

New `creatercopygroup` and `setrcopygroup` options for asynchronous streaming mode

New “`async`” replication mode

With the addition of Asynchronous Streaming Remote Copy with HPE 3PAR OS 3.2.2, there is a new replication mode associated with the “`creatercopygroup`” command, the “`async`” option has been added to specify asynchronous streaming mode is desired. The `setrcopygroup` command also supports the new “`async`” replication mode for Remote Copy groups, where the storage administrator wants to set asynchronous streaming mode. Switching replication modes for a group does require that the group first be stopped from the “`stoprecoverygroup`” command. The group is first stopped, the mode is then changed with the “`setrcopygroup`” command, and the group is then restarted with the “`startrecoverygroup`” command.

New “`snap_freq`” subcommand for Asynchronous Streaming Remote Copy groups

When using the “`setrcopygroup`” command for a Remote Copy group running in asynchronous streaming mode there is a new “`snap_freq`” subcommand that specifies how frequently Remote Copy should automatically take coordinated snapshots of the VLUN members of the group. For groups in asynchronous streaming mode, Remote Copy will, on a regular basis (default is once an hour), create its own coordinated snapshots. These snapshots are used as resync points should all of the replication links fail and a delta resync become necessary.

Remote Copy will utilize these coordinated snapshots to get the primary and secondary arrays back into synchronization. The default interval for creating these automatic coordinated snapshots is once an hour but this interval can be increased or decreased with the “`snap_freq`” subcommand to the “`setrcopygroup`” command. It is recommended that care be taken if the “`snap_freq`” frequency is reduced down from the default one-hour interval to a smaller more frequent value as creating snapshots too frequently (especially for Remote Copy groups containing a lot of VLUNs) can add a substantial load to the array.

New “`period`” subcommand behavior for Asynchronous Streaming Remote Copy groups

For Remote Copy groups set for asynchronous streaming mode, the “`period`” subcommand operates differently than it does for groups in periodic asynchronous mode. With Periodic Asynchronous Remote Copy, the “`period`” subcommand specifies how frequently the VLUNs on the primary array should be delta-resynched with the VLUNs on the secondary array. A value of zero specifies they will never be delta-resynched (a manually initiated delta-resynch is required).

When a Remote Copy Group is started in asynchronous streaming mode (“async” mode), the “period” subcommand serves two purposes. First, if Remote Copy determines that it needs to suspend Asynchronous Streaming Remote Copy group(s), in reaction to an issue such as replication link throughout degradation for example, the “period” subcommand specifies the order in which Remote Copy will choose groups to suspend. Groups with larger “period” value definitions will be suspended first and will to be restarted last vs. groups with smaller “period” values. This provides the storage administrator with some QoS control on how Remote Copy behaves if the replication environment degrades for some reason.

The second purpose of the “period” subcommand is to specify how long, in minutes, Remote Copy should wait before trying to automatically restart an asynchronous streaming group. If, after waiting, the defined “period” Remote Copy determines that resources are still insufficient to restart the group it will wait another full period before trying to restart the group again. If a group is defined to have a period value of zero, it will be the first group to be suspended if necessary and Remote Copy will not attempt to restart this group automatically, a manual restart of the group will be required.

Note

An Asynchronous Streaming Remote Copy group with a “period” value of zero (0) will be suspended first and will not automatically restart. A manual restart of the groups will be required if it becomes suspended.

Remote Copy volume group VLUN limitations

Remote Copy volume groups define collections of volumes that are required to have write order consistency across them to ensure data I/O integrity. Use cases include multiple volumes written to by a database application or multiple LUNs used by a file system. Unless write order consistency is maintained across all the volumes associated with a database instance, recovery from a disaster by the application cannot be guaranteed at the target array. Remote Copy volume groups ensure this write order consistency. There are limits to the number of VVs that can be replicated via Remote Copy and limits on the number of VVs supported per Remote Copy volume group. The specific limits depend on the HPE 3PAR OS release and the replication mode chosen. Check table 3 for details on the total number of volumes and total volumes that can be included per Remote Copy volume group for HPE 3PAR OS release 3.2.2.

Table 3. HPE 3PAR OS 3.2.2 limits on number of replicated volumes per Remote Copy volume group based on replication mode and system node count

REPLICATION MODE(S)	2-NODE SYSTEM	4- TO 8-NODE SYSTEMS	MAXIMUM RC VOLUME GROUPS	MAX VVS PER RC VOLUME GROUP
Synchronous mode	2400	2400	2400	300
Asynchronous streaming mode	2400	2400	512	300
Combined synchronous and asynchronous streaming modes	2400	2400	2400 (512 may be for asynchronous streaming)	300
Periodic asynchronous mode	6000	6000	6000	300
Combined synchronous, asynchronous streaming, and periodic asynchronous modes	2400	2400	2400 (512 may be asynchronous streaming)	300

When mixing replication modes between a pair of HPE 3PAR arrays, the maximum number of volumes supported is defined by the replication mode with the lowest limit. For example, if you are mixing synchronous, asynchronous streaming, and periodic asynchronous modes between two arrays you do not get 2400 synchronous or asynchronous streaming volumes and 6000 periodic asynchronous volumes. You get a maximum (total) of 2400 volumes replicated between those two arrays because both synchronous and asynchronous streaming are limited to 2400 volumes.

If the array is in an MxN configuration, the maximum replicate volumes still has the same limits, a total of either 2400 or 6000. This means in an MxN topology, if an array has a Remote Copy relationship with four other arrays the maximum number of VVs it can replicate (total) is limited by the replication mode(s) it is using. If that array is replicating via synchronous mode with two arrays, asynchronous streaming with one array, and periodic asynchronous with two arrays, it can replicate a maximum (total) of 2400 VLUNs and that includes bidirectional replication.

System sizing considerations when using Remote Copy

When sizing HPE 3PAR arrays that will be used to replicate data using Remote Copy, care must be taken to account for the additional IOPS and workload that Remote Copy will impose on the arrays. This means that in addition to being sized to service the native workload generated by servers connected to it, the secondary array in a Remote Copy relationship must be sized to also service the additional write I/Os for the replicated data coming from the primary array. This is true for all replication modes—synchronous, asynchronous streaming, and periodic asynchronous.

Additionally, both arrays must be sized to account for the IOPS resulting from snapshots created by Remote Copy. In periodic asynchronous mode, snapshots are created on both the primary and the secondary arrays during every “period” interval. This means if a 10-minute “period” is specified for a Periodic Asynchronous Remote Copy group, the primary and secondary arrays will both create snapshots for all the VVs in that Remote Copy group every 10 minutes. This takes place before sending delta changes from the primary array to the secondary array.

The snapshots on the primary array will result in copy on write I/Os (COW I/Os) for all new host writes to the VVs in the Remote Copy group on the primary array. This array must be sized to account for the additional back-end write I/Os generated by these snapshots. On the secondary array, a snapshot of all the base VVs in the Remote Copy group is taken prior to the delta changes being sent from the primary array. As the delta changes coming from the primary array are applied to the base VVs on the secondary array, COW I/Os will occur on the snapshots for these volumes. The secondary array must be sized to account for the additional back-end write I/Os generated by the snapshots.

If you are using asynchronous streaming mode, Remote Copy will create coordinated snapshots at every “snap_freq” interval (default is once an hour). These coordinated snapshots will result in additional back-end I/Os on both the primary and secondary array. These additional I/Os must be accounted for when sizing the arrays. In synchronous mode, Remote Copy only creates snapshots if the replication links fail or if a Remote Copy group is suspended. So the impact of snapshots is not as pronounced although the user-initiated coordinated snapshots need to be considered.

If the CPG specified on either array for the Remote Copy snapshots is a CPG using nearline drives, performance will suffer as the COW I/Os to the nearline drives may not be serviced fast enough due to their limited performance. For this reason HPE recommends the administrator does not use nearline drives for snapshots that Remote Copy creates.

Note

HPE recommends that nearline drives should not be used for the snapshots that Remote Copy creates.

Data replication link sizing

In any live data replication based disaster-tolerant solution, the network used to replicate data between the primary and the disaster recovery sites is a key part of the solution. It may be the case that the solution is active-active where both sites run production and are expected to back one another up, so there is no “primary” and “secondary” site—both sites are primary for some RC groups and secondary for other RC groups. The replication network’s speed has a direct effect on the RPO and RTO of the solution, as well as an effect on the total recurring cost of the solution. It must be sized properly to ensure smooth operation of the entire solution, at all times, at the best cost to the enterprise.

Most solutions replicating data synchronously require a replication link where the speed is very close to the maximum data generation rate of the data generation (sometimes a solution sized for the 95th or 98th percentile will work but not always). For example, if there is a peak data generation rate of 400 MB that lasts a few minutes during the day, yet the average data generation rate for the rest of the day is closer to 250 MB, a replication link of 400 MB will be required to prevent server I/O latency spikes when the 400 MB I/O spike occurs.

The link must be sized to the maximum I/O rate spike, if not, queued data waiting to get onto the link will manifest itself as additional I/O latency. Think of it like a freeway with only two lanes versus a freeway with six lanes. For a given spike in traffic, you will wait in line much longer to get onto the two-lane freeway than to get onto the six-lane freeway, simply because the six-lane freeway can move more traffic.




Figure 4. Effect of replication link speed on synchronous replication when the data generation rate exceeds the replication link speed

For solutions using periodic asynchronous replication, the replication link can be sized to the average data generation rate for the period in which the most data is generated. For example, if the solution is set to a period of 10 minutes, if the 24-hour day is broken down into 10 minute periods and in the worst case 400 MB on average is generated during one of these periods the replication link needs to be sized to 400 MB to ensure the data from the worst case period can be replicated within 10 minutes before the next period starts. If the data cannot be replicated within the period length an error will be logged in the system log indicating that the remote copy group did not meet its resync interval.

With Asynchronous Streaming Remote Copy properly sizing the replication link becomes especially important as the solution is expected to provide for an RPO measured in seconds rather than minutes or hours. Delivering an RPO measured in a handful of seconds requires a replication link whose speed is very close to the write rate of the data that is being generated so that RC can ship data at close to the rate it is being generated.

This means that like a synchronous replication solution, a solution based on Asynchronous Streaming Remote Copy will require replication links that are sized very close to the maximum write data generation rate expected. At its first delivery, Asynchronous Streaming Remote Copy holds data that is being replicated in cache, on both the primary and secondary arrays. It is important to ensure that this data does not consume too much cache on either array.

If it does, Asynchronous Streaming Remote Copy will suspend Remote Copy groups in reaction to the overconsumption of array cache. Having properly sized data replication links will ensure constant and smooth streaming of data between the arrays, which in turn ensures the target RPO is met.




Figure 5 (a). Effect of replication link speed on Periodic Asynchronous Remote Copy cache consumption




Figure 5 (b). Effect of replication link speed on Periodic Asynchronous Remote Copy cache consumption

Remote Copy topologies

HPE 3PAR Remote Copy software can be deployed in different topologies depending on customer needs. Remote Copy supports M-to-N topologies. It also supports Synchronous Long Distance (SLD) that combines synchronous and periodic asynchronous replication to replicate a Remote Copy group to two separate target arrays. The M-to-N topology is very flexible and allows for some very complex Remote Copy solutions to be deployed.

One-to-one topology

Remote Copy’s simplest topology is a one-to-one configuration (it’s an M-to-N where $M = N = 1$), where a pair of HPE 3PAR StoreServ arrays replicate volumes between one another (figure 6). This topology supports disaster-tolerant and cluster-based scenarios between two geographically distinct data centers if desired. HPE 3PAR StoreServ Storage supports bidirectional replication in a one-to-one topology. Synchronous, Asynchronous Streaming, and Periodic Asynchronous replication modes may be used simultaneously in a one-to-one configuration, albeit for different RC groups and on different Remote Copy physical links.

The three modes can only be run simultaneously between two HPE 3PAR StoreServ arrays if they are both running HPE 3PAR OS 3.2.2 and later. When mixing different replication modes between two StoreServ arrays different physical transports must be used for each mode. For example, RCIP for the Periodic Asynchronous Remote Copy groups, RCFC for Synchronous Remote Copy groups, and FCIP for Asynchronous Streaming Remote Copy groups.





Figure 6. One-to-one bidirectional remote copy

M-to-N topology

In an M-to-N topology any HPE 3PAR StoreServ array in the topology can have a bidirectional Remote Copy relationship with up to four other HPE 3PAR StoreServ arrays. The relationship between any pair of arrays can be simultaneously synchronous and periodic asynchronous, it can be bidirectional and use RCIP, RCFC, or FCIP link layer connections. RCFC and FCIP both leverage FC ports on the StoreServ for replicating data, but with FCIP, a FC-IP router is used to extend the SAN fabric over an IP network. Asynchronous streaming is supported between any pair of arrays in an M-to-N topology. However, any given array can replicate in asynchronous streaming mode with at most one other array and the Remote Copy link layer must be RCFC or FCIP. In an M-to-N topology, the Remote Copy groups are only replicated between a single pair of arrays, there is no support for a Remote Copy group to be replicated from one primary array to two separate secondary arrays. Figure 7 and 8 show example M x N topologies.



Note: Different volume groups are being replicated (1:N)

Figure 7. Three StoreServ arrays in an M-to-N configuration (only one of these Remote Copy relationships can be in asynchronous streaming mode)




Figure 8. Five StoreServ arrays in an any-to-any M-to-N configuration (notice that all relationships can be bidirectional) with asynchronous streaming mode supported between two array pairs

Synchronous long distance topology

The Remote Copy SLD topology allows volumes in a Remote Copy volume group to be replicated from one primary StoreServ array to two different secondary StoreServ arrays. It does this by replicating data synchronously between two StoreServ arrays, the “source” and “sync target” arrays. SLD simultaneously replicates the same data via periodic asynchronous mode between the source and a third StoreServ array, the disaster recovery, or “async target” array. Right now, Asynchronous Streaming Remote Copy is not supported in SLD topologies but will be supported in the future.

Note

Check with your HPE representative for information on the support of asynchronous streaming mode in a SLD topology.

The user has the option of treating the two sync arrays in an active-active manner, failing over between them if and when a data center failure dictates a failover is necessary and resuming operations on the “sync target array”. This provides a failover solution that delivers an RPO equal to zero due to the synchronous nature of the replication that occurs between the sync arrays. On failover to a sync target array, the passive periodic asynchronous link between that array and the async target array is made active. Any data that was replicated on the sync target but that has not yet made it to the async target array is sent following the failover. This brings the async target array up to date with the last write that occurred on the sync target. Operations then continue in the sync target data center and it continues to replicate data to the async target array.

The user also has the option, when a data center failure dictates that a failover is necessary, of failing over to and resuming operations on the async target array. This can be done once data that was replicated on the sync target but that has not yet made it to the async target is replicated to the async target array. Once the async target array is consistent with the state of the sync target array, operations continue on the async target array with no data loss (RPO = 0).

This failover results in the periodic asynchronous link between the sync target and async target array being reversed so updates to the async target array are replicated back to the sync target array, albeit in periodic asynchronous mode. Used in this manner, an asynchronous long distance topology can deliver an RPO of zero at the async target site except in cases where a regional disaster has rendered both the source and sync target arrays down simultaneously.

Starting with the release of HPE 3PAR OS 3.1.2 bidirectional replication between the two “sync arrays” is supported. This means that Remote Copy can support multiple SLD configurations across the three arrays that are set up in an SLD topology (figure 9). Also, we can have other separate Remote Copy groups that are not part of an SLD configuration replicating synchronously between the two sync arrays in the topology.




Figure 9. SLD mode: Long-distance replication with zero data loss

Remote Copy performance design

HPE 3PAR StoreServ Storage is built around a massively parallel architecture where all workloads are distributed across all available system resources. Maximum performance levels are delivered by striping volumes widely across many disk spindles and dividing the workload for any given volume across all available controller nodes. This frees the administrator from the requirement to plan, design, or fine-tune volume layouts since all provisioning tasks automatically take advantage of this massively parallel architecture.

HPE 3PAR Remote Copy software is implemented over multiple replication links (at least two) for availability and performance (load balancing). All Remote Copy operations, even if for a single volume, are balanced over all available links of the same type (RCIP, RCFC, or FCIP) configured for the same replication mode (synchronous, asynchronous streaming, or periodic asynchronous) between two arrays (same RC target). By utilizing all available links, Remote Copy enhances the efficiency, performance, and scalability of the solution. Should one link fail, all operations continue seamlessly on the remaining links and alerts are generated to enable the administrator to address the link failure.

In order to further maximize the performance of Remote Copy operations, multiple connections are opened on the Remote Copy link. While ensuring write order consistency, Remote Copy will use multiple connections (sockets) to ensure that independent writes (for example, from different volume groups or while synchronizing a periodic asynchronous mode volume group) will not block each other. Furthermore, performance over high-latency TCP links is enhanced by using multiple connections over the same link.

HPE is sensitive to the needs of IT administrators and has architected Remote Copy to efficiently use the resources allotted to the solution. The “setrcopytarget tunelinks” command can be used to specify the link speed and steady state latency to RCIP allowing Remote Copy to automatically choose the correct number of connections to properly tune link throughput. RCIP Copy IP will adjust the number of sockets it uses on the network to maximize throughput based on the inputs to this command. The “setrcopytarget tput” command is used to specify the maximum amount of bandwidth RCIP is allowed to use in its IP connection for RCIP-based solutions. This gives the storage administrator some level of QoS to prevent Remote Copy from consuming the entire throughput of the available IP link.

Solution benefits

Simplified implementation

In order to simplify implementation HPE 3PAR Remote Copy software was designed to be end-user configurable and maintainable.

- Supports synchronous, asynchronous streaming, and periodic asynchronous modes so that array administrators don't have to implement different vendor solutions depending on distance requirements, data loss tolerance levels, and RTO and RPO.
- Supports replication over Fibre Channel (RCFC), IP-extended SAN fabrics (FCIP), as well as native IP (RCIP) connectivity. This again ensures that the array administrator has flexibility and choice within a single product and isn't required to introduce additional third-party hardware for distance extension or protocol conversion.
- Supported on all HPE 3PAR StoreServ models, so organizations can save drastically on hardware costs by using a mixture of midrange and high-end arrays depending on specific needs. HPE 3PAR StoreServ Storage is the only platform that currently supports multi-site and multi-mode disaster recovery using midrange arrays.¹
- End-user configurable, allowing the administrator to easily configure and manage the remote replication solution directly. This simplicity allows the administrator to respond to changing business needs quickly, without requiring additional third-party consulting or lengthy planning cycles.

Simplified management

When a storage administrator manages a remote replication solution, simplicity is essential to ensuring the safe recovery of data. HPE 3PAR Remote Copy software delivers simplicity throughout the product. For example, Remote Copy is managed with a hierarchical layering of manageable objects and relationships to help organize the replication process logically. This allows simple management of entire arrays or individual groups of volumes where appropriate. In addition, Remote Copy provides the ability to start and stop all Remote Copy groups with a single operation, which greatly enhances the ease of use. These objects and relationships are as follows:

- **Remote Copy links:** This is the vehicle by which data is sent and received between Remote Copy targets. A link is specifically defined from one physical interface (FC or IP) on one HPE 3PAR StoreServ to the equivalent physical interface another system.
- **Remote Copy pairs (or targets):** This is the relationship between two HPE 3PAR StoreServ systems. For any given system, the remote or secondary system is referred to as the remote copy target. It is the relationship between the local system or primary array and the target system or secondary array that constitutes a Remote Copy pair. Remote Copy commands generally refer to the target, which is a named object that is created when you first configure Remote Copy.
- **Remote Copy volume groups (or consistency groups):** These are logically related groups of volumes for which there is a cross-volume ordering of writes. Actions taken on a volume group are taken on all volumes in the group simultaneously, ensuring write order consistency. Multiple volume groups can be associated with each Remote Copy target.

¹ Please check with your HPE representative for details on support for Remote Copy between the versions of the HPE 3PAR StoreServ OS in your environment.

Figure 10 represents a schematic view of the layering of volume groups, targets, and links. The following statements apply to the relationships illustrated in this diagram:

- Remote Copy links are defined between systems over a physical link connection (either FC or IP). These links can be viewed and managed easily, giving the HPE 3PAR administrator a simple view of the “health” of the remote replication link.
- A Remote Copy target is the relationship between the primary system and the secondary system, and is layered over the pair of links. In this fashion, the links form a highly available transport for the Remote Copy target relationship.
- One or more volume groups are added to a Remote Copy target.




Figure 10. Remote Copy links, targets, and volume groups

Failover management

When a disaster occurs that requires failover to the remote site, it is important that the failover be easy to manage. A failover scenario could be at the level of the entire physical site or it might be limited to a single server or even a single application. Remote Copy allows for failure of entire Remote Copy targets, which results in failover of all volumes that are being replicated between the HPE 3PAR StoreServ arrays. Similarly, failover of individual Remote Copy volume groups for cases where a single server or application needs to be failed over is also supported.

Remote Copy makes failover management simple through the ability to failover one or all of the groups with a single command.

Transparent failover with HPE 3PAR Peer Persistence

Peer Persistence is a high-availability configuration between two data centers in which the ESXi hosts or Windows® hosts are setup in a metro cluster configuration with access to storage arrays in both sites. Storage volumes created on one storage array are replicated to the other array using Synchronous Remote Copy to ensure that the volumes are in sync at all times.

Peer Persistence software takes advantage of the asymmetric logical unit access (ALUA) capability that allows paths to a SCSI device to be marked as having different characteristics. With ALUA the same LUN can be exported from both arrays simultaneously, but only the paths to the side accepting write I/Os to the volume will be marked as active. The paths to the secondary side volume will be marked as standby preventing any host from performing I/O using those paths. In the event of a non-disruptive array volume migration scenario, the standby paths are marked as active and host traffic to the primary storage array is redirected to the secondary storage array without impact to the hosts. Figure 11 shows a recommended Peer Persistence configuration.




Figure 11. Transparent failover with Peer Persistence

With Peer Persistence VMware® vMotion or Hyper-V Live Migration can also be used for load balancing across data centers. With Peer Persistence, movement of VMs across data centers is completely transparent to the application, thereby enabling a truly load-balanced VMware environment.

Refer to the Peer Persistence white paper that can be located at hp.com/go/3PAR for more information.

Cross-product interoperability

HPE 3PAR Remote Copy software is common to and fully interoperable across the entire HPE 3PAR StoreServ product family, which consists of four different classes of HPE 3PAR StoreServ systems. Since all classes of systems are built on the same platform and run the same software, one can choose, for example, to replicate from up to four HPE 3PAR StoreServ 7200 systems at various data centers to a single HPE 3PAR StoreServ 10000, HPE 3PAR StoreServ 7400, or HPE StoreServ 20000 array at a failover site.

Alternatively, a customer with an HPE 3PAR StoreServ 10000 and modest remote replication needs might choose to use HPE 3PAR StoreServ 8000 at the remote site to back up a smaller subset of applications hosted on the primary array. HPE 3PAR StoreServ Storage allows the customer the flexibility to properly size the remote site storage based on the needs of the solution.

Simple verification of replicated data

Remote Copy on the target StoreServ array can be used to create read or write snapshots of the VVs in the Remote Copy volume group for verification purposes. This simple and efficient operation allows periodic validation of the data without the remote site downtime that is associated with other replication solutions. Furthermore, the same approach can be used to facilitate offsite and serverless backup. A snapshot of the Remote Copy volume group is presented to a backup server where tape backup or archives are maintained. This achieves the goal of offsite tape storage without the tapes ever being created at the original primary site.

Additionally, coordinated snapshots can be leveraged to provide a transactionally consistent copy of a database for a Remote Copy group being replicated via asynchronous streaming or periodic asynchronous by quiescing the database and then creating the coordinated snapshot. The snapshot can then be used for backup or other activities with the knowledge that it was created at a particular point in time, complete up to the last committed transaction.

Access to production data without impacting replication

Remote Copy users have the flexibility to replicate a read or write virtual copy (snapshot) of a replicated storage volume to the same or another destination of HPE 3PAR StoreServ system. Customers with separate development or data mining data centers can leverage the enhanced capability to provide these data centers with access to production data without interrupting the replication process on the base storage volume.

Recovery managers

HPE 3PAR StoreServ Storage also extends the value of periodic asynchronous mode by integrating Remote Copy with HPE 3PAR StoreServ Recovery Manager Software for Oracle, Exchange, SQL, and VMware. This integration allows users the capability of further protecting their Oracle, Exchange, SQL, or VMware environments with remote replication using HPE 3PAR Remote Copy software, all via an easy-to-manage interface.

Integration with cluster products

HPE 3PAR StoreServ Remote Copy replication is also tightly integrated with the HPE Cluster Extension (CLX), Metrocluster, and Continentalcluster products.

Integration with Cluster Extension provides true integrated disaster-tolerant solutions that extend Microsoft® Cluster Server clusters for Windows 2003, 2008, and Windows 2012 (November 2012) beyond a single data center to metro distances using Remote Copy to provide the data replication.

Integration with HPE Metrocluster and Continentalcluster products provides integrated disaster-tolerant solutions cluster for both HPE-UX and Linux® at both metro and continental distances.

All of these integrated cluster solutions provide true, complete disaster-tolerant solutions based on HPE 3PAR StoreServ Remote Copy.

Contact your HPE sales representative for more details.

Summary

HPE 3PAR Remote Copy software provides the data sharing and data protection solution that today's IT departments require with ground-breaking simplicity and efficiency. By simplifying the deployment of remote replication, Remote Copy empowers the storage administrator to deploy new disaster recovery solutions quickly and to adapt to changing business needs.

Remote Copy integration with, Cluster Extension, Metrocluster, and Continental cluster, provide disaster tolerant solutions that offer administrators solutions that effectively protect their data and quickly recover from outages with limited administrative interaction.

In addition, HPE 3PAR StoreServ Storage is the first storage platform to support multi-site capability on midrange arrays, enabling cloud service providers and enterprise customers alike to reduce their equipment costs.

Learn more at

hp.com/go/3PAR



Sign up for updates

★ Rate this document



© Copyright 2011–2013, 2015 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for HPE 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. HPE shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle is a registered trademark of Oracle and/or its affiliates. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

4AA3-8318ENW, November 2015, Rev. 5

Overview

HPE 3PAR StoreServ 8000 Storage


The HPE 3PAR StoreServ 8000 Storage offers enterprise Tier 1 storage at a midrange price. HPE 3PAR StoreServ 8000 Storage delivers the performance advantages of a purpose-built, flash-optimized architecture without compromising resiliency, efficiency, or data mobility. The new HPE 3PAR Gen5 Thin Express ASIC provides silicon-based hardware acceleration of thin technologies, including inline deduplication, to reduce acquisition and operational costs by up to 75% without compromising performance. With unmatched versatility, performance, and density, HPE 3PAR StoreServ 8000 Storage gives you a range of options that support true convergence of block and file protocols, all-flash array performance, and the use of spinning media to further optimize costs. HPE 3PAR StoreServ 8000 Storage offers rich, Tier-1 data services, quad-node resiliency, seamless data mobility between systems, high availability through a complete set of persistent technologies, and simple and efficient data protection with a flat backup to HPE StoreOnce Backup appliances. Four models are available: 8200, 8400, 8440, and 8450. You can start small and grow without painful upgrades down the road.

NOTE: For more information about the value of HPE 3PAR StoreServ 8000 Storage refer to the HPE 3PAR StoreServ 8000 Datasheet http://h20195.www2.hp.com/v2/GetDocument.aspx?docname=4AA5-9493ENW&doctype=data-sheet&doclang=EN_US&searchquery=&cc=us&lc=en

HPE 3PAR StoreServ 8000 is storage made effortless.



**HPE 3PAR StoreServ 8000 Storage
(2-Node Storage Base)**



**HPE 3PAR StoreServ 8000 Storage
(4-Node Storage Base)**

What's New

- Direct Attach Copper (DAC) Cable Support

Overview

Host OS Support

Citrix® XenServer® | HP-UX® | IBM® AIX® | Microsoft® Windows® Server, including Microsoft® Hyper-V™ | Apple Mac OS

OpenVMS* | Oracle® Linux® (UEK and RHEL compatible kernels) | Oracle® Solaris | Ubuntu | VMware vSphere™

Red Hat® Enterprise Linux® | Red Hat® Enterprise Virtualization

SUSE® Linux Enterprise | SUSE® Linux Virtualization | IBM Virtualization | Oracle VM

For the latest information on supported operating systems refer to Single Point of Connectivity Knowledge for HPE Storage Products (SPOCK): <http://www.hpe.com/storage/spock>

Overview

Summary	8200	8400	8440	8450
Number of Controller Nodes	2	2 or 4	2 or 4	2 or 4
HPE 3PAR Gen5 ASICs	2	2 or 4	2 or 4	2 or 4
Processors	2 x 6-core 2.2 GHz	2-4 x 6-core 2.2 GHz	2-4 x 10-core 2.4 GHz	2-4 x 10-core 2.4 GHz
Total Cache	832 GiB	1664 GiB	8384 GiB	384 GiB
Flash Cache (optional)	768 GiB	1536 GiB	8000 GiB	Not Applicable
On-Node Cache	64 GiB	128 GiB	384 GiB	384 GiB
Total Cache per node pair	832 GiB	832 GiB	4192 GiB	192 GiB
Flash Cache per node pair	768 GiB	768 GiB	4000 GiB	Not Applicable
On-Node Cache per node pair	64 GiB	64 GiB	192 GiB	192 GiB
Maximum Host Ports	12 ports	24 ports	24 ports	24 ports
16Gb/s Fibre Channel Host Ports	4 - 12 ports	4 - 24 ports	4 - 24 ports	4 - 24 ports
10Gb/s iSCSI Host Ports	0 - 4 ports	0 - 8 ports	0 - 8 ports	0 - 8 ports
10Gb/s FCoE Host Ports	0 - 4 ports	0 - 8 ports	0 - 8 ports	0 - 8 ports
1Gb/s Ethernet Adapter	0 - 8 ports	0 - 16 ports	0 - 16 ports	0 - 16 ports
10Gb/s Ethernet Adapter	0 - 4 ports	0 - 8 ports	0 - 8 ports	0 - 8 ports
Maximum Initiators Supported	2048	4096	4096	4096
Built-in 1GbE Ports ¹	2	2 - 4	2 - 4	2 - 4
2U Controller Node Drive Capacity	24	24	24	24
Number of Hard Disk Drives ²	6 - 240	6 - 576	6 - 960	Not Applicable
Number of Solid State Drives	6 - 120	6 - 240	6 - 480	6 - 480
Max Raw Capacity (approx.) ³	750 TiB ⁴	2400 TiB	3000 TiB	1843 TiB
Usable File Capacity ⁵	2 - 128TiB	2 - 256TiB	2 - 256TiB	2 - 256TiB

Capacity Details	8200	8400	8440	8450
RAID Levels	RAID 0, 1, 5, 6			
RAID 5 Data to Parity Ratios	2:1 - 8:1			
RAID 6 Data to Parity Ratios	4:2, 6:2, 8:2, 10:2, 14:2			
Drive Capacities (SSDs) ⁶	400 SSD, 480 SSD, 920 SSD, 480 nonAFC SSD, 1920 SSD, 3840 SSD, 7680 SSD			
Drive Capacities (HDD)	300 15K SAS ⁷ , 600 15K SAS 600 10K SAS, 1200 10K SAS, 1800 10K SAS 2000 7.2K NL ⁸ , 4000 7.2K NL, 6000 7.2K NL, 8000 7.2K NL			Not Applicable
Number of Add-on Drive Enclosures ⁹	0 - 9 enclosures	0 - 22 enclosures	0 - 38 enclosures	0 - 18 enclosures

NOTE: Specifications are subject to change without notice.

¹ Two built-in 1GbE ports per node pair can be used either for Remote Copy (RCIP) or for File Persona.

² Not applicable to HPE 3PAR StoreServ 8200 All-Flash Starter Kit and HPE 3PAR StoreServ 8400 All-Flash Starter Kit

³ Maximum raw capacity currently supported with any and all drive types. The minimum supported raw capacity is equal to 8 * Min drive size available.

⁴ For storage capacity, 1 GiB = 2³⁰ bytes and 1 TiB = 1,024 GiB

Overview

⁵ Usable file capacity supported for HPE 3PAR File Persona Software Suite

⁶ SSDs are Solid State Drives

⁷ SAS drives are Serial Access SCSI Drives

⁸ NL drives are Nearline (7200 RPM) Enterprise SAS drives

⁹ Each Drive Enclosure holds up to 24 drives in 2U for small form factor (2.5") drives and 4U for large form factor (3.5") drives

HP 3PAR StoreServ 8000 All-Flash Starter Kit

HPE 3PAR StoreServ 8200 All-Flash Starter Kit

The HPE 3PAR StoreServ 8200 All-Flash Starter Kit is an all-flash version of the HPE 3PAR StoreServ 8200 that provides all-flash acceleration at entry-level price. The kit includes the HPE 3PAR StoreServ 8200 Storage System Base equipped with 8 x 480GB SFF nonAFC SSD drives, the OS Suite, and Virtual Copy software. It also includes 3 year 24x7 Proactive Care Support. The HPE 3PAR StoreServ 8200 All-Flash Starter Kit has the same Drive LTUs cap (48) and maximum number of SSDs (120) as the HPE 3PAR StoreServ 8200.

HPE 3PAR StoreServ 8200 All Flash Starter Kit

MOS95A

The HPE 3PAR StoreServ 8200c All-Flash Starter Kit is an orderable part number that includes the following individual parts.

Quantity	Component	
1	HP 3PAR StoreServ 8200 2-node Field Integrated Storage Base	K2Q36A
8	HP 3PAR 8000 480GB SAS nonAFC SFF SSD	K2P88A
1	HP 3PAR 8200 Operating System Suite Base LTU	L7B45A
8	HP 3PAR 8200 Operating System Suite Drive LTU	L7B46A
1	HP 3PAR 8200 Virtual Copy Base LTU	L7B57A
8	HP 3PAR 8200 Virtual Copy Drive LTU	L7B58A
1	OS Media Kit	BD362AAE / BD363AAE

HPE 3PAR StoreServ 8400 All-Flash Starter Kit

The HPE 3PAR StoreServ 8400 All-Flash Starter Kit is an all-flash version of the HPE 3PAR StoreServ 8400 that provides all-flash acceleration at entry-level price in a 4-node scalable system. The kit includes the HPE 3PAR StoreServ 8400 2-Node Storage System Base equipped with 8 x 480GB SFF nonAFC SSD drives, the OS Suite, and Virtual Copy software. It also includes 3 year 24x7 Proactive Care Support. The HPE 3PAR StoreServ 8400 All-Flash Starter Kit has the same Drive LTUs cap (168) and maximum number of SSDs (240) as the HPE 3PAR StoreServ 8400.

HPE 3PAR StoreServ 8400 All-Flash Starter Kit

MOT18A

The HPE 3PAR StoreServ 8400 All-Flash Starter Kit is an orderable part number that includes the following individual parts.

Quantity	Component	
1	HP 3PAR StoreServ 8400 2-node Field Integrated Storage Base	H6Y96A
8	HP 3PAR 8000 480GB SAS nonAFC SFF SSD	K2P88A
1	HP 3PAR 8400 Operating System Suite Base LTU	L7B69A
8	HP 3PAR 8400 Operating System Suite Drive LTU	L7B70A
1	HP 3PAR 8400 Virtual Copy Base LTU	L7B81A
8	HP 3PAR 8400 Virtual Copy Drive LTU	L7B82A
1	OS Media Kit	BD362AAE / BD363AAE

HP 3PAR StoreServ 8000 All-Flash Starter Kit

NOTE: It is possible to add more SSDs, drive enclosures (and any other 8000 accessory), and software to the HPE 3PAR StoreServ 8200/8400 All-Flash Starter Kit quote. Each of the nested components will appear as separate parts on the quote and will be priced accordingly

HP 3PAR StoreServ 8000 Converged File/Block Starter Kit

HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit

The HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit is a converged file/block version of the HPE 3PAR StoreServ 8200 that is twice as easy to manage, has 40% better density, and is more attractively priced than other comparable unified storage arrays. The kit includes the HPE 3PAR StoreServ 8200 Storage System Base equipped with 8 x 600GB 10K SFF HDDs and 12 x 2TB 7.2K LFF HDDs, a LFF SAS Drive Enclosure, 2 x 2-port 10GbE adapters, the OS Suite, Replication Suite, and 16 x File Persona 1TB LTUs. It also includes 3 year 24x7 Proactive Care Support. The HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit has the same Drive LTUs cap (48) and maximum number of HDDs/SSDs (240/120) as the HPE 3PAR StoreServ 8200.

Model

HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit

MOT74A

The HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit is an orderable part number that includes the following individual parts.

Quantity	Component	
1	HP 3PAR StoreServ 8200 2N Fld Int Base	K2Q36A
1	HP 3PAR 8200 OS Suite Base LTU	L7B45A
20	HP 3PAR 8200 OS Suite Drive LTU	L7B46A
1	HP 3PAR 8200 Replication Suite Base LTU	L7B49A
20	HP 3PAR 8200 Replication Ste Drive LTU	L7B50A
16	HP 3PAR 8200 File Persona Ste 1TB LTU	BD440A
8	HP 3PAR 8000 600GB SAS 10K SFF HDD	K2P99A
12	HP 3PAR 8000 2TB SAS 7.2K LFF HDD	K2P95A
1	HP 3PAR 8000 LFF(3.5in) SAS Drive Encl	H6Z27A
2	HP 3PAR 8000 2-pt 10Gb Eth Adapter	E7Y70A
1	OS Media Kit	BD362AAE / BD363AAE

NOTE: It is possible to add more HDDs/SSDs, drive enclosures (and any other 8000 accessory), and software to the HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit quote. Each of the nested components will appear as separate parts on the quote and will be priced accordingly

Service and Support and Warranty Information

Warranty

3 Year, On-site Warranty Service for hardware components. 7x24 4-hour remote response with next business day on-site response.

The warranty on all HPE 3PAR StoreServ 8000 Solid State Drives is 5 years, parts only. Please refer to the HPE 3PAR StoreServ 8000 Drives section for the complete list of SSD SKUs. The warranty on all other HPE 3PAR StoreServ 8000 drives (SAS performance and Nearline SAS) is 3 years, parts only. Hewlett Packard Enterprise warrants only that the Software media will be free of physical defects for a period of ninety (90) days from delivery.

For more information about Hewlett Packard Enterprise's Global Limited Warranty and Technical Support, visit: <http://www.hpe.com/storage/warranty>

NOTE: All currently available HPE 3PAR StoreServ SSDs carry a five-year warranty offering unconditional replacement in case of drive failure, media wear-out, or both.

Service and Support

Protect your business beyond warranty with HPE Support Services

HPE Technology Services delivers confidence, reduces risk and helps customers realize agility and stability. Our integrated portfolio of Services for storage help customers reduce costs, optimize data, streamline storage management, and improve backup and recovery. HPE Support Services enable you to choose the right service level, length of coverage and response time as you purchase your new storage solution, giving you full entitlement for the support for need for your IT and business.

Connect your devices

Unlock all of the benefits of your technology investment by connecting your products to Hewlett Packard Enterprise. Achieve up to 77%¹ reduction in down time, near 100%² diagnostic accuracy and a single consolidated view of your environment. By connecting, you will receive 24x7 monitoring, pre-failure alerts, automatic call logging, and automatic parts dispatch. HPE Proactive Care Service and HPE Datacenter Care Service customers will also benefit from proactive activities to help prevent issues and increase optimization. All of these benefits are already available to you with your server storage and networking products, securely connected to Hewlett Packard Enterprise support.

1- IDC

2 – HP CSC reports 2014 - 2015

Optimized Care

HPE Proactive Care Advanced* - 24x7 coverage, three year Support Service

This services helps achieve a higher return on your product investment with personalized support from a local assigned Account Support Manager who will share best practice advice and personalized recommendations designed to help improve availability and performance to increase stability and reduce unplanned downtime. Leverage your system's ability to connect to Hewlett Packard Enterprise for pre-failure alerts, automatic call logging and parts dispatch. For business critical incidents, this service offers critical event management to reduce mean time to resolution. This recommendation provides 24x7 coverage with four-hour response for hardware and collaborative support that offers two-hour callback for supported software issues. Collaborative software management is included with independent software vendors unless you have your software support from Hewlett Packard Enterprise where we own all cases from start through to resolution.

<http://www.hpe.com/h20195/V2/GetDocument.aspx?docname=4AA5-3259ENW&cc=us&lc=en>

Standard Care

HPE Proactive Care* with 24x7 coverage, three year Support Service + 20 services credits for 1st year

Service and Support and Warranty Information

HPE Proactive Care gives customers an enhanced call experience plus helps preventing problems and maintains IT stability by utilizing personalized proactive reports with recommendations and advice when your products are connected to Hewlett Packard Enterprise. This Service combines three years' proactive reporting and advice with our 24x7 coverage, four hour hardware response time when there is a problem. <https://www.hpe.com/h20195/v2/GetPDF.aspx/4AA3-8855ENW.pdf>

Basic Care

HPE Proactive Care* with 24x7 coverage, three year Support Service

HPE Proactive Care gives customers an enhanced call experience plus helps preventing problems and maintains IT stability by utilizing personalized proactive reports with recommendations and advice when your products are connected to Hewlett Packard Enterprise. This Service combines three years' proactive reporting and advice with our 24x7 coverage, four hour hardware response time when there is a problem. <https://www.hpe.com/h20195/v2/GetPDF.aspx/4AA3-8855ENW.pdf>

HPE Services Support Credits

Services Support Credits offer flexible services and technical skills to meet your changing IT demands. With a menu of service that is tailored to suit your needs, you get additional resources and specialist skills to help you maintain peak performance of your IT. Offered as annual credits, you can plan your budgets while proactively responding to your dynamic business.

NOTE: *HPE Proactive Care and HPE Proactive Care Advanced require that the customer connect their devices to make the most of these services and receive all the deliverables.

Related Services

HPE 3PAR StoreServ Storage Installation and Startup Service - Hewlett Packard Enterprise installs and tests your hardware and software onsite, including configuration. We deliver a custom tailored storage deployment, properly integrated into your environment. <http://h20195.www2.hpe.com/v2/GetPDF.aspx/4AA5-8035ENW.pdf>

HPE Storage Transformation Workshop - Explore data management transformation journey to business-aligned visions, aligning your specific situation and Hewlett Packard Enterprise's experiences. <http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-9541ENW.pdf>

HPE Storage Modernization Service - Modernize your storage to take better advantage of physical or virtualized server environments. <http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA3-4620ENW.pdf>

HPE StoreServ Integration Service - Integrate your new HPE 3PAR StoreServ system so that it is agile, performs effectively, and scales to rapid growth. <http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-9254ENW.pdf>

HPE 3PAR StoreServ Data Migration - Proven expertise and tools help you migrate data across your data center or around the globe.

HPE 3PAR StoreServ Data Migration Service data sheet - US English (A4)

HPE Storage Transformation Workshop - Explore data management transformation journey to business-aligned visions, aligning your specific situation and Hewlett Packard Enterprise's experiences. <http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-9541ENW.pdf>

HPE Storage Modernization Service - Modernize your storage to take better advantage of physical or virtualized server environments.

Service and Support and Warranty Information

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA3-4620ENW.pdf>

HPE StoreServ Integration Service - Integrate your new HPE 3PAR StoreServ system so that it is agile, performs effectively, and scales to rapid growth.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-9254ENW.pdf>

HPE Storage Data Migration Service - Hewlett Packard Enterprise expertise and tools help you migrate data across your data center or around the globe. Take the burden of migration off your shoulders and put it in the capable hands of expert HPE storage migration consultants. Our proven approach helps orchestrate the complete data migration and consolidation process while maintaining consistent data availability during the transfer process.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA3-0774ENW.pdf>

HPE QuickStart Service for HPE 3PAR StoreServ Storage - Choose the most effective, appropriate methods for configuring and migrating to a HPE 3PAR

platform. <http://www8.hpe.com/us/en/services/services-detail.html?compURI=tcm:245-826727&pageTitle=Consulting-Services&contentView=business>

HPE EVA to HPE 3PAR Acceleration Service - The HPE EVA to HPE 3PAR Acceleration Service can help guide you or even execute data migration activities on your behalf that can not only optimize, but provide OPEX and CAPEX savings as a result of your journey from HPE EVA to HPE 3PAR StoreServ. This service provides customers with an alternative DIY ("do-it-yourself") data migration option with guidance from TS Storage migration specialists. With the help of migration experts, this service allows customers to execute an HPE EVA to HPE 3PAR StoreServ data migration at their own pace and lowered cost.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-4234ENW.pdf>

HPE 3PAR StoreServ 8000 Software Installation and Startup Service - Designed to provide a smooth startup, HPE 3PAR 8000 Software Installation and Startup Service provides deployment of your HPE 3PAR 8000 storage software, helping to ensure proper installation in your storage environment as well as helping you increase the benefit from your storage investment. Complementing your new HPE 3PAR 8000 storage software, HPE 3PAR 8000 Software Installation and Startup Service provides the necessary activities required to help you deploy your licensed HPE 3PAR 8000 software products into operation.

<http://h20195.www2.hpe.com/v2/GetPDF.aspx/4AA5-8036ENW.pdf>

HPE SAN Deployment Service - Hewlett Packard Enterprise delivers complete design and implementation services for Fibre Channel, FCoE, FCIP, SAS, and iSCSI storage area network (SAN) connectivity components. This service includes three levels of support based on the type of environment, ranging from simple to more complex. A trained service specialist guides the implementation of the storage switches in the SAN environment according to Hewlett Packard Enterprise quality standards.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/5981-8527EN.pdf>

HPE 3PAR Replication Software Suite Installation and Startup Service - Provides implementation of the HPE 3PAR Virtual Copy, Remote Copy, Peer Persistence and Cluster Extension components of the HPE 3PAR Replication Software Suite product. The service is designed to help get HPE 3PAR 8000 Replication Software Suite up and running quickly and to provide a demonstration of the product's key features using sample or test data only. HPE 3PAR 8000 Replication Suite includes HPE 3PAR Remote Copy, HPE 3PAR Virtual Copy, HPE 3PAR Peer Persistence, and HPE 3PAR Cluster

Service and Support and Warranty Information

Extension. The HPE 3PAR Cluster Extension Implementation service is available separately.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-2570ENW.pdf>

HPE Data Replication Solution Service for HPE 3PAR Virtual Copy - HPE Data Replication Solution Service for HPE 3PAR Virtual Copy Software helps create, manage and configure local replication data mirroring and snapshot capabilities of HPE 3PAR StoreServ storage systems. This service enables snapshots and mirroring to facilitate data restores, minimize downtime for backups, perform application testing, support data mining use with decision-support tools.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA3-8107ENW.pdf>

HPE Data Replication Solution Service for HPE 3PAR Remote Copy - HPE Data Replication Solution Service for HPE 3PAR Remote Copy Software configures real-time data mirroring between local and remote HPE 3PAR StoreServ storage systems to safeguard critical business information. Provides scalable deployment of HPE 3PAR Remote Copy Software with real-time data mirroring between a local and a remote HPE 3PAR StoreServ storage system.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA3-8627ENW.pdf>

HPE 3PAR Adaptive Optimization Policy Implementation Service - HPE 3PAR Adaptive Optimization Policy Implementation Service provides analysis, recommendations, and implementation of HPE 3PAR Adaptive Optimization policies to enable storage tiering using data collected from the HPE 3PAR Storage system over time. With the assistance of a designated IT storage administrator, a service specialist works with the customer to implement Adaptive Optimization policies on the HPE 3PAR Storage system to help deliver service level optimization for virtual and cloud data centers while reducing cost, increasing agility and minimizing risk.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA2-3842ENW.pdf>

HPE Storage Virtual Volume Design and Implementation Service - When redeploying an HPE StorageWorks Disk Array, the HPE Virtual Volume Design and Implementation Service provides the necessary activities required to design and implement a new virtual volume configuration.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA2-3764ENN.pdf>

HPE Thin Volume Conversion Service - Provides evaluation and execution of conversion from standard to thin provisioned virtual volumes for HPE 3PAR Storage. A service specialist advises the customer on HPE 3PAR Thin Provisioning best practices, provides evaluation of potential disk capacity savings if target virtual volumes are converted, and plans and implements thin conversion processing. The service leverages 3PAR thin provisioning capabilities to help optimize storage capacity, reduce cost, increase agility and maintain performance.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA2-3842ENW.pdf>

HPE Performance Analysis Service for HPE Disk Arrays - The service provides data collection, detailed I/O analysis and enhancement recommendations for HPE 3PAR StoreServ Storage disk arrays, HPE EVA P6000 Storage disk arrays and HPE XP P9000 Storage disk arrays. HPE Performance Analysis Service for HPE Storage Disk Arrays provides a single engagement concerning the performance of a single HPE 3PAR StoreServ Storage disk array, HPE EVA P6000 Storage disk array and HPE XP P9000 Storage disk array

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/5982-6668EN.pdf>

Service and Support and Warranty Information

HPE 3PAR Performance and Capacity Trending Service - HPE 3PAR Performance and Capacity Trending Service provides data collection, analysis, and reports with key performance and capacity metrics for your HPE 3PAR StoreServ array. Through this service, you will receive a specified number of reports describing long-term trends in performance and capacity usage, and have the option to purchase additional reports. You will also receive briefing sessions highlighting The Hewlett Packard Enterprise findings and recommendations.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA5-8792ENW.pdf>

HPE 3PAR Health Check Service - The HPE 3PAR Health Check service is delivered as a single engagement, providing data collection, analysis, report creation, and a briefing session concerning the performance of a single HPE 3PAR StoreServ Storage System. This health check service is best for HPE 3PAR StoreServ Storage Systems that have been installed and are in normal production mode. It can also be used to establish a baseline for future reference to improve the effective use of your storage system.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-3225ENW.pdf>

HPE 3PAR Storage Rebalance Service - The HPE 3PAR Rebalance Service helps balance data across an HPE 3PAR StoreServ Storage array to take advantage of the capabilities of the array architecture. The service provides analysis, planning, and implementation of data movement and/or physical movement of drive magazines within the array.

<http://h20195.www2.hpe.com/V2/GetPDF.aspx/4AA4-0280ENW.pdf>

Customer Self Installation

Customer Self Installation (CSI)

Customers have the option of self-installing HPE 3PAR StoreServ 8000 Storage system. The Customer Self Installation option is available for HPE 3PAR StoreServ 8000 Storage systems that meets the following criteria:

- 2-node configuration (8200, 8400 2N, 8440 2N, 8450 2N)
- Maximum of 4 additional drive enclosures
- Single rack (the physical Service Processor can be in a separate rack)
- CTO configurations (factory integrated)
- BTO configurations (field integrated) without additional host adapters

Customer technical profile

In order to successfully install the HPE 3PAR StoreServ 8000 Storage system the installer should:

- Have a good understanding and knowledge of Storage Area Networks, Fiber Channel fundamentals and a basic understanding of TCP/IP and other networking protocols (DNS/NTP).
- Have a good understanding of server virtualization technology, in particular of Hypervisors such as VMware ESXi and Microsoft Hyper-V.
- Be able to maintain and install server hardware and Microsoft Windows and/or Linux Operating Systems.
- Have experience creating Storage LUNs, presenting/exporting LUNs to a server and formatting the LUNs to make them usable for applications.
- Be able to troubleshoot hardware and software issues using logs and documentation.

If the installer doesn't meet the profile or is not comfortable with the self-installation process, Hewlett Packard Enterprise recommends engaging the Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Channel Partner to purchase HPE deployment services.

Customer responsibilities

The Customer will:

- Ensure that the host and SAN environment is supported and compliant with HPE recommendations and best practices. Host and SAN Implementation Guides are available at <http://www.hp.com/go/hpsc>. Support Matrix are available on SPOCK (HP Storage Single Point of Connectivity Knowledge) <http://www.hpe.com/storage/spock>.
- Resolve any problems with their SAN and host environment, prior to installing the HPE 3PAR StoreServ 8000 Storage.

Customer Self Installation documentation

Prior to installing the HPE 3PAR StoreServ 8000 Storage system, the installer should thoroughly review the following documentation.

- HPE 3PAR StoreServ 8000 Storage Self-Install Guide: <http://www.hpe.com/support/3PAR8000CSI>
- HPE 3PAR StoreServ 8000 Storage Series Cabling Configuration Guide: <http://www.hpe.com/support/3PAR8000Cabling>
- HPE 3PAR StoreServ 8000 Storage Installation video: <http://www.hpe.com/support/3PAR8000CSIVideo>
- Forum on HPE 3PAR StoreServ 8000 Self-Install: <http://www.hpe.com/forum/3PAR8000CSIHELP>

The Customer Self Installation option is available only for initial installs, not for upgrades. Customer Self Upgrade (CSU) is optional for HPE 3PAR OS software. Customer Self Repair (CSR) information is available at this

link: http://h18033.www1.hp.com/hpe/support/selfrepair/ww/replace_part.asp?myinc=s008

NOTE: Customers performing a self-install (according to rules identified above) will not void their warranties and will be fully supported.

Configuration Information

Step 1 - Choose a Base configuration

HPE 3PAR StoreServ 8000 configurations start with the selection of the Base. The Base includes controller nodes, bays for small form factor drives, and PCIe slots for host adapter cards. SAN Kits are also considered base configurations.

HPE 3PAR StoreServ 8000 Base Configurations



Factory Integrated in HPE rack	HP 3PAR StoreServ 8200 2-node Storage Base	K2Q35A
	HP 3PAR StoreServ 8400 2-node Storage Base	H6Y95A
	HP 3PAR StoreServ 8400 4-node Storage Base	H6Z01A
	HP 3PAR StoreServ 8440 2-node Storage Base	H6Y97A
	HP 3PAR StoreServ 8440 4-node Storage Base	H6Y98A
	HP 3PAR StoreServ 8450 2-node Storage Base	H6Z17A
	HP 3PAR StoreServ 8450 4-node Storage Base	H6Z23A
Factory Integrated in HPE rack in a Storage Centric Configuration	HP 3PAR StoreServ 8200 2-node Storage Base for Storage Centric Rack	K2Q37A
	HP 3PAR StoreServ 8400 2-node Storage Base for Storage Centric Rack	H6Z12A
	HP 3PAR StoreServ 8400 4-node Storage Base for Storage Centric Rack	H6Z03A
	HP 3PAR StoreServ 8440 2-node Storage Base for Storage Centric Rack	H6Z09A
	HP 3PAR StoreServ 8440 4-node Storage Base for Storage Centric Rack	H6Z14A
	HP 3PAR StoreServ 8450 2-node Storage Base for Storage Centric Rack	H6Z20A
	HP 3PAR StoreServ 8450 4-node Storage Base for Storage Centric Rack	H6Z25A
Field Integrated	HP 3PAR StoreServ 8200 2-node Field Integrated Storage Base	K2Q36A
	HP 3PAR StoreServ 8400 2-node Field Integrated Storage Base	H6Y96A
	HP 3PAR StoreServ 8400 4-node Field Integrated Storage Base	H6Z02A
	HP 3PAR StoreServ 8440 2-node Field Integrated Storage Base	H6Z07A
	HP 3PAR StoreServ 8440 4-node Field Integrated Storage Base	H6Z13A
	HP 3PAR StoreServ 8450 2-node Field Integrated Storage Base	H6Z18A
	HP 3PAR StoreServ 8450 4-node Field Integrated Storage Base	H6Z24A

- A minimum of one (1) configuration base must be ordered for each array.
- The HPE 3PAR StoreServ 8000 base configuration includes (2 or 4) controller nodes, (24) small form factor drive bays per node pair, (4) built-in 16 Gb/sec FC ports per node pair, (4) 16Gb shortwave FC SFP per node pair, (2) PCIe adapter slots for host adapter cards per node pair (one slot per node), (2) 2m SAS cables per node pair, (1) mounting rail kit per node pair, and power cords.
- The HPE 3PAR StoreServ 8000 2-node base configuration also includes (2) 1U rack filler panels to reserve 2U of rack space above the 2-node Storage Base for a future upgrade to a 4-node configuration
- All base configurations include (1) built-in 1GbE port for management and (1) 1GbE port for either Remote Copy over IP or File Persona, per node.

Configuration Information

- The Storage Centric rack versions of HPE 3PAR StoreServ 8000 are for CTO (factory Configure-To-Order) only. With a Storage Centric configuration, the storage system gets placed in the center of the rack so that future expansion of that storage system becomes easier.
- In Storage Centric configurations, non-3PAR components, with the exception of certain StoreFabric Storage Networking switches, if added to the same order, get placed in a separate rack.

The following diagrams show a HPE 3PAR StoreServ 8000 4N Storage System in Storage Centric and non-Storage Centric configurations



StoreServ 8000 4-node Storage system in a storage centric configuration in a HPE Intelligent Series Rack

StoreServ 8000 4-node Storage system in a non-storage centric configuration in a HPE Intelligent Series Rack

Configuration Information

HPE 3PAR StoreServ 8000 Upgrade Controller Node Pair

Use the HPE 3PAR StoreServ 8000 Upgrade Node Pair to convert an existing previously installed HPE 3PAR StoreServ 8000 2-node Storage Base into a 4-node configuration.

HP 3PAR StoreServ 8400 Upgrade Node Pair	H6Z06A
HP 3PAR StoreServ 8440 Upgrade Node Pair	H6Z08A
HP 3PAR StoreServ 8450 Upgrade Node Pair	H6Z19A

- H6Z06A is used to upgrade a HPE 3PAR StoreServ 8400 2-node Storage Base into a 4-node configuration.
- H6Z08A is used to upgrade a HPE 3PAR StoreServ 8440 2-node Storage Base into a 4-node configuration.
- H6Z19A is used to upgrade a HPE 3PAR StoreServ 8450 2-node Storage Base into a 4-node configuration.
- One (1) pair of controller nodes beyond the base configuration is supported on the 3PAR StoreServ 8000.
- The 3PAR StoreServ 8000 Upgrade Node Pair includes (2) controller nodes, (24) small form factor drive bays, (4) built-in 16 Gb/sec FC ports, (4) 16Gb shortwave FC SFP, (2) PCIe adapter slots for host adapter cards (one slot per node), (2) 2m SAS cables, (4) node link cables, (1) mounting rail kit, and power cords
- The upgrade node pair includes (2) built-in 1GbE ports for management and (2) 1GbE ports for either Remote Copy over IP or File Persona.

Step 2 - Choose Host Adapter

Host adapters can be ordered separately to be installed in the field or they can be factory configured into controller nodes. Host adapter cards provide the array with additional FC ports, with 10Gb/s iSCSI/FCoE ports, or with 1GbE/s and 10Gb/s Ethernet ports. The additional FC ports can be used for connection to hosts or used to connect to other HPE 3PAR StoreServ Storage systems in a Remote Copy relationship. The iSCSI/FCoE ports permit host connection in iSCSI and FCoE environments. The Ethernet ports can be used only with the HPE 3PAR File Persona Software Suite for File services connectivity.

HPE 3PAR StoreServ 8000 Host Adapters

HP 3PAR StoreServ 8000 4-port 16Gb Fibre Channel Adapter	H6Z00A
HP 3PAR StoreServ 8000 2-port 10Gb iSCSI/FCoE Adapter	H6Z10A
HP 3PAR StoreServ 8000 4-port 1Gb Ethernet Adapter	H6Z05A
HP 3PAR StoreServ 8000 2-port 10Gb Ethernet Adapter	E7Y70A

- The host adapter cards are optional because the Storage Base products and the Upgrade Controller Node Pair include built-in FC ports.
- Ethernet Adapters (H6Z05A and E7Y70A) can be used exclusively with the HPE 3PAR File Persona Software Suite.
- The 16Gb/s Fiber Channel Adapter (H6Z00A) includes (4) 16Gb/s shortwave FC SFP+. The 10Gb/s iSCSI/FCoE Adapter (H6Z10A) includes (2) 10Gb/s shortwave SFP+. The 10Gb/s Ethernet Adapter (E7Y70A) includes (2) 10Gb/s SR SFP+

Configuration Information

- Each node in a node pair (a node pair is composed of the two controller nodes in a single 2U enclosure) must have the same number and type of adapters: FC, iSCSI/FCoE, and Ethernet adapters may not be intermixed in a node pair.
- The 4 ports of the FC adapter can be individually configured to connect to a host or to a remote array in an RC configuration.
- The two ports of the iSCSI/FCoE adapter can be individually configured by the user as iSCSI or FCoE.

Adapter Configurations permitted on HPE 3PAR StoreServ 8000 2-node systems

Product	Nodes	A	B	C	D				
2-node Storage Base	Node 0	Empty	FC	iSCSI/FCoE	Eth				
	Node 1	Empty	FC	iSCSI/FCoE	Eth				
Upgrade Node Pair	Node 2	Empty	FC	iSCSI/FCoE	Eth				
	Node 3	Empty	FC	iSCSI/FCoE	Eth				

Adapter Configurations permitted on HPE 3PAR StoreServ 8000 4-node systems

Product	Nodes	A	B	C	D	E	F	G	H
4-node Storage Base	Node 0	Empty	FC	iSCSI/FCoE	Eth	Empty	Empty	Empty	FC
	Node 1	Empty	FC	iSCSI/FCoE	Eth	Empty	Empty	Empty	FC
	Node 2	Empty	Empty	Empty	Empty	FC	iSCSI/FCoE	Eth	iSCSI/FCoE
	Node 3	Empty	Empty	Empty	Empty	FC	iSCSI/FCoE	Eth	iSCSI/FCoE
		I	J	K	L	M	M	O	P
	Node 0	iSCSI/FCoE	FC	Eth	iSCSI/FCoE	Eth	FC	iSCSI/FCoE	Eth
	Node 1	iSCSI/FCoE	FC	Eth	iSCSI/FCoE	Eth	FC	iSCSI/FCoE	Eth
	Node 2	FC	Eth	FC	Eth	iSCSI/FCoE	FC	iSCSI/FCoE	Eth
	Node 3	FC	Eth	FC	Eth	iSCSI/FCoE	FC	iSCSI/FCoE	Eth



NOTE: If the configuration includes Remote Copy over Fibre Channel it is recommended that optional Fiber Channel HBAs are purchased as the built in ports will not offer sufficient connectivity.

NOTE: Ethernet Adapters can be used only with the HPE 3PAR File Persona Software Suite for File services connectivity

Step 3 - Choose Drive Enclosures

Add drive enclosures to expand the configuration and to add large form factor drives to the configuration. Drive enclosures can be ordered separately for installation in the field, or they can be factory configured in a rack. Drive enclosures are optional. Because the Storage Base products and the Upgrade Node Pair include small form factor drive bays, the minimum configuration does not require any additional drive enclosures. For larger configurations, attach drive enclosures. Each drive enclosure includes 24 drive bays. The two drive enclosure types can be intermixed in a single array. The HPE 3PAR StoreServ 8200 supports up to nine (9) added drive enclosures. The HPE 3PAR StoreServ 8400 supports up to twenty-two (22) added drive enclosures. The HPE 3PAR StoreServ 8440 supports up to thirty-eight (38) added drive enclosures. The HPE 3PAR StoreServ 8450 supports up to eighteen (18) added drive enclosures.

Configuration Information



HPE 3PAR StoreServ 8000 SFF(2.5in) SAS Drive Enclosure **HPE 3PAR StoreServ 8000 LFF(3.5in) SAS Drive Enclosure**

Drive Enclosures

HP 3PAR StoreServ 8000 SFF(2.5in) SAS Drive Enclosure	H6Z26A
HP 3PAR StoreServ 8000 LFF(3.5in) SAS Drive Enclosure	H6Z27A
HP 3PAR StoreServ 8000 SFF(2.5in) Field Integrated SAS Drive Enclosure	E7Y71A
HP 3PAR StoreServ 8000 LFF(3.5in) Field Integrated SAS Drive Enclosure	E7Y72A

- Each drive enclosure includes 24 drive bays, (2) IO modules, (2) 1m SAS cables, (1) mounting rail kit, and power cables.
- The 2U SAS drive enclosure provides 24 SFF drive bays arranged in a single row.
- The 4U drive enclosure provides 24 LFF drive bays, arranged in four (4) columns of six (6) slots each.
- Drive enclosures are connected in daisy chains from the SAS ports of the controller nodes.
- The best practice is to balance the drive enclosures across the SAS ports, remembering that the controller node enclosures include (24) drives attached to the SAS port labeled DP-1.
- The best practice when including LFF and SFF drive enclosures in the same array is to arrange them in the rack so that all of the 2U enclosures that belong to one node pair are together and all of the 4U enclosures for that node pair are together. When connecting the backend SAS cables, intermix the 2U and 4U SAS enclosures on each SAS port.
- With a four node configuration, the best practice is to attach the same number of drive enclosures and drive types to each node pair.
- To achieve highest availability in multi-enclosure configurations, configure a minimum of two (2) enclosures per node pair for RAID 1, a minimum of four (4) enclosures per node pair if RAID 5 is included, and a minimum of three (3) enclosures per node pair if RAID 6 is included. Include enclosures containing node pairs in the count with the 2U SAS enclosures.
- Drive bays that are not filled with a drive must be covered with a drive blank to preserve proper air flow.
- If future capacity upgrades are expected, include enough Drive Enclosures so that there are some empty bays in each enclosure after all drives are added.

Configuration Information

Step 4 - Choose Drives

Drives are orderable at the time the array is purchased, or can be added in the future when additional capacity is required. HPE 3PAR StoreServ 8000 drives are sold as single drives. Note that these drives are only compatible with the HPE 3PAR StoreServ 8000 SAS Drive Enclosures.

HPE 3PAR StoreServ 8000 SAS Drives

HPE 3PAR SSDs	HPE 3PAR StoreServ 8000 400GB SAS MLC SFF(2.5in) Solid State Drive	N9Y06A
	HP 3PAR StoreServ 8000 480GB SAS MLC SFF(2.5in) Solid State Drive	K2Q95A
	HP 3PAR StoreServ 8000 480GB SAS cMLC SFF(2.5in) Solid State Drive	K2P88A
	HP 3PAR StoreServ 8000 480GB SAS cMLC LFF(3.5in) Solid State Drive	K2Q96A
	HP 3PAR StoreServ 8000 1.92TB SAS cMLC SFF(2.5in) Solid State Drive	K2P89A
	HP 3PAR StoreServ 8000 3.84TB SAS cMLC SFF(2.5in) Solid State Drive	K2P91A
	HPE 3PAR StoreServ 8000 7.68TB SAS SFF (2.5in) Solid State Drive	P9L83A
HPE 3PAR SAS HDDs (Performance HDDs)	HP 3PAR StoreServ 8000 300GB SAS 15K SFF(2.5in) Hard Drive	K2P97A
	HP 3PAR StoreServ 8000 600GB SAS 15K SFF(2.5in) Hard Drive	K2P98A
	HP 3PAR StoreServ 8000 600GB SAS 10K SFF(2.5in) Hard Drive	K2P99A
	HP 3PAR StoreServ 8000 1.2TB SAS 10K SFF(2.5in) Hard Drive	K2P93A
	HP 3PAR StoreServ 8000 1.8TB SAS 10K SFF(2.5in) Hard Drive	K2P94A
HPE 3PAR NL SAS HDDs	HP 3PAR StoreServ 8000 2TB SAS 7.2K SFF(2.5in) Hard Drive	M0S92A
	HP 3PAR StoreServ 8000 2TB SAS 7.2K LFF(3.5in) Hard Drive	K2P95A
	HP 3PAR StoreServ 8000 4TB SAS 7.2K LFF(3.5in) Hard Drive	K2P87A
	HP 3PAR StoreServ 8000 6TB SAS 7.2K LFF(3.5in) Hard Drive	K2P96A
	HPE 3PAR StoreServ 8000 8TB SAS 7.2K LFF(3.5in) Hard Drive	P9B44A

- For each drive type installed in the array, the minimum recommended initial quantity is eight (8) drives per node pair for SSD and SAS performance HDDs, and twelve (12) drives per node pair for Nearline HDDs.
NOTE: 8 drives support RAID 1 and RAID 5. For RAID 6 choose 12 drives.
- Minimum upgrade quantity is 4 drives per node pair or 2 drives per enclosure, whichever is larger. Best practice is to run Autonomic Rebalance (also known as tunesys) after adding the drives.
- RAID 6 is strongly recommended for Nearline drives.
- All node enclosures must contain either zero (0) or an even number of the same type of drives (FC, NL, SSD).
- All drive enclosures must contain an even number of drives, with a minimum of two.
- A best practice is to add equal numbers of drives to all enclosures compatible with the drive type being added.
- With a four node configuration, the best practice is to attach the same number and type of drives to each node pair.
- Small Form Factor (SFF)-specific configuration practices

Configuration Information

- SFF drives may be loaded into the Storage Base enclosures, the Upgrade Node Pair enclosure and the 2U SAS drive enclosure.
- SFF drives must be loaded in pairs of identical drives, beginning with the leftmost slot, slot 0, and filling to the right, leaving no empty slots between drives.
- Large Form Factor (LFF)-specific configuration practices
 - LFF drives may be loaded into the 4U SAS drive enclosure.
 - LFF drives must be loaded in pairs of identical drives starting at the bottom of a column, leaving no empty slots between drives in the column.
 - Intermixing SSDs and spinning media in a LFF drive enclosure is allowed as long as each drive type is installed in even pairs in the same column.
 - It is permitted to have empty columns between columns containing drives. Different columns do not have to contain the same number of drives.
 - An all LFF drive configuration is permitted, leaving the Storage Base Enclosure empty.

HPE 3PAR FIPS Encrypted SSD/HDD	HP 3PAR StoreServ 8000 920GB SAS MLC SFF(2.5in) FIPS Encrypted Solid State Drive	K2P90A
	HP 3PAR StoreServ 8000 1.92TB SAS cMLC SFF(2.5in) FIPS Encrypted Solid State Drive	K2R27A
	HPE 3PAR StoreServ 8000 3.84TB SAS cMLC SFF(2.5in) FIPS Encrypted Solid State Drive	M0T66A
	HP 3PAR StoreServ 8000 600GB SAS 15K SFF(2.5in) FIPS Encrypted Hard Drive	K2P92A
	HP 3PAR StoreServ 8000 1.2TB SAS 10K SFF(2.5in) FIPS Encrypted Hard Drive	K2P85A
	HPE 3PAR StoreServ 8000 2TB SAS 7.2K LFF(3.5in) FIPS Encrypted Hard Drive	N9Y04A
	HPE 3PAR StoreServ 8000 4TB SAS 7.2K LFF(3.5in) FIPS Encrypted Hard Drive	N9Y05A
	HP 3PAR StoreServ 8000 6TB SAS 7.2K LFF(3.5in) FIPS Encrypted Hard Drive	K2P86A
	HPE 3PAR StoreServ 8000 8TB SAS 7.2K LFF(3.5in) FIPS Encrypted Hard Drive	P9B45A
HPE 3PAR Encryption License	HPE 3PAR 8200 Data Encryption LTU	L7B67A
	HPE 3PAR 8200 Data Encryption E-LTU	L7B67AAE
	HPE 3PAR 8400 Data Encryption LTU	L7B91A
	HPE 3PAR 8400 Data Encryption E-LTU	L7B91AAE
	HPE 3PAR 8440 Data Encryption LTU	L7C15A
	HPE 3PAR 8440 Data Encryption E-LTU	L7C15AAE
	HPE 3PAR 8450 Data Encryption LTU	L7C39A
	HPE 3PAR 8450 Data Encryption E-LTU	L7C39AAE

- An encrypted HPE 3PAR StoreServ array, i.e. any HPE 3PAR StoreServ array that has the HPE 3PAR Data Encryption license activated or intended to be activated, must have only self-encrypted drives installed.
- A non-encrypted HPE 3PAR StoreServ array can have a mix of encrypted and non-encrypted drives.

Configuration Information

- Customers have option to turn on encryption, non-disruptively, at any time, even after data has been written to the system.
- FIPS 140-2 Validated Self-Encrypting Drives (SEDs) have been certified by the U.S. National Institute of Standards and Technology (NIST) and Canadian Communications Security Establishment (CSE) as meeting the Level 2 security requirements for cryptographic modules as defined in the Federal Information Processing Standards (FIPS) 140-2 Publication
- Strengthen the DAR solution with an optional FIPS 140-2 Level-2 validated external key manager. Supports KMIP 1.1 for key management communications
- Supports HPE Enterprise Secure Key Manager 4.0 and SafeNet KeySecure k460 and k150 centralized key management
- A data encryption license (LTU) is required to enable encryption on the array. One encryption license is required for each encrypted array.
- Once encryption is enabled on the HPE 3PAR StoreServ Storage, it cannot be disabled.
- The local key manager is included in the HPE 3PAR OS. There is not a separately orderable part number for the local key manager.

Step 5 - Choose Service Processor Implementation

The HPE 3PAR Service Processor remotely monitors the HPE 3PAR StoreServ 8000 and enables remote servicing of the array. The key capabilities of the Service Processor are to:

- Enable rapid, proactive responses to issues
- Provide a secure communication channel between the customer's data center and HPE 3PAR Central for:
 - Remote Online Software Upgrade --Upgrade software with no application disruption
 - Remote Diagnostics --Maintain key diagnostic information centrally on a historical basis
 - Remote Serviceability--Provide fast predictive response and remediation

Each HPE 3PAR StoreServ 8000 requires its own Service Processor. The Service Processor functions as the communication interface between a customer's IP network and HPE 3PAR Central by managing all service-related communications. The Service Processor leverages the industry-standard HTTP over Secure Sockets Layer (HTTPS) protocol to secure and encrypt data communication. The Service Processor can be deployed either as a virtual Service Processor (VSP) or a physical Service Processor.

Virtual Service Processor

A virtual Service Processor is included free with the base HPE 3PAR Operating System. The virtual Service Processor can be installed on a customer-provided VMware or Microsoft Hyper-V system that meets the following specifications:

- Virtualization operating system
 - VMware vSphere ESXi 4.1, 5.0, 5.1, 5.5
 - Microsoft Hyper-V Server 2008 R2, 2012 or 2012 R2
- Server features
 - 2 GB RAM (minimum for the VSP Virtual Machine)
 - 256 GB free disk space (minimum for the VSP Virtual Machine)
 - NOTE: VSP storage must not reside on the array it is managing.**
 - DVD ROM or DVD RW
 - 1 Gb Ethernet port
 - For VSP on VMware, the server must be listed in the VMware Compatibility Guide

Configuration Information

- For VSP on Hyper-V, the server must be listed on the Windows Server Catalog

Physical Service Processor

The physical Service Processor is a dedicated storage appliance located within the storage rack providing close proximity to the HPE 3PAR StoreServ 8000 Storage. The physical Service Processor is fully supported and maintained by HPE Services. The physical Service Processor has serial port connectivity that provides maintenance access for trouble shooting capabilities.

If a VMware server is not available to run the virtual Service Processor, the physical Service Processor is the alternative choice for remote monitoring and remote service. The physical Service Processor is available in two version: with Single Power Supply and with Redundant Power Supply.

Service Processor	HP 3PAR StoreServ SPS Service Processor	K2R28A
	HP 3PAR StoreServ RPS Service Processor	K2R29A

HPE 3PAR Policy Server

HPE 3PAR Policy Server works to implement customer-configurable remote service access policies. Installed on a customer-provided host, Policy Server provides the customer with ultimate flexibility and control to allow or deny outbound communication or remote service connections to and from an HPE 3PAR StoreServ Storage system. Policy Server also serves as the centralized point for collecting and storing audit logs of all diagnostic data transfers and authorized remote service connections to and from all configured HPE 3PAR Storage systems. HPE 3PAR Policy Server provides the

- The customer has complete control over policy administration.
- A centralized policy administration for all HPE 3PAR Storage systems is provided.
- A centralized audit log to facilitate security audits is provided.
- Up to 100 3PAR systems can be managed with a single 3PAR Policy Server license
- Policy Server 6.15 can be run on a Virtual Machine. For the latest information on supported hypervisors, refer to Single Point of Connectivity Knowledge for HPE Storage Products (SPOCK): <http://www.hpe.com/storage/spock>

HP 3PAR Policy Manager Software LTU

TE400B

Step 6 - Choose Cables for host connection and remote copy connection

Cables are required on the HPE 3PAR StoreServ 8000 Storage for drive enclosure connections and for host connectivity. Copper SAS cables are required for connecting the drive enclosures to the nodes on the same rack and for daisy chaining between adjacent drive enclosures. Storage Base products, the Upgrade Node Pair and the Drive Enclosures all include Copper SAS cables. SAS Active Optical Cables are required if an HPE 3PAR StoreServ 8000 needs to be expanded into an adjacent rack, to connect drive enclosures in adjacent racks to the nodes in the base rack. OM4 Fiber Cables are required for host connectivity, Remote Copy and Peer Motion. The copper 1GbE cables are used for Remote Copy over IP and for connection to the Management Port.

Cables

SAS Active Optical Cables	HP 10m Mini SAS High Density Active Optical Cable	E7V95A
	HP 25m Mini SAS High Density Active Optical Cable	E7V96A

OM4 Cables

HP Premier Flex LC/LC Multi-mode OM4 2 fiber 1m Cable	QK732A
HP Premier Flex LC/LC Multi-mode OM4 2 fiber 2m Cable	QK733A
HP Premier Flex LC/LC Multi-mode OM4 2 fiber 5m Cable	QK734A
HP Premier Flex LC/LC Multi-mode OM4 2 fiber 15m Cable	QK735A
HP Premier Flex LC/LC Multi-mode OM4 2 fiber 30m Cable	QK736A

Configuration Information

	HP Premier Flex LC/LC Multi-mode OM4 2 fiber 50m Cable	QK737A
Copper 1GbE cables	HP 4.3m/14ft CAT5 RJ45 M/M Ethernet Cable	C7536A
	HP 7.6m/25ft CAT5 RJ45 M/M Ethernet Cable	C7537A
	HP 15.2m/50ft CAT5 RJ45 M/M Ethernet Cable	C7542A
Optical splitters	HP Multi Fiber Push On to 4 x Lucent Connector 5m Cable	K2Q46A
	HP Multi Fiber Push On to 4 x Lucent Connector 15m Cable	K2Q47A
Direct Attach Copper Cables* HPE 3COM (H3C)		
	HP X240 10G SFP+ to SFP+ 1.2m Direct Attach Copper Cable	JD096C
	HP X240 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	JD097C
	HP X240 10G SFP+ to SFP+ 5m Direct Attach Copper Cable	JG081C
	HP X240 QSFP+ 4x10G SFP+ 1m DAC Cable	JG329A
	HPE X240 QSFP+ 4x10G SFP+ 3m DAC Cable	JG330A
	HPE X240 QSFP+ 4x10G SFP+ 5m DAC Cable	JG331A
HPE ProCurve		
	HP X242 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	J9281B
	HP X242 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	J9283B
Cisco		
	HPE C-series 3M Passive Copper SFP+ Cable	K2Q21A
	HPE C-series 5M Passive Copper SFP+ Cable	K2Q22A

NOTE: DAC cable support for 3PAR 8000 platforms requires HPE 3PAR OS version 3.2.2 MU3 or higher. DAC cables are supported for 10GbE speeds with iSCSI, FcoE, and File protocols.

Step 7 - Choose Racking Options

The HPE 3PAR StoreServ 8000 is compatible with most industry standard 4-post EIA 19 inch racks with square mounting holes, including the HPE Intelligent Series Rack and the HPE 10000 G2 Series Rack. The HPE 3PAR StoreServ 8000 can be factory configured and shipped in a rack or shipped without a rack for field integration into an existing rack. The rack used for factory integration is the HPE Intelligent Series Rack.

HPE Rack and Rack Options

Factory Select a rack to house your HPE 3PAR StoreServ 8000.

Integration **NOTE:** The HPE Intelligent Series Rack is the only series supported for factory configuration.

Primary Configuration Rules The HPE 3PAR StoreServ 8000 will be configured into an HPE Intelligent Series Rack with the appropriate power distribution units (PDUs). If other products such as servers or back-up products are included in the cab, a different PDU will be added (if required) or can be chosen from a list of appropriate offerings shown in the configuration tool. The HPE Intelligent Series Rack must be purchased for factory configuration. Additional 3PAR StoreServ 8000 controller node enclosures and drive enclosures may be ordered for multiple subsystem integration at the factory. The 3PAR StoreServ 8000 is also supported in HPE 10000 G2 Series racks for field installation. When calculating available U-space, assume that no space will be placed between the mounted

Configuration Information

components. For redundancy, order PDUs in quantities of two. Refer to the Configuration and User Guide in the Information Library at the Rack Solutions webpage.

HPE Intelligent Series Racks

NOTE: The number of components that will fit in a rack varies and is determined by the interior U-space of the rack.

HP 47U 600mm x 1075mm Enterprise Shock Rack	BW912A
HP 47U 1075mm Side Panel Kit	BW915A
HP 42U 600mm x 1075mm Enterprise Shock Rack	BW904A
HP 42U 1075mm Side Panel Kit	BW906A
HP 36U 600mm x 1075mm Enterprise Shock Rack	BW896A
HP 36U 1075mm Side Panel Kit	BW898A
HP 42U 600mm x 1200mm Enterprise Shock Rack	BW908A
HP 42U 1200mm Side Panel Kit	BW909A

For more information on the HPE rack offerings, please see the following URL:

<http://h18004.www1.hpe.com/products/servers/platforms/rackandpower.html>

For more information on rack options, see: **<http://www.hpe.com/products/rackoptions>**.

For more information on PDUs, see:

<http://h18004.www1.hpe.com/products/servers/proliantstorage/power-protection/pdu.html>.

HPE PDU Pivot Kit

HP EVA PDU Pivot Kit

AG730A

Used to reclaim 2U of space in an HPE Intelligent Series Rack with HPE 3PAR StoreServ 8000 configurations. This kit allows the PDUs to be placed in the back of the rack without requiring any rack U space.

NOTE: The use of the PDU Pivot Kit is strongly recommended and is the default option when orders are configured, as it will save 2U of valuable rack space.

NOTE: 0D1 will appear after this part number to indicate factory integration where appropriate.

Non-HPE rack and power requirements

The Storage Bases, the Upgrade Node Pair, and the Drive Enclosures include mounting rails that are compatible with industry standard 4-post EIA 19 inch racks with square mounting holes. For detailed information on determining compatibility of a non-HPE rack, please review the information included in the HPE 3PAR StoreServ 8000 StoreServ 8000 Site Planning Guide

Step 8 - Choose Software

Hewlett Packard Enterprise provides an extensive selection of features for HPE 3PAR StoreServ Storage. All of the features available on the HPE 3PAR StoreServ 20000 Storage system are also available on the HPE 3PAR StoreServ 8000 Storage system, the result of a common architecture that spans from small and medium businesses to the largest global enterprise. For convenient ordering, the 3PAR StoreServ 8000 provides the features in Suites.

Two types of software licensing methods are employed with the HPE 3PAR StoreServ 8000. Some software titles are licensed per system; other titles are licensed by drive. With system-based licensing one license covers the whole array, independent of configuration or capacity. Drive-based licensing, in contrast, is licensed per installed drive. A software title with drive-based licensing includes two Licenses to Use (LTU), a Base LTU that enables the software feature for the system and a Drive LTU that licenses the use of one drive. For each software title, purchase one Base LTU per title per array, and one Drive LTU, up to a cap, for every drive, that is installed in the array, independent of drive type. For the HPE 3PAR StoreServ 8200 the Drive LTUs cap at 48. For the HPE

Configuration Information

3PAR StoreServ 8400 the Drive LTUs cap at 168. For the HPE 3PAR StoreServ 8440 the Drive LTUs cap at 320. For the HPE 3PAR StoreServ 8450 the Drive LTUs cap at 168. After reaching the cap, you do not need to purchase any more Drive LTUs for that title.

For more information regarding HPE 3PAR software see <http://h18006.www1.hp.com/storage/solutions/3par/software.html>

For more information regarding HPE 3PAR 8000 software SKUs

see: <http://h20195.www2.hp.com/v2/GetPDF.aspx/c04199812.pdf>

HPE 3PAR Operating System Software Suite (Required)

Required for all new HPE 3PAR StoreServ systems, this foundational software suite gives you everything you need to get up and running quickly and efficiently. Powered by HPE 3PAR ASIC, HPE 3PAR StoreServ's Thin Technologies which include HPE 3PAR Thin Provisioning, HPE 3PAR Thin Persistence, HPE 3PAR Thin Conversion and HPE 3PAR Thin Deduplication, form the base of this software suite. Performance acceleration is assured with HPE 3PAR Adaptive Flash Cache by reducing application response time. Network simplification and security are covered with VLAN tagging. Simplified management is offered by HPE 3PAR Operating System, HPE 3PAR StoreServ Management Console, HPE 3PAR Host Explorer and HPE SmartStart software designed to get you off to a quick start with your new HPE 3PAR StoreServ system; HPE 3PAR System Reporter and HPE 3PARInfo software are designed to track performance and capacity utilization trends for multiple HPE 3PAR StoreServ Systems. Other highlights of this suite include HPE 3PAR Full Copy, autonomic rebalancing capabilities that help you optimize the use of future capacity expansions, and support for standard multipathing software for high availability in clustered environments. A one year license for online import is included to enable migration from HPE EVA, EMC CLARiON CX4, EMC VNX/VNX2, EMC VMAX (VMAX, VMAX SE, VMAX 10K, 20K, 40K) or HDS NSC, USP, USP V, USP VM.

HPE 3PAR Replication Software Suite

This suite bundles HPE 3PAR Virtual Copy with HPE 3PAR Remote Copy software, both also sold separately for all HPE 3PAR StoreServ models. HPE 3PAR Virtual Copy software protects and shares data affordably with rapid recovery using reservation-less, non-duplicative, copy-on-write snapshots. HPE 3PAR Remote Copy offers simple and cost effective data protection for efficient multi-tenant disaster recovery.

Also, included in this bundle is Peer Persistence which ensures transparent automatic failover over metropolitan distances using Remote Copy Synchronous mode. The Suite also includes HPE 3PAR Cluster Extension Software which enables automatic failover across data centers using Remote Copy Asynchronous mode.

HPE 3PAR Data Optimization Software Suite

This software bundle combines HPE 3PAR Dynamic Optimization, HPE 3PAR Adaptive Optimization, HPE 3PAR Priority Optimization and HPE 3PAR Peer Motion software together. HPE 3PAR Dynamic Optimization delivers the required service levels for the lowest possible cost throughout the data lifecycle. HPE 3PAR Adaptive Optimization improves storage utilization by enabling cost-optimized storage tiering. HPE 3PAR Priority Optimization assures service levels with QoS controls for mission critical applications. HPE 3PAR Peer Motion enables load balancing at will wherein, movement of data and workloads between arrays is initiated without impacting applications, users or services. The four software titles bundled in this suite are also sold separately for all HPE 3PAR StoreServ models.

HPE 3PAR File Persona Software Suite

This software suite enables rich set of file protocol services, core file data services and an Object Access API to extend the spectrum of primary storage workloads natively addressed by HPE 3PAR StoreServ 8000 Storage. With this solution, the architectural benefits of HPE 3PAR StoreServ 8000 Storage can be extended to use cases such as: home directories and user shares; content management and collaboration; data preservation and governance; and custom cloud applications. This software suite is licensed based on aggregate usable file capacity in increments of 1TB software LTU.

Configuration Information

HPE 3PAR Security Software Suite	This software suite bundles HPE 3PAR Virtual Domains and HPE 3PAR Virtual Lock software. With this suite, you can segregate access and deliver robust storage services for different applications and user groups with additional security attached to the retention of storage volumes.
HPE Smart SAN	HPE Smart SAN for HPE 3PAR makes end-to-end SAN configuration and management simple and autonomic, reducing the probability of errors through SAN automation. It is an application embedded in the SAN components (array, hosts and switches) that enables the 3PAR to automate configuration for settings and policies across the SAN. Smart SAN features enable customers to automate peer zoning, resulting in the creation of fewer zones, and enables configuration of zones in minutes, not hours.
HPE 3PAR Application Software Suite for Hyper-V	Protect your Microsoft Hyper-V environment with HPE 3PAR Recovery Manager for Microsoft Hyper-V and the HPE 3PAR VSS Provider software, included in this software bundle.
HPE 3PAR Application Software Suite for Exchange	This bundle gives you the essentials for use with Microsoft Exchange, including HPE 3PAR Recovery Manager for Exchange and the HPE 3PAR VSS Provider software.
HPE 3PAR Application Software Suite for Oracle	Everything you need for protecting Oracle databases, including HPE 3PAR Recovery Manager for Oracle and Oracle space reclamation capabilities.
HPE StoreOnce Recovery Manager Central	By combining the performance of snapshots with the protection of backups, this software integrates HPE 3PAR StoreServ with HPE 3PAR StoreOnce Backup Systems to provide a converged availability and flat backup service that augments traditional backup processes. With this automated, non-intrusive software, the simplicity and performance of local and remote snapshots can be combined with the reliability and cost effective retention of deduplicated backups
HPE StoreFront Remote SaaS Portal	The HPE StoreFront Remote SaaS Portal provides proactive tools and integrated data collection from the HPE 3PAR StoreServ Storage arrays that call home to deliver unique insights and analytics all in one dashboard. Identify capacity and performance issues early through intuitive capacity and performance trend analysis and forecasting. These valuable analytics help maximize asset utilization and optimize the datacenter with recommendations and remedial actions when issues arise. Users can log into http://www.storefrontremote.com to claim their arrays and get access for free.

Step 9 - Choose File Controller

With HPE 3PAR StoreServ File Controller you get an efficient, bulletproof, and effortless way to provide file storage from HPE 3PAR StoreServ Storage. 3PAR StoreServ File Controller saves you time and money by supporting hundreds to thousands of concurrent users and diverse file workloads. It also has non-intrusive data deduplication that provides an average 50-60% in space savings. It provides security through features such as built-in encryption, sophisticated access controls, online snapshots, and the ability to run endpoint protection and backup software onboard so that data is protected at rest and in flight. HPE 3PAR StoreServ File Controllers are clustered file gateway configurations with transparent failover and online rolling maintenance updates that deliver continuous availability of data to users, servers, and applications. With a straightforward and consistent management experience, it also provides robust capabilities for demanding 24 x 7 file storage environments.

Configuration Information

For more details on the specifications and data services offered by the HPE 3PAR StoreServ File Controller please refer to the following link: <http://www8.hpe.com/h20195/v2/GetHtml.aspx?docname=c04637524>

Technical Specifications

Physical Dimensions	Width in/mm	Depth in/mm	Height in/mm/U	Weight lb/kg
36U 1075mm Intelligent Series Rack	23.54/597.9	44.3/1125.2	68.84/1748.6	428/195
42U 1075mm Intelligent Series Rack	23.54/597.9	44.3/1125.2	79/2006.6	451/205
42U 1200mm Intelligent Series Rack	23.54/597.9	51.19/1300.2	79/2006.6	531/241
47U 1075mm Intelligent Series Rack	23.54/597.9	44.3/1125.2	88.53/2248.7	483/220
HPE 3PAR StoreServ 8200 2N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	47.7/21.6
HPE 3PAR StoreServ 8200 2N Storage Base (with two host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	48.7/22.1
HPE 3PAR StoreServ 8400 2N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	47.7/21.6
HPE 3PAR StoreServ 8400 2N Storage Base (with two host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	48.7/22.1
HPE 3PAR StoreServ 8400 4N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	6.925/175.9/4	97.8/44.4
HPE 3PAR StoreServ 8400 4N Storage Base (with four host adapters, no drives)	19/483	26.6/676.1*	6.925/175.9/4	99.8/45.3
HPE 3PAR StoreServ 8440 2N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	47.7/21.6
HPE 3PAR StoreServ 8440 2N Storage Base (with two host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	48.7/22.1
HPE 3PAR StoreServ 8440 4N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	6.925/175.9/4	97.8/44.4
HPE 3PAR StoreServ 8440 4N Storage Base (with four host adapters, no drives)	19/483	26.6/676.1*	6.925/175.9/4	99.8/45.3
HPE 3PAR StoreServ 8450 2N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	47.7/21.6
HPE 3PAR StoreServ 8450 2N Storage Base (with two host adapters, no drives)	19/483	26.6/676.1*	3.46/87.95/2	48.7/22.1
HPE 3PAR StoreServ 8450 4N Storage Base (no host adapters, no drives)	19/483	26.6/676.1*	6.925/175.9/4	97.8/44.4
HPE 3PAR StoreServ 8450 4N Storage Base (with four host adapters, no drives)	19/483	26.6/676.1*	6.925/175.9/4	99.8/45.3
HPE 3PAR StoreServ 8000 SFF(2.5in) SAS Drive Enclosure (without drives)	19/483	24.8/630.7	3.46/87.95/2	33.5/15.2
HPE 3PAR StoreServ 8000 LFF(3.5in) SAS Drive Enclosure (without drives)	19/483	24.9/631.4	6.89/175/4	42.9/19.5
SFF SAS drive with carrier	0.78/19.9**	896/227.7	3.42/86.85**	0.7/0.32 (varies by type)
LFF SAS drive with carrier	4.36/110.84**	8.67/220.26	1.18/30.0**	1.9/0.86 (varies by type)
Service Processor				37.48/17

Technical Specifications

Power Requirements

Input Voltage (VAC)

100 - 240 VAC

Frequency (Hz)

50 - 60

Component	Idle (watts / BTU/hr)	Transactional (watts / BTU/hr)
Node Pair (8200 or 8400), no drives, no add-on host adapters	236 / 803	398 / 1357
Node Pair (8440 or 8450), no drives, no add-on host adapters	344 / 1173	363 / 1238
4-port 16Gb/s Fibre Channel Adapter	18.61 / 63.5	19.13 / 65.3
2-port 10Gb/s iSCSI/FCoE Adapter	34 / 115.8	40 / 136.4
2-port 10Gb/s Ethernet Adapter	5.69 / 19.4	5.71 / 19.5
4-port 1Gb/s Ethernet Adapter	1.97 / 6.7	1.97 / 6.7
8000 SFF(2.5in) SAS Drive Enclosure, no drives	150 / 512 (average)	150 / 512 (average)
8000 LFF(3.5in) SAS Drive Enclosure, no drives	164 / 559 (average)	164 / 559 (average)
300GB 15K Small Form Factor HDD	6.7 / 22.9	6.9 / 23.1
600GB 15K Small Form Factor HDD	7.0 / 24.0	7.3 / 25.1
600 GB 10K Small Form Factor HDD	6.3 / 21.4	7.4 / 25.2
1.2TB 10K Small Form Factor HDD	6.2 / 21.1	8.2 / 27.9
1.8TB 10K Small Form Factor HDD	7.3 / 24.8	7.5 / 25.6
2TB 7.2K Small Form Factor NL HDD	6.1 / 20.8	7.2 / 24.6
2TB 7.2K Large Form Factor NL HDD	7.5 / 25.6	10.6 / 36.1
4TB 7.2K Large Form Factor NL HDD	9.1 / 31.1	13.1 / 44.6
6TB 7.2K Large Form Factor NL HDD	11.9 / 40.7	14.3 / 48.74
8TB 7.2K Large Form Factor NL HDD	10.8 / 36.8	14.9 / 50.7
400GB Small Form Factor SSD	3.3/11.3	5.8/19.8
480GB Small Form Factor SSD	2.2 / 7.5	5.5 / 18.7
480GB Small Form Factor nonAFC SSD	3.2 / 11	7.1 / 24.4
480GB Large Form Factor nonAFC SSD	1.8 / 6.1	6.7 / 22.9
920GB Small Form Factor FIPS SSD	2.2 / 7.5	5.5 / 18.7
1920GB Small Form Factor SSD	3.5 / 11.9	8.9 / 30.5
3840GB Small Form Factor SSD	3.4 / 11.6	11.0 / 37.5
7680GB Small Form Factor SSD	3.4 / 11.6	13.7 / 46.74

Technical Specifications

Environmental Specifications⁴

Operating Temperature	41° to 104° F (5° to 40° C) - Reduce rating by 1° F for each 1000 ft altitude (1.8° C/1,000 m)
Shipping Temperature	-30° to 60°C (-22 to 140°F). Maximum rate of change is 20°C/hr (36°F/hr)
Altitude (ft/m) max.	10,000 ft / 3,048 m
Shipping Altitude (ft/m) max.	40,000ft/ 12,192 m
Humidity	10% to 90% non-condensing
Shipping Humidity	10% to 90% non-condensing
Operating Vibration	0.25 G, Sine, 5-500 Hz, 0.1 Grms, Random 10-100Hz
Non-operating Vibration	0.5 G, 5 - 500 Hz, Sine
Operating Shock	2 G, 11ms, half-sine
Non-operating Shock	10 G, 11ms, half-sine
Maximum Air Flow	Storage Base and Upgrade Node Pair - 109 CFM per enclosure 8000 SFF(2.5in) SAS Drive Enclosure - 105 CFM 8000 LFF(3.5in) SAS Drive Enclosure - 109 CFM

Electromagnetic Compatibility	CISPR 22:2008/ EN55022:2010 Class A CISPR 24:2010/ EN 55024:2010 IEC 61000-3-2:2005/ EN 61000-3-2:2006 +A1:2009 +A2:2009 IEC/ EN 61000-3-3:2008 AS/NZS CIPSR 22: 2009 Class A CNS 13438:2006 Class A 47 CFR Part 15 Subpart b Class A ICES-003 Issue 5 Class A V-3/2014.04 RRA Notice No. 2014-8 (2014.06.23) & 2014-37 (204.06.23) Class A RRA Notice No. 2014-9 (2014.06.23) & 2014-38 (2014.06.23)
-------------------------------	---

Acoustics Sound pressure level measured per ISO 7779 specs during normal operating fan conditions, from a minimum of 3,000 RPM to a maximum of 10,000 RPM	Fan Speed (RPM)	8200/8400	8400 4N	8440 2N	8440 4N	8000 2U	8000 4U
		2N Storage Base	Storage Base	Storage Base	Storage Base	SAS Drive Enclosure	SAS Drive Enclosure
	Minimum	63.8	67.2	72	74	62.6	61.3
	Maximum	93.4	96.5	93	97	85.4	88

Safety	IEC 60950-1:2005 (2nd Edition); Am 1:2009 EN 60950-1:2006 +A11:2009+A12 EN 62479:2010 CNS 14336-1 2nd Edition UL 60950-1 2nd Ed. CAN/CSA C22.2 No. 60950-1
--------	---

NOTE: ⁴ Specifications are subject to change without notice.

Certifications / Markings	cTUVus Mark TUV T-mark (EN 60950) CE Mark FCC Class A IC Class A	KCC GOST-R C-Tick WEEE China RoHS EU RoHS
---------------------------	--	---

Technical Specifications

VCCI Class A
BSMI Class A

Summary of Changes

Date	Version History	Action	Description of Change
15-Aug-2016	From Version 9 to 10	Added	Added DAC cables support
07-Jun-2016	From Version 8 to 9	Changed	Added new drives (7.68TB SSD, 8TB NL FE HDD) Added Smart SAN
15-Apr-2016	From version 7 to 8	Changed	Updated SSD descriptions
8-April-2016	From version 6 to 7	Changed	Changes made throughout the QuickSpecs
31-March-2016	From version 5 to 6	Changed	Added new drives (2TB NL FE HDD, 8TB NL HDD) Added CSI section Added optical splitter cables (K2Q46A, K2Q47A)
16-Feb-2016	From Version 4 to 5	Changed	Changes made throughout the QuickSpecs
04-Dec-2015	From Version 3 to 4	Changed	Added HPE 3PAR StoreServ 8200 Converged File/Block Starter Kit. Added new drives (400GB SSD, 3.84TB FE SSD, 4TB NL FE HDD).
02-Oct-2015	From Version 2 to 3	Added	Added The HPE StoreFront Remote SaaS Portal
28-Sept-2015	From Version 1 to 2	Changed	Changes made throughout the QuickSpecs
24-Aug-2015	Version 1	Created	Create first version



Sign up for updates

© Copyright 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.



c04607918 - 15247 - Worldwide - V10 - 15-August-2016