



CATE Storage and Backup Engineering
Citigroup Architecture & Technology Engineering

NetApp FAS3200 Series Storage Controller (Filer)

CSBE-15012-RN

Release Notes

Version 3.2

©Citigroup Inc. 2015, All Rights Reserved.

This document contains material confidential to Citigroup. The contents are protected by specific employee and contractor agreements regarding confidentiality. Reproduction for distribution outside of Citigroup is prohibited.

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	15-APR-2011	Richard Sopp	Stephen Hopkins	Peer review
0.2	20-APR-2011	Richard Sopp	CTI NAS Ops	Initial draft release
1.0	28-APR-2011	Richard Sopp	CTI NAS Ops	Initial production release
1.1	10-JUN-2011	Richard Sopp	n/a	Updated to ONTAP 8.0.1P5
1.2	15-JUL-2011	Richard Sopp	n/a	Updated SP firmware
1.3	29-NOV-2011	Richard Sopp	n/a	Updated Service Image
1.4	21-DEC-2011	Richard Sopp	n/a	Capped SP firmware at v1.2.2 as v1.2.3 requires ONTAP 8.0.2P3 or later
2.0	03-JAN-2012	Richard Sopp	n/a	Updated for ONTAP 7.2.6P2 and ONTAP 8.0.2P4
2.1	10-JAN-2012	Richard Sopp	Tim Armstrong Stephen Hopkins Jackson Tan	CTI RFT review
2.2	30-JAN-2012	Richard Sopp	Stakeholders	Production release
2.3	04-JUN-2012	Richard Sopp	n/a	Added bug 536445
2.4	22-JUN-2012	Richard Sopp	n/a	Updated SP firmware to v1.3
2.5	17-SEP-2012	Richard Sopp	n/a	Reverted SP firmware to version 1.2.3 Added procedures for FCMT0 flag issue and validating NV-MEM battery version
2.6	01-MAY-2013	Richard Sopp	n/a	Added bug 519766
2.7	24-MAY-2013	Richard Sopp	n/a	Update SP firmware to v1.3.2
2.8	4-NOV-2013	Tan Jackson	Richard Sopp	Added FAS3220 and FAS3250 platforms
2.9	29-NOV-2013	Tan Jackson	n/a	RFT release
3.0	12-MAY-2014	Richard Sopp	n/a	Production Release
3.0.1	12-MAY-2014	Richard Sopp	n/a	Update SP firmware to v1.3.3
3.1	30-MAY-2014	Richard Sopp	n/a	Production Release
3.1.1	27-APRIL-2015	Tan Jackson	Richard Sopp	Update SP firmware v1.3.3P2
3.1.2	18-MAY-2015	Tan Jackson	n/a	Release for Testing

3.2	1-JUNE-2015	Tan Jackson	n/a	Release for Production
-----	-------------	-------------	-----	------------------------

2 Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, CITIGROUP and the Arc Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

3 Contact

Please address any comments or questions regarding the information contained in this document via email to the “CATE Storage and Backup Engineering” distribution list or alternatively to gdse@citi.com.

4 Latest Versions

It is recommended that the latest version of the document be used from the CATE Storage and Backup Engineering website <https://catecollaboration.citigroup.net/domains/platstor/gdse/Pages/default.aspx>. Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5 Regional Reviewers

The purpose of this section is to list all the stakeholders who reviewed this document during the Request for Comment process.

Region	Reviewed by
Asia Pacific	No response received
EMEA	No response received
Japan	No response received
Latin America / Mexico	No response received
North America	No response received

Table 1 - Regional stakeholder review

Table of Contents

1	Version Control.....	2
2	Legal Notice	4
3	Contact.....	4
4	Latest Versions.....	4
5	Regional Reviewers.....	4
6	Introduction.....	6
6.1	Operating System dependencies	6
6.2	Specification comparison between 3210, 3220, 3240, 3250 and 3270 models	8
6.3	Specification comparison between 3250, 3170 and 6250 models.....	11
6.4	Important considerations when choosing a system configuration.....	12
7	Best Practices for ONTAP 8 / 7-mode	13
7.1	Connecting to Power Distribution Units	13
7.2	Firmware requirements.....	13
7.3	The Service Processor (SP)	14
7.4	The Management Port (e0M)	14
7.5	The Alternate Control Path port (e0P)	15
7.6	PCIe card slot assignments.....	16
7.7	Clustering	17
7.8	Storage.....	18
7.9	Network	23
7.10	Tape connectivity (Fiber Channel)	24
7.11	Virtual Servers (vFilers)	24
7.12	Flash Cache.....	25
7.13	Memory management for replication threads.....	25
7.14	Shared chassis resources.....	25
7.15	Maintenance best practices	27
8	Known Issues	28
8.1	Issues fixed in this release	29
9	Elevations.....	31
9.1	Power Distribution Unit Connectivity	31
9.2	Shared Chassis elements.....	31
9.3	Standalone controllers (un-clustered).....	32
9.4	Clustered controllers.....	34
10	References	37
11	Appendices.....	38
11.1	Service Image (BIOS) upgrade procedure.....	38
11.2	Service Processor Firmware upgrade procedure	43
11.3	Service Processor Firmware downgrade procedure	46
11.4	Procedure to validate NV-MEM battery firmware version prevents bug 536445	46
11.5	NV-MEM battery firmware upgrade procedure (bug 536445)	46
11.6	Procedure to validate the status of the Service Processor FCMT0 flag	47
11.7	Procedure to reset FCMT0 flag	48
11.8	Specification comparison between 3210, 3220, 3240, 3250 and 3270 models (for ONTAP prior to 8.1.4)	50
11.9	Root Volume Recommendation for ONTAP 7.3.6/ ONTAP 8.0.2.....	51
11.10	Maximum Volume (Single) for ONTAP 7.3.6/ ONTAP 8.0.2	51
11.11	Concurrent replication process maximums for ONTAP 7.3.6/ ONTAP 8.0.2	51
12	Document Version Control	52
13	Test Results	54
13.1	Physical Tests.....	54
13.2	Research Tests.....	68

6 Introduction

NetApp Filers are one of the two NAS platforms deployed globally within Citi. They are connected to the IP network and used to store various types of data via the CIFS (Windows) and NFS (AIX, ESX, Linux and Solaris) protocols. The storage controller or Filer is the main server hardware component that contains CPUs, Memory, NICs, SAS HBAs and cache. It runs the NetApp Data ONTAP operating system and connects the NetApp disk subsystem to the Citi IP network.

The FAS3200 series, consisting of the 3210, 3220, 3240, 3250 and 3270 models, is the latest generation of NetApp's mid-range storage platform offering well suited to the bulk of NAS workloads within Citi. Direct comparisons with previous mid-range models are somewhat challenging due to the different configuration permutations and flexibility offered with the introduction of the IO Expansion Module (IOXM) which provides the ability increase the number of PCI expansion slots on the system. In an apples to apples comparison with its 3100 series predecessor the 3200 series offers a 50% decrease in form factor (3U, down from 6U), a 33% decrease in weight (36.3kg/80 lbs, down from 54.89kg/121 lbs) and a 25% decrease in power and thermal requirements. Depending on workload type and model the FAS3200 series offers anywhere from 1.2x to 1.8x the performance of the previous 3100 series platform, and a 15% increase in the maximum storage capacity.

6.1 Operating System dependencies

The following operating system version (ONTAP) dependencies apply when considering use of the FAS3210, 3220, 3240, 3250 and 3270. The following table demonstrates which hardware platform is supported on ONTAP 8/7-mode, which is the preferred ONTAP release for the FAS3200 series.

Supported Platform	ONTAP 7.3.x Support	ONTAP 8/7-mode Support	Recommended ONTAP 8.x version (Citi)
FAS3210	Y	Y	8.1.4P6
FAS3220	N	Y	8.1.4P6
FAS3240	Y	Y	8.1.4P6
FAS3250	N	Y	8.1.4P6
FAS3270	Y	Y	8.1.4P6

Table 2 - Hardware platform by Minimum Operating System Version Support Matrix

6.1.1 ONTAP 8 / 7-mode

The FAS3200 series is approved for use within Citi as part of an ONTAP 8.x 7-mode where the price / performance ratio of this mid-range hardware platform is required.

It is recommended, where possible, that FAS3200 series system should be deployed with ONTAP 8.x in order to support 64-bit aggregates thereby simplifying capacity management and scaling of the disk pool. Deploying a FAS3200 series with ONTAP 7.3.6 will negate the opportunity to leverage 64-bit aggregates (introduced as part of ONTAP 8.x) when the systems is installed and will add operational complexity at a later date when 32-bit aggregates will require online conversion to the 64-bit format.

6.1.2 ONTAP 8 / Cluster mode (c-mode)

The FAS3200 series is **not approved** for use within Citi as part of an ONTAP 8.x cluster mode (c-mode) system at this time. This restriction will be lifted when cluster mode is made available by CATE as a certified solution.

6.1.3 ONTAP GX

NetApp does not support the FAS3200 series with the ONTAP GX release family, and it is not possible to install this version of ONTAP on this hardware platform. Logically it follows that the FAS3200 series is **not approved** for use within Citi running ONTAP GX.

6.2 Specification comparison between 3210, 3220, 3240, 3250 and 3270 models

The previous section provided a comparison of environmental specifications between the FAS3100 and FAS3200 series whereas this section focuses more on hardware specifics including the differences in compute potential between each model in the FAS3200 series.

Feature	3210	3220	3240	3250	3270
CPU Vendor	Intel	Intel	Intel	Intel	Intel
Codename	Wolfdale	Harpertown	Harpertown	Harpertown	Wolfdale
Branding	Xeon	Xeon	Xeon	Xeon	Xeon
Model	E5220	L5410	L5410	L5410	E5240
Cores	2	4	4	4	2
Clock rate (core)	2.33 GHz	2.33 GHz	2.33 GHz	2.33 GHz	3.0 GHz
Cache	6MB (L2)	12 MB (L2)	2 x 6MB (L2)	12 MB (L2)	6MB (L2)
Total CPU count ¹	2	2	2	4	4
Total core count ²	4	8	8	16	8
Controller architecture	64-bit	64-bit	64-bit	64-bit	64-bit
RAM	5GB	12GB	8GB	20GB	20GB ³
NV-MEM	512MB	1.6GB	1GB	2GB	2GB
Number Native FC ports (speed)	2 (4GB SFP)	2 (4GB SFP)	2 (4GB SFP)	2 (4GB SFP)	2 (4GB SFP)
Number Native 1GB Ethernet ports	2	2	2	2	2
Number Native 10GB Ethernet ports	0	0	0	0	0
Expansion slots	2 PCIe	2 PCIe (without IOXM) 6 PCIe (with IOXM)	2 PCIe (without IOXM) 6 PCIe (with IOXM)	6 PCIe (with IOXM)	2 PCIe (without IOXM) 6 PCIe (with IOXM)
FlashCache (maximum)	256GB	512GB	512GB	1TB	1TB
Maximum raw disk capacity (TB) (HA)	720	1920	2400	2880	3840
Maximum number of SAS/BSAS (SATA) drives (HA)	240	480	600	720	960
Maximum raw 64-bit aggregate size (TB)	75	90	90	105	105

Table 3 - Detailed comparison of FAS3210, FAS3220, FAS3240, FAS3250 and FAS3270

¹ Each figure is a per Cluster amount. These figures should be halved for a standalone configuration.

² Each figure is a per Cluster amount. These figures should be halved for a standalone configuration.

³ NVMEM on FAS/V3270 leverages a portion of the 20GB of controller memory, resulting in 18GB memory for Data ONTAP. Previously published memory capacities for the FAS/V3200 series controllers were based on usable memory estimates and have been changed as of 11/05/2012 to represent the physical memory in the controller.

The following attributes are common to all models and all configurations:

Width	19" IEC rack-compliant (17.61", 44.7cm)	
Depth	24" (60.7cm) - (26.5" w/cable management) (67.3 cm)	
Operating Temperature, Altitude, and Relative Humidity	10° C to 40° C (50° F to 104° F); at <= 3,000 m (at <= 10,000' feet) elevation; 20% to 80% relative humidity, noncondensing (28° C wet bulb temperature)	
Non-operating Temperature and Relative Humidity	-40° C to 65° C (-40° F to 149° F); 5% to 95% relative humidity, noncondensing, in original container	
Operating Acoustic Noise	≤ 61 dBA sound pressure (LpA) @ normal operating conditions (at 22°C and at sea level)	
Min cooling clearances:	Front: 10in (25.4cm)	Rear: 12in (30.5cm)
Min service (maintenance) clearances:	Front: 30in (76.2cm)	Rear: 30in (76.2cm)

6.2.1 Configuration types, rules and model naming conventions

The following table is designed to summarize the different supported configuration types within the FAS3200 series as well as the naming conventions associated with each configuration.

FAS3200 series model			FAS 3210	FAS 3220	FAS 3240	FAS 3250	FAS 3270
Stand-alone (single chassis)	Single controller / no expansion	Supported	Y	Y	N	N	N
		Designation	FAS3210	FAS3220	n/a	n/a	n/a
	Expanded single controller	Supported	N	Y	Y	Y	Y
		Designation	n/a	FAS3220E	FAS3240E	FAS3250E	FAS3270E
HA-Pair ⁴ (single and dual chassis)	Single chassis / no expansion	Supported	Y	Y	Y	N	Y
		Designation	FAS3210A	FAS3220A	FAS3240A	n/a	FAS3270A
	Dual chassis / expanded controllers	Supported	N	Y	Y	Y	Y
		Designation	n/a	FAS3220AE	FAS3240AE	FAS3250AE	FAS3270AE

Table 4 - Names of all valid FAS3200 configurations

⁴ Formerly known as a cluster

The following list summarizes some of the key aspects of the configuration table above:

- Expansion modules (IOXM) are not supported on the FAS3210
- Dual chassis with expansion is only available with the FAS3220AE, FAS3240AE, FAS3250AE and FAS3270AE
- The 3250 does not come in a single chassis / dual controller configuration with no expansion modules (unlike it's predecessors the 3240 and 3270)
- Confusingly the naming conventions of these different models and configurations are different to the FAS6200 high-end range (which neglects to append an 'E' to designate the presence of an IOXM)

6.3 Specification comparison between 3250, 3170 and 6250 models

The FAS3250 has been billed as a platform transition path for a FAS3170 that is either experiencing performance bottlenecks and requires additional compute / cache resources or a system that requires more connectivity (either Ethernet or FiberChannel) and has exhausted all available PCI expansion slots on the FAS3170. In certain situations the FAS6220 is offered as the successor to the FAS3170 in the high-end series. The purpose of the following table is to provide a convenient cross reference of these different platforms and their attributes as they tend to all overlap in a certain segment of the FAS product family.

Feature	3170	3250	6220
CPU			
Vendor	AMD	Intel	Intel
Codename	Santa Rosa	Harpertown	Nahelem-EP
Branding	Opteron	Xeon	Xeon
Model	2218	L5410	E5507
Cores	2	4	4
Threads	8	4	4
Clock rate (core)	2.6 GHz	2.33 GHz	2.26 GHz
Cache	2x1MB (L2)	12 MB (L2)	4 MB (SMART)
CPU Count	2	2	2
Controller architecture	64-bit	64-bit	64-bit
RAM	16GB	20GB	48GB
NV-RAM	2GB	2GB (NV-MEM)	4GB
Number Native FC ports (speed)	4 (4GB)	2 (4GB SFP)	4 (8GB SFP+)
Number Native 1GB Ethernet ports	2	2	2 (1Gbps) 4 SFP+ (10Gbps)
Number Native 10GB Ethernet ports	0	0	4 (SFP+)
Expansion slots	4	6 PCIe (with IOXM)	6 PCIe ⁵
FlashCache (maximum)	1TB	1TB	3TB
Maximum raw disk capacity (TB) (HA)	2520	2880	4800
Maximum number of SAS/BSAS (SATA) drives (HA)	840	720	1200
Maximum raw 64-bit aggregate size (TB)	105	105	162

Table 5 - Detailed comparison of FAS3170, FAS3250 and FAS6220

⁵ The FAS6220 controller has 6x physical expansion slots (2x vertical slots and 4x horizontal slots). While the horizontal slots support a wide mix of adapters, the vertical slots are limited to the NVRAM, 8Gb FC, and 6Gb SAS adapters that cannot be used in horizontal slots.

6.4 Important considerations when choosing a system configuration

Outside of the CPU and memory specifications on each FAS3200 model, the next factor in terms of influencing the choice of system configuration is the number of PCIe slots required for disk and network connectivity, as well as FlashCache. The following best practices, section [6.4.2](#), is designed to provide the reference material necessary to calculate the number of PCIe expansions slots required for a given configuration.

6.4.1 Populate the system upfront with all the necessary PCIe adapters

While new cards can be inserted in a system at any time after the system has been commissioned for production use, the non-disruptive installation process will require a cluster takeover and giveback (or system shut down for a standalone system), which is sometimes difficult to obtain business unit approval for the necessary change control.

Instead think about the requirements of the system over the life of the system if you know the system requirements will increase after the first few years of use. For example consider a system that requires just two SAS HBAs per controller to start with, but whose storage requirements will increase at the end of the second year of use requiring more SAS shelves than can be supported by the two HBAs. Instead of procuring the additional HBAs later in the life cycle of the system, installing them upfront will reduce the requirement for additional maintenance windows at a later date.

6.4.2 FAS3210 as a CIFS and/or NFS File Server

The FAS3210A warrants a special mention as it has a unique set of attributes and may appeal as a cost effect small system for remote offices that desire a more redundant SAS connectivity configuration than the single path configuration offered by the FAS2040. It could also be used in strategic data centers where business units require physical system isolation and cannot store their data on a traditional shared NAS storage device.

As the entry level model, contention for PCIe v2 expansions slots is significant as there are only two on each controller in the cluster. The recommended use of these expansions slots is 1 x quad port SAS HBA and 1 x quad port NIC in order to provide hardware resiliency when paired with the on-board SAS and Ethernet ports. Having exhausted the expansion slots on the system any discussion of FlashCache is a moot point, however for such a small system the fundamentals of disk and network connectivity should be considered first making a SATA + FlashCache configuration unviable on this model.

Raw storage capacity of the FAS3210 is limited to 10 SAS shelves (240 drives) and it is recommended that these be configured as a single SAS stack to ensure that SAS connectivity meets the best practices described in section [7.8.1](#).

7 Best Practices for ONTAP 8 / 7-mode

This section presents the recommended configurations and scaling for the 3200 series when running Data ONTAP 8.x/7-mode.

7.1 Connecting to Power Distribution Units

The FAS3200 series chassis contains two power supply units (PSUs). To ensure system availability, connect each PSU to a separate Power Distribution Unit (PDU) within the data center. Refer to section [9.1](#) for a diagram depicting this relationship.

7.2 Firmware requirements

There are two hardware components on the FAS3200 series that are capable of receiving a field firmware upgrade. To upgrade the firmware follow the procedures as published on the NetApp support website (support.netapp.com). The service processor (SP) firmware upgrade can be performed without impacting the Filer's ability to serve data as the SP is for management traffic only. However the system board firmware upgrade will require a cluster failover maintenance event in order to reboot the Filer to the latest level of firmware.

Component	Firmware Type	Minimum Firmware	Recommended Firmware
System board	Service Image (BIOS)	Version 5.2.1 File: 30802413_7mode Release date: 26 Apr 2013	Version 5.2.1 File: 30802413_7mode Release date: 26 Apr 2013
Service Processor	Service Processor Image	Version: 1.3.3P2 File: 30602977_ontap Release date: 19-Dec-14	Version: 1.3.3P2 File: 30602977_ontap Release date: 19-Dec-14 ⁶

Table 6 - Recommended Firmware Levels

For SAS and ACP firmware minimum and recommended levels refer to the [DS4243](#) and [DS2246](#) disk shelf release notes as appropriate.

Note ESH2, ESH4 and AT-FCX shelf firmware levels are not applicable to the FAS3200 series as the legacy DS14 enclosures are not certified by CATE for use with this hardware platform.

As a FPGA device (field-programmable gate array) the FlashCache card could also be considered as running a "type" of firmware, however these firmware images are included as part of an ONTAP release and will not be distributed or upgraded separately (unlike system board, service processor and disk shelf firmware).

⁶ File listed for the image is for installation from the Data ONTAP prompt.

7.3 The Service Processor (SP)

The service processor (SP) is NetApp's next generation Remote LAN Management (RLM) and Baseboard Management Controller (BMC) offering remote management functionality for the Filer. The SP performs Field Replaceable Unit (FRU) reporting, FRU tracking; monitors and reports on internal sensors; reports on and can set LED states; and actively manages fan speed. The SP only uses SSH and does not support Telnet or RSH.

The SP and e0M management interfaces share the same physical RJ45 port on the center of the storage controller's rear panel which is indicated by the wrench or spanner icon above it. The two interfaces can share the same port due to a small network switch embedded in on the controller's system board.




Figure 1 - SP (remote management device) and e0M network connections

Refer to the Data ONTAP Solutions Guide for the procedure to configure the Service Processor and for the network reference architecture.

7.4 The Management Port (e0M)

All models in the FAS3200 series feature an e0M interface that uses the same physical RJ45 port as the SP located on the center of the storage controller's rear panel which is indicated by the wrench or spanner icon above it. The e0M port allows for the separation of data and management traffic on the storage system and supports all the management protocols which are available on other interfaces on the storage system (SSH, HTTPS and SNMP).

Refer to the Data ONTAP Solutions Guide for the procedure to configure the e0M management port and for the network reference architecture.

7.5 The Alternate Control Path port (eOP)

Unlike the 3100 series, the 3200 series includes an on-board out-of-band disk subsystem management port, designated eOP, meaning that additional NICs or Ethernet ports do not need to be ordered to connect the private ACP (alternate control path) network.

Although connectivity of the eOP port to the NetApp SAS shelves is not required for the storage system to serve data, implementation of the private ACP network between all controllers and SAS shelves in a storage system is strongly recommended for system resiliency, hardware recovery, disk enclosure log collection, and the ability to remotely manage the disk enclosure thereby reducing the need for a physical visit to the data center to reboot a 'wedged' enclosure (for example).

When you configure ACP on one of the system's network interfaces, you must supply a private domain name that conforms to RFC1918. The recommended approach when configuring eOP is to use the system default domain, 192.168.0.0/22 (255.255.252.0).. If you are configuring ACP for disk shelves attached to an active/active configuration, you must supply the same ACP domain name and network mask for both systems.

Do not connect the ACP port to a routed network, and do not configure switches or hubs between the ACP port and the designated Ethernet port. Doing so is not supported and will cause interruptions in service.

For full details on the ACP network and cabling requirements refer to the CATE [DS4243](#) and [DS2246](#) Release Notes.

Consult the latest ONTAP Solutions Guide for best practices regarding the configuration of the eOP interface, including entries in the `/etc/rc` and `/etc/hosts` files.

7.6 PCIe card slot assignments

NetApp publishes System Configuration Guides on their hardware universe site (hwu.netapp.com). This site enables users to choose the hardware specifications that users wish to view and compare them to other platforms. Users can check on details for which PCIe cards should be in which expansion slots on the system. As part numbers for these PCIe cards change frequently this information cannot be embedded in this document as these Release Notes would continually require updating. The information for ONTAP 8.1.4 is available on <http://hwu.netapp.com/Adapter/Index>⁷

- Select the system configuration guide from the above [link](#) that relates to the particular FAS3200 series model you are interested in
- When the site loads select “Adapters” from the menu on top and do a search by ONTAP and Model
- From the filters options, select the ONTAP version, platform, adapter types and specifications. Click on “Show Results”
- For the 3200 series the following assignment priorities are universal:
 - Populate the expansion slots as required according to the expansion adapter requirements of your system
- If your candidate expansion slot is already populated with an expansion adapter then locate the next slot of highest available preference for the adapter in question. If this is already occupied then repeat this step and so on (Refer to *Priority Slot Assignment* column on site)
- If you are installing more than one card of a certain type in the system populate all the slots for that card type using the process in the steps above before moving on to the next card type in the list
- If you run out available expansion slots for the number of expansion adapters you are planning to put in the system then the storage controller model is under specification for your requirements and you should look at the next controller model up the series
- If you reach the end of the adapter list or have allocated all the expansion adapters you were considering putting in the system then the assignment process is complete.

⁷ The Hardware Universe (HWU) application is NetApp’s new system configuration solution, replacing the System Configuration Guides

7.7 Clustering

Clustering offers the ability to failover network, compute and disk resources from one storage controller to another.

High availability is delivered in the FAS3200 series with on-board NV-MEM (DIMMS) as opposed to an integrated NV-RAM / cluster interconnect expansion adapter. At the physical media level the cluster interconnect requires TwinAx cabling, which should be ordered with the system (see section [7.7.1](#) for special circumstances with the FAS3210A, 3220A, 3240A and 3270A systems⁸), and at the protocol level runs iWarp unlike the previous mid-range series. The connection speed of the interconnect is 10 GBPS. The interconnect ports are designated c0a and c0b and can be found to the immediate right of the on-board SAS ports on the storage controller.

7.7.1 FAS3210A, FAS3220A, FAS3240A and FAS3270A (Single Chassis Cluster)

On the FAS3200 series where the clustered configuration is delivered as a single chassis FAS3210A or FAS3220A or FAS3240A or FAS3270A there are no TwinAx cluster interconnection cables required. All cluster interconnectivity is achieved via the shared chassis backplane.

In terms of naming the storage controllers within each chassis follow a standard approach throughout the environment to ensure ease of communication between data center site support personnel and the storage administration team. The recommended approach is to have the storage controller at the top of the chassis with a lower designation than the one in the bottom bay such that Filer1 would be in the top bay and Filer2 in the bottom bay (as an example).

7.7.2 FAS3220AE, FAS3240AE, FAS3250AE and FAS3270AE (Dual Chassis Cluster)

On the FAS3220, FAS3240, FAS3250 and FAS3270 where the clustered configuration is delivered as a dual chassis FAS3220AE, FAS3240AE, FAS3250AE and FAS3270AE (respectively) two physical 10GbE TwinAx interconnection cables are required. Errors were found when testing a cluster interconnect cross connection on this hardware platform and as such it is strongly recommended that cluster interconnects are cabled in parallel such that port c0a on node A connects to port c0a on node B, and port c0b on node A connects to port c0b on node B.

In terms of naming the storage controllers in a dual chassis configuration follow a standard approach throughout the environment to ensure ease of communication between data center site support personnel and the storage administration team. When both chassis are in the same cabinet the recommended approach is to have the storage controller towards the top of the cabinet with a lower designation than the one towards the bottom of the cabinet, such that Filer1 would be above Filer2 in the same cabinet.

The 10GbE TwinAx Cluster interconnect cables are available in 0.5 meter, 2 meter and 5 meter lengths which facilitates the placement of chassis in different cabinets (likely to be adjacent cabinets with a 5 meter cable length limitation).

⁸ FAS3250 does not come with High Availability in a single chassis – i.e. FAS3250A is not part of product offering from NetApp

7.8 Storage

The FAS3200 series is equipped with two on-board SAS ports that are serviced by a single ASIC. To ensure maximum system availability, follow the SAS cabling and HBA requirements outlined in section [7.8.1](#). Consult the appropriate NetApp system configuration guide for your 3200 series model to determine the preferred placement of SAS HBAs within the system.

7.8.1 Cabling

Detailed instructions for SAS cabling can be found in the [DS4243](#) and [DS2246](#) CATE Release Notes. The following table is designed to assist in determining how many SAS HBAs are required per controller based on the number of SAS stacks required. The calculations are based on the constant that there are 24 disks in both the DS4243 and DS2246 shelves, and that you can have a maximum of 10 shelves in each SAS stack.

FAS3200 model	Max # drives	Max # SAS stacks	Required # SAS HBAs
FAS3210	240	1	1
FAS3220	480	2	2
FAS3240	600	2.5	2
FAS3250	720	3	2
FAS3270	960	4	2

Table 7 – SAS HBA requirements per storage controller

Summary:

- Connectivity for one SAS stack can be achieved with 1 HBA and using 1 of the on-board SAS ports
- For connectivity to more than one SAS stack, two SAS HBAs per controller is the minimum CATE supported SAS connectivity configuration.
- Connectivity for up to 4 SAS stacks can be achieved by 2 SAS HBAs per controller.
- For the PCIe expansion slot assignments for the SAS HBAs consult section [7.6](#)
- For the FAS3240 note that the maximum SAS stack count is 2.5 stacks. This means there can only be 5 shelves in the third SAS stack on the system.

At a high level the following cabling principles should be observed to maintain the highest level of system availability and apply to irrespective of whether the configuration is a HA Pair or a standalone storage controller configuration:

- Avoid HBAs as a single point of failure. Do not connect both paths of one SAS stack back to a single HBA. For each path of the two paths from a storage controller to a SAS stack, ensure that each path is connected to the storage controller via a separate HBA or on-board SAS port.
- Avoid ASICs as a single point of failure. On a quad port HBA there are two ASICs, one for ports A and B and one for ports C and D. If the previous principle is adhered to this rule will naturally follow as well. Ensure that both paths from a storage controller to a SAS stack are not serviced from the same ASIC. The on-board SAS ports on a FAS3200 series are serviced by a single ASIC therefore when implementing SAS stack counts greater than one CATE recommends two SAS HBAs and foregoing the use of the on-board SAS ports.

7.8.2 Capacity

The following tables are provided to assist in determining the hardware components necessary for a given raw storage capacity requirement.

Do not mix shelf types in the same SAS stack. The DS4243 has a 3GB SAS backplane and the DS2246 has a 6GB SAS backplane. Adding a DS2246 to a stack that it comprised of DS4243s will force the DS2246 backplane to negotiate down to 3GB.

7.8.2.1 NDMP considerations

At some point within the lifecycle of the data on the storage system it will require backing up. The most common form of NAS backup in Citi is NDMP. There are scaling limitations with the use of NDMP that may artificially limit the raw storage capacity of your FAS3200 system. Storage Administrators are requested to check with their local backup teams to confirm the supported maximum capacity of the storage system permitted to allow a successful NDMP backup schedule. Remember that production systems often replicate to a DR system where the data is then backed up via NDMP so these scaling requirements apply to both the production system as well as the DR system.

If NDMP backups are not required as part of the intended design, consult the following tables for maximum supported storage system capacity for each model in the FAS3200 series.

7.8.2.2 DS4243 shelf

Support for the SAS based DS4243 enclosure is available as of the first Citi certified release of ONTAP 8.0.1 and above. For full details on the DS4243 storage administrators are encouraged to review the [DS4243 release notes](#).

System Capacity (Multipath-HA and single head [dual path] configurations)					
Platform	3210	3220	3240	3250	3270
Maximum Number of Disks - All drive types & sizes	240	480	600	720	960
Storage Raw Maximum (TB) - SAS 300GB	70	140	175	210	281
SAS 450GB	105	210	263	316	421
SAS 600GB	140	281	351	421	562
BSAS (SATA) 1TB	240	480	600	720	960
BSAS (SATA) 2TB	480	960	1200	1440	1920
BSAS (SATA) 3TB	720	1440	1800	2160	2880
Maximum DS4243 shelf count					
All drive types & sizes (a)	10	20	25	30	40
Maximum shelves per stack (b)	10	10	10	10	10
Disks per shelf (c)	24	24	24	24	24
Maximum disks per stack (b x c)	240	240	240	240	240
Implied maximum SAS stack count (a / b)	1	2	2.5	3	4
Implied maximum required SAS HBA count per controller ⁹	1	2	2	2	2
Aggregate Maximum Size (32-bit) (TB)	16	16	16	16	16
Aggregate Maximum Size (64-bit) (TB) ¹⁰	75	90	90	105	105

Table 8 - DS4243 supported RAW storage capacities

⁹ Implies that the SAS stack is shared between the storage controller pair in an HA cluster

¹⁰ Requires ONTAP 8.1 or later. Earlier versions of ONTAP 8 support smaller maximum aggregate sizes

For the PCIe expansion slot assignments for the SAS HBAs consult section [7.6](#)

7.8.2.3 DS2246 shelf

Support for the SAS based DS2246 enclosure is available as of the first Citi certified release of ONTAP 8.0.1 and above. For full details on the DS2246 storage administrators are encouraged to review the [DS2246 release notes](#).

System Capacity (Multipath-HA and single head [dual path] configurations)					
Platform	3210	3220	3240	3250	3270
Maximum Number of Disks - All drive types & sizes	240	480	600	720	960
Storage Raw Maximum (TB) - SAS 450GB	105	210	263	316	421
SAS 600GB	140	281	351	421	562
SAS 900GB	210	421	527	632	843
Maximum DS2246 shelf count					
All drive types & sizes (a)	10	20	25	30	40
Maximum shelves per stack (b)	10	10	10	10	10
Disks per shelf (c)	24	24	24	24	24
Maximum disks per stack (b x c)	240	240	240	240	240
Implied maximum SAS stack count (a / b)	1	2	2.5	3	4
Implied maximum required SAS HBA count per controller ¹¹	1	2	2	2	2
Aggregate Maximum Size (32-bit) (TB)	16	16	16	16	16
Aggregate Maximum Size (64-bit) (TB)	75	90	90	105	105

Table 9 – DS2246 supported RAW storage capacities

For the PCIe expansion slot assignments for the SAS HBAs consult section [7.6](#)

7.8.2.4 DS14 shelf

The legacy DS14 enclosures are not certified by CATE for use with the FAS3200 series.

¹¹ Implies that the SAS stack is shared between the storage controller pair in an HA cluster

7.8.3 RAID Configuration

The RAID configuration of the system will depend on use case (General File Services, VMware using NFS, Citi Grid, Archive, etc). Consult the ONTAP Solutions Guide for the relevant version of Data ONTAP to be installed on FAS3200 series for recommendations on RAID configuration. For a specific use case a reference design document may also be published on the CATE Storage and Backup Engineering SharePoint site. The content of any Design Document supersedes the contents of the general Solutions Guide.

7.8.4 Root Volume Sizing



The 3200 series, in conjunction with ONTAP 8, requires a significantly larger root volume (system volume) than its predecessors. As such depending on the model of 3200 series being considered and the capacity of the drives to be installed on the system a 1+2 dual parity RAID group may not be sufficient to store the root volume. This will impact the total useable space of the system being ordered so consideration of the following table is important¹².

All volumes must be created as Flexvols to optimize the provisioning model of NetApp NAS Filers. They must also be created as thick volumes (space guarantee set to volume). The following table lists the recommended root volume size per hardware model. The Citi recommended value is essentially double the NetApp recommendation to allow enough space for log files.

Hardware model	NetApp Recommended Root Vol Size	Citi Recommended Root Vol Size	Number of disks (inc parity)						
			300 GB	450 GB	600 GB	900 GB	1 TB	2 TB	3 TB
3210	145GB	290GB	4	3	3	3	3	3	3
3220	250GB	500GB	5	4	4	3	3	3	3
3240	195GB	390GB	4	4	3	3	3	3	3
3250	250GB	500GB	5	4	4	3	3	3	3
3270	250GB	500GB	5	4	4	3	3	3	3

Table 10 - Root volume size recommendations with ONTAP 8.1 or later

¹² Even if deploying a FAS3200 series with ONTAP 7.3.6P2, follow the sizing guidelines presented here to ensure enough storage capacity for the root aggregate to simplify future upgrades to ONTAP 8.1.x

7.8.5 Maximum number of volumes

Within the FAS3200 series there are slight differences in the maximum number of flexible volumes that can be supported by a storage system. These details are outlined in the following table.

Hardware model	Maximum number of FlexVols per storage controller	Maximum number of FlexVols per HA-Pair
3210	500 ¹³	1000
3220	500	1000
3240	500	1000
3250	500	1000
3270	500	1000

Table 11 - Maximum number of volumes

Remember the more volumes you add to a NetApp storages system, the longer the cluster failover time will take as there are more logical objects and services to transfer to the partner system.

¹³ With ONTAP 8.1 or later. Prior to ONTAP 8.1 only 200 volumes were supported per storage controller (400 in a HA-pair)

7.9 Network

In addition to the e0M and e0P ports discussed earlier in this document the FAS3200 series has two on-board network interfaces focused on the client data serving functions of the Filer.

The on-board ports consist of two 1 GbE ports (e0a and e0b). Depending on the total network bandwidth required by the client workload and the availability of 10GbE network infrastructure, the network interface configuration requirements will vary and extra NIC expansion adapters maybe required. Consult the appropriate [NetApp System Configuration Guide](#) for your 3200 series model to determine the preferred placement of additional NICs within the system.

7.9.1 Cabling

The following table is designed to assist in determining how many additional NICs are required per storage controller based upon the desired bandwidth and ONTAP interface group configuration. As the recommendation PCIe expansion slot is determined by referring to the [NetApp Systems Configuration Guide](#), any additional cards will just be referred to by generic references such as NIC1 or NIC2 in the following table. The table also shows how all the available ports will be aggregated into interfaces in order to achieve the required bandwidth and component resiliency. Complete details on cabling an optimal network configuration can be found in the ONTAP Solutions Guide.

Bandwidth	Required # NICs	Interface Group 1	Interface Group 2
1 GbE redundant (Active/Passive)	1	e0a (onboard)	NIC 1 / port a
2 GbE redundant (Active/Passive)	2	NIC1 / port a NIC 2 / port a	NIC1 / port c NIC 2 / port c
4 GbE redundant (Active/Passive)	2	NIC1 / ports a and c NIC 2 / ports a and c	NIC1 / ports b and d NIC 2 / ports b and d
10 GbE redundant (Active/Passive)	2	NIC 1 / port a	NIC 2 / port b

Table 12 – NIC requirements per storage controller

The use of 10Gb NICs is supported by NetApp and certified for use by CATE Storage Engineering where 10Gb network infrastructure is deployed. The latest 10Gb cards benefit from a TSO (TCPIP Segment Offload) capability and the TSO functionality is available even when the ports are joined together in active/passive (single-mode) interface group.

7.9.2 Port Bundling

The port bundling standard for ONTAP 8.x 7-mode continues to be the static mode Etherchannel configuration used historically within Citi. The readiness of LACP within ONTAP is currently being assessed but due to some concerning initial findings has been de-coupled from the ONTAP 8 certification project while the issues regarding LACP on NetApp are addressed.

7.10 Tape connectivity (Fiber Channel)

Much of the NetApp NAS tape backup at Citi is implemented using NDMP over IP and so the need to directly connect tape drives to a Filer has diminished as newer technologies have emerged offering efficiencies in terms of infrastructure optimization.

If direct Fiber Channel connection to tape infrastructure is required for your environment refer to section [7.6](#) for the PCIe slot assignments of the FC HBAs required for tape connectivity. In addition ensure the following principles are followed to maintain optimal availability of the entire environment:

- Where available (unused) and eligible PCI expansion slots allow, install two FC Tape HBAs in each storage controller for component redundancy and connect each HBA via a separate path to the tape infrastructure (applies to both directly attached and tape SAN implementations)
- Where a single Tape HBA per storage controller is the only option due to PCIe slot contention operational owners and the business must accept the risk that the lone HBA is a single point of failure in the infrastructure. In an attempt to compensate for this risk the recommendation is to order the quad port Tape Storage Adapter so that redundant ASICs can be used (ports a and b are on one ASIC and ports c and d are on the other ASIC). Do not order the dual port Tape Storage Adapter as this uses a single ASIC and compounds the single point of failure problem.

7.11 Virtual Servers (vFilers)

ONTAP 8.1 continues support for same number of vFilers per controller as was introduced with ONTAP 7.2. You can configure up to 64 usable vFilers on a single storage controller (up to 130 in a cluster), depending on the memory capacity of the hosting filer. The 65th vFiler unit is vFiler0 which represents the physical hosting Filer, and is created automatically when Multistore is licensed on the system. The maximum limits for number of vFilers are shown in the table below, and highlights the maximum value supported by NetApp and the value certified by CATE Storage Engineering. Starting with the FAS3200 series there is no difference between the number of vFilers per system supported by CATE and NetApp.

Filer memory	Applicable FAS3200 models	NetApp	Citi
2GB or more	FAS3210, FAS3220, FAS3240, FAS3250 and FAS3270	65	65

Table 13 - Supported vFiler quantities per storage controller

7.12 Flash Cache

Flash Cache, the read only cache expansion card, is supported with ONTAP 8.1 (7-mode). The FAS3200 series accepts a combination of the 512GB or 1TB capacity card depending on the model of the Filer. The total number of Flash Cache adapters per controller differs according to the model of the system in the series.

FAS3200 series model	Per card capacity	Max adapters per controller	Max capacity per controller (TB)
FAS3210	NA	NA	NA
FAS3220	512GB	1	1
FAS3240	512GB	1	1
FAS3250	512GB (1TB)	2 (1)	2
FAS3270	512GB (1TB)	2 (1)	2

Table 14 - Flash Cache capacity per PCI card per FAS3200 series model

The use of FlashCache cards will likely depend upon the use case of the storage system. In most cases a design document will be available on the [CATE SharePoint Site](#) containing recommendations for type, amount and configuration of FlashCache for a given workload.

7.13 Memory management for replication threads

The final factor for comparison relates to how memory is managed when a nearstore personality license is installed, and as such how many concurrent snapmirror and snapvault operations can be supported by a device.

Each value listed in the following table indicates the maximum number of concurrent replication operations allowed for a specific model, when using a single type of replication. The values are exclusive, and not cumulative. For example, if you are using the maximum number of Volume SnapMirror replication operations for a system as a source, you cannot use any more replication operations of any type for the system.

Memory management without nearstore license	3210	3220	3240	3250	3270
Snapvault streams with ONTAP 8.1	64	64	64	64	64
Source volume snapmirror streams with ONTAP 8.1	50	50	50	50	50
Destination volume snapmirror streams with ONTAP 8.1	50	50	50	50	50
Memory management with nearstore license	3210	3220	3240	3250	3270
Snapvault streams with ONTAP 8.1	120	320	320	320	320
Source volume snapmirror streams with ONTAP 8.1	50	50	50	50	50
Destination volume snapmirror streams with ONTAP 8.1	100	100	100	100	100

Table 15 - Concurrent replication process maximums

7.14 Shared chassis resources

For the FAS3210, FAS3220, FAS3240, FAS3250 and FAS3270, irrespective of whether it is configured as a single controller or in a clustered configuration, there are various aspects of the shared chassis that act independently of the Filer itself. This behavior is similar to earlier mid-range hardware platforms provided by NetApp where the power supplies and fans were integrated into the storage controller (Filer).

There are two power supply units (PSUs) in the chassis of FAS3200 series and they power both controllers and all the fans. Each controller or IOXM shares the three fans which are powered by the shared PSUs. In order to avoid the problem of which controller manages the PSUs and what happens if a rogue controller decides to turn off the PSUs, neither controller "manages" the PSUs. The PSUs manage themselves.

The PSUs have their own environmental monitoring devices and will shutdown of their own accord if something goes wrong internal to the PSU. The PSUs manage their own thermal environment (internal redundant fans) and monitor their own voltages. If anything goes wrong and the PSU is unable to deliver power properly, it will shut down like it does today in existing NetApp hardware platforms.

The system is notified of such an event through the event message system (EMS) bus and will send alerts via the usual mechanisms to notify the users if such an event occurs.

The FAS3240AE and FAS3270AE rely on separate chassis that act independently as there is no shared backplane involved.

7.15 Maintenance best practices

There are a number of important considerations in terms of overall system availability when planning the exchange of hardware FRUs (Field Replaceable Units) to swap out a faulty component.

7.15.1 Maintaining power flow in shared chassis configuration

When performing maintenance in a shared chassis configuration (FAS3210A, FAS3220A, FAS3240A or FAS3270A), remember to leave the power supplies turned on in order to provide power to the partner node. Power will be disconnected from the storage controller requiring maintenance when the controller is removed from the chassis on its rails.

7.15.2 Time limit when replacing a fan module

A faulty fan module can be removed without powering down the system. You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

If the system is in a single-enclosure HA pair or dual-enclosure HA pair configuration, only the controller module associated with that fan shuts down after the two-minute window. The partner controller module continues to serve data without interruption.

7.15.3 Console messages when replacing a power supply

Replace the power supply within two minutes of removing it from the chassis. If this two minute window is exceeded the system continues to function, but Data ONTAP will send many messages to the console about the degraded power supply until the power supply is replaced.

8 Known Issues

The following list contains known issues with the FAS3200 series platform.

Bug #	Description	Workaround
519766 Sev 1	FAS3200 series systems shipped prior to 1 st June 2012 are prone to a condition where the system will encounter a PCI Machine Check Error. These errors are encountered at a greater likelihood than expected and will cause the storage system to panic (unplanned cluster failover event) in order to maintain data integrity	Strategic fix The strategic fix involves replacing the system motherboard, the PCI bus expansion module [IOXM] (if installed), and any SAS HBAs installed with equivalent versions manufactured after 1 st June 2012. This can be achieved using the standard hardware maintenance procedures documented on the NetApp support site and failing over file sharing services to the other node in the HA-Pair in order to maintain overall system availability

Table 16 - Known issues with the FAS3200 series platform

8.1 Issues fixed in this release

The following table summarizes known issues which are resolved with the latest certified firmware versions for the FAS3200 platform.

Bug #	Description	Fixed in
617701 Sev 2	FAS/V32XX NVMEM battery current reported as high after SP firmware 1.3 upgrade. Data ONTAP treats this error as CRITICAL and it will cause the controller to halt	SP F/W 1.3P1 and above
634251 Sev 2	32xx systems NVMEM battery current reported as high after SP firmware 1.3 upgrade on Data ONTAP 8.x. Data ONTAP treats this error as CRITICAL and it will cause the controller to halt	SP F/W 1.3P1 and above
610670 Sev 3	Warn low and critical low messages from sensor "Battery 3.3V". Data ONTAP will incorrectly produce alerts are the condition	SP F/W 1.3P1 and above
476392 Sev 3	System Fans on FAS32XX systems run high speed after a SP FW update / reset. The system will produce more noise running.	SP F/W 1.3 and above
546048 Sev 2	Service Processor (SP) fails to come up after "sp reboot" command. The service processor will be unavailable on the network	SP F/W 1.3 and above
617664 Sev 3	NVMEM battery FRU data reports the previous firmware revision after applying the FC-MTO workaround. Incorrect battery firmware information is reported	SP F/W 1.3.1 and above
457467 Sev 2	Service Processor causes BIOS to fail to obtain Field Replaceable Unit (FRU) data. Any transactions issues by the BIOS will fail after boot	SP F/W 1.3.1 and above
546010 Sev 2	Chassis Fan FRU fails on 32xx systems. Causes spurious false readings to be passed to Data ONTAP and incorrect alerts to be generated	SP F/W 1.3.2 and above
702464 Sev 3	After the SP firmware is upgraded to version 1.3.2 on FAS32xx systems, all system fans stay at high speed. Causes noise	SP F/W 1.3.2P1 and above
626959 Sev 3	Service Processor (SP) firmware update occasionally can cause system shutdown in Data ONTAP 7.3.5 and later releases within the 7.3 release family. Unplanned system shutdown (rare event)	SP F/W 1.3.3 and above
653557 Sev 3	Entering the "sp setup" command at the LOADER prompt does not update IPv4 configuration settings. Loss of management port connectivity	SP F/W 1.3.3 and above
649168 Sev 3	SP to system console does not time out after inactivity Breaks CISS requirements for inactivity timeout	SP F/W 1.3.3 and above
544285 Sev 4	SP not configurable with network cable not attached Prevents pre-configuring interface before physical port is connected	SP F/W 1.3.3 and above
547618 Sev 2	Certain words in Service Processor (SP) configuration parameters can result in the SP not connecting to the network Loss of management port connectivity	SP F/W 1.3.3 and above
862792 Sev 1	SP CPU usage eventd/ trapd is high and eventd/ trapd crashes when SP SNMP is enabled	SP F/W 1.3.3P2 and above

853063	Problem installing user help files in older releases	SP F/W 1.3.3P2 and above
853081	Fix array size issue in KCS Interface	SP F/W 1.3.3P2 and above
853143	Add resource monitor task to Service Processor	SP F/W 1.3.3P2 and above
853410	Reboot SP if IPMI tasks are being continually restarted	SP F/W 1.3.3P2 and above
853418	Fix memory leak in eventd/ trapd	SP F/W 1.3.3P2 and above
822180	Service Processor can trigger watchdog reset after 500+ days of uptime The Service Processor (SP) contains a memory leak affecting the IPMI software, eventually causing the SP to become unresponsive and possibly triggering a watch dog reset of Data ONTAP after running for approximately 500 days or longer	SP F/W 1.3.3P2 and above The resource leak will still slowly occur, but the SP will gracefully restart if resources run low.
536445 Sev 2	FAS32x0 system warning message: NVMEM battery is low on power and should be replaced as soon as practical. A FAS32x0 NVMEM battery can end up in a state where it is no longer charging.	SP F/W 1.3.2P1 and above

Table 17 - Known issues fixed with latest certified firmware versions

9 Elevations

9.1 Power Distribution Unit Connectivity

In order to ensure system availability it is imperative that each PSU in the chassis is connected to a separate PDU.




Figure 2 - PDU to PSU cabling

9.2 Shared Chassis elements

The power supply units on the left and rights of the integrated chassis provide a redundant power supply to the entire chassis. The rest of the chassis is comprised of two large bays which are occupied by controllers (Bay A and B), IO expansion modules (Bay B only) or are left empty (Bay B only) depending on the configuration of the system.




Figure 3 - Anatomy of the FAS3200 chassis

9.3 Standalone controllers (un-clustered)

9.3.1 FAS3210

The following diagram shows the FAS3210 (single chassis standalone controller).
The IO Expansion Module (IOXM) is not available for the FAS3210 model.




Figure 4 – FAS3210 rear elevation

Markup	Description
	Shared chassis power supplies
	Filer head (controller) 1
	PCIe expansion slots (2)
	On-board SAS ports (2)
	On-board fiber channel ports (2)
	On-board serial console port (RJ45)
	On-board gigabit Ethernet ports (2)
	On-board TwinAx 10 GbE cluster interconnect ports (2)
	e0M management port / service processor
	Alternate Control Path for SAS shelves (out-of-band shelf management)
	Chassis bay 2 (empty)




Figure 5 – FAS3210, FAS3240E, FAS3270E front elevation

9.3.2 FAS3220E, FAS3240E, FAS3250E and FAS3270E

The following diagram shows the FAS3220, FAS3240, FAS3250 or FAS3270 (single chassis standalone controller). Both models have the same rear elevation.

The FAS3220E, FAS3240E, FAS3250E and FAS3270E are only available with the IO Expansion Module (IOXM).




Figure 6 – FAS3240E or FAS3270E rear elevation

Markup	Description
	Shared chassis power supplies
	Filer head (controller) 1
	IO Expansion module (IOXM)
	PCIe expansion slots (6)
	On-board SAS ports (2)
	On-board fiber channel ports (2)
	On-board serial console port (RJ45)
	On-board gigabit Ethernet ports (2)
	On-board TwinAx 10 GbE cluster interconnect ports (2)
	e0M management port / service processor
	Alternate Control Path for SAS shelves (out-of-band shelf management)

9.4 Clustered controllers

9.4.1 FAS3210A, FAS3220A, FAS3240A, FAS3270A

The following diagram shows the FAS3210A, FAS3220A, FAS3240A and FAS3270A (single chassis clustered controller) configurations. The IO Expansion Module (IOXM) is not available for these models.




Figure 7 – FAS3210A, FAS3220A, FAS3240A and FAS3270A rear elevation

Markup	Description
	Shared chassis power supplies
	Filer head (controller) 1
	Filer head (controller) 2
	PCIe expansion slots (4)
	On-board SAS ports (4)
	On-board fiber channel ports (4)
	On-board serial console port (RJ45) (2)
	On-board gigabit Ethernet ports (4)
	On-board TwinAx 10 GbE cluster interconnect ports (4)
	e0M management port / service processor (2)
	Alternate Control Path for SAS shelves (out-of-band shelf management) (2)

9.4.2 FAS3240AE, FAS3240AE, FAS3250AE and FAS3270AE

The following diagram shows the FAS3220AE, FAS3240AE, FAS3250AE or FAS3270AE (dual chassis clustered controller). All models have the same rear elevation.

The FAS3220AE, FAS3240AE, FAS3250AE and FAS3270AE are only available with the IO Expansion Module (IOXM).




Figure 8 – FAS3220AE, FAS3240AE, FAS3250AE or FAS3270AE rear elevation

Markup	Description
	Chassis 1
	Chassis 2
	Shared chassis power supplies (chassis 1)
	Shared chassis power supplies (chassis 2)
	Filer head (controller) 1
	Filer head (controller) 2
	PCIe expansion slots (6 per head / 12 in the cluster)
	On-board SAS ports (2 per head / 4 in the cluster)
	On-board fiber channel ports (2 per head / 4 in the cluster)
	On-board serial console port (RJ45)
	On-board gigabit Ethernet ports (2 per head / 4 in the cluster)
	On-board TwinAx 10 GbE cluster interconnect ports (2 per head / 4 in the cluster)
	eOM management port / service processor
	Alternate Control Path for SAS shelves (out-of-band shelf management)




Figure 9 – FAS3200 series front elevation (bezel)

10 References

This section details reference documentation that will be used during this project

Document	Location / Description
FAS3200 Hardware Documentation Library*	https://support.netapp.com/documentation/productlibrary/index.html?productID=30425
System Configuration Guides*	https://hww.netapp.com/Controller/Index?platformTypeId=2032
Flash Cache/ Flash Cache 2*	http://www.netapp.com/us/products/storage-systems/flash-cache/flash-cache-tech-specs-more.aspx http://www.netapp.com/us/products/storage-systems/flash-cache/flash-cache-tech-specs.aspx#Platforms
Data Protection Online Backup and Recovery Guide*	https://library.netapp.com/ecm/ecm_download_file/ECMP1141750 https://library.netapp.com/ecm/ecm_download_file/ECMM1278416 https://library.netapp.com/ecm/ecm_download_file/ECMM1277793

* Indicates a login to the vendor's support site (<http://support.netapp.com>) is required. Due to recent issues with the support site you may have to login first and then copy and paste the above URLs into your browser.

11 Appendices

11.1 Service Image (BIOS) upgrade procedure

For reference purposes the FAS3200 BIOS update procedure is shown below.

The example demonstrates an upgrade from 5.1.0 to 5.1.1.

The complete procedure is available on NetApp's support site at the following location:

http://support.netapp.com/NOW/download/tools/serviceimage/30801900_7mode/

1. Confirm the current BIOS version on the boot media and currently loaded on the system:

a. For the boot media:

```
Filer> version -b
/cfcard/x86_64/freebsd/image1/kernel: OS 8.0.1P5
/cfcard/x86_64/freebsd/image2/kernel: OS 8.0.1P5
/cfcard/x86_64/diag/diag.krn: 5.4.7
/cfcard/x86_64/firmware/excelsio/firmware.img: Firmware 1.9.0
/cfcard/x86_64/firmware/DrWho/firmware.img: Firmware 2.5.0
/cfcard/x86_64/firmware/SB_XV/firmware.img: Firmware 4.4.0
/cfcard/x86_64/firmware/SB_XVI/firmware.img: Firmware 5.1.0
/cfcard/x86_64/firmware/SB_XVIII/firmware.img: Firmware 7.0.1
/cfcard/common/firmware/zdi/zdi_fw.zpk: PAM II Firmware 1.10 (Build
0x201012200653)
/cfcard/common/firmware/zdi/zdi_fw.zpk: X1936A FPGA Configuration
PROM 1.0 (Build 0x200706131558)
```

b. For the currently loaded BIOS version:

```
Filer> sysconfig -a
NetApp Release 8.0.1P5 7-Mode: Thu May 19 07:59:27 PDT 2011
System ID: 1573798315 (11c2-fas3240-01); partner ID:
1573798381 (11c2-fas3240-02)
System Serial Number: 700000458174 (11c2-fas3240-01)
System Rev: F0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Full Connectivity
slot 0: System Board 2.3 GHz (System Board XVI F0)
Model Name: FAS3240
Part Number: 111-00693
Revision: F0
Serial Number: 5000564915
BIOS version: 5.1.0
Loader version: 3.0
Processors: 4
Processor ID: 0x1067a
Microcode Version: 0xa0b
Processor type: Intel(R) Xeon(R) CPU
L5410 @ 2.33GHz
Memory Size: 8192 MB
Memory Attributes: Bank Interleaving
Hoisting
Chipkill ECC
NVMEM Size: 1024 MB of Main Memory Used
CMOS RAM Status: OK
```

2. Download the latest BIOS service image from your distribution server and confirm install:

a. Download

```
Filer> software install
http://10.40.42.115/netapp/firmware/service_image/fas3240/v5.1.1/30801900.zip
software: You can cancel this operation by hitting Ctrl-C in the next 6 seconds.
software: Depending on system load, it may take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to 30801900.zip
software: 100% file read from location.
software: /etc/software/30801900.zip has been copied.
software: installing software, this could take a few minutes...
software: installation of 30801900.zip completed.
Mon Nov 21 12:48:58 EST [11c2-fas3240-01:
cmds.software.installDone:info]: Software: Installation of
30801900.zip was completed.
```

b. Confirm image file is present in /etc/software

```
Filer> software list
801P5_q_image.tgz
30801764.zip
801P4_q_image.tgz
SP_FW.zip
30801900.zip
```

3. Set the console to advanced privilege

```
Filer> priv set advanced
```

4. Install the new BIOS service image on the boot media and confirm installation:

a. Install the service image

```
Filer*> download -d
download: Reminder: upgrade both the nodes in the Cluster
download: Reminder: You may need to upgrade Volume SnapMirror
destination
download: filers associated with this filer. Volume SnapMirror can
not mirror
download: if the version of ONTAP on the source filer is newer than
that on
download: the destination filer.
download: You can cancel this operation by hitting Ctrl-C in the next
6 seconds.
download: Depending on system load, it may take many minutes
download: to complete this operation. Until it finishes, you will
download: not be able to use the console.
Mon Nov 21 12:51:04 EST [11c2-fas3240-Mon Nov01:
download.request:notice]: Operator requested download initiated
21 12:51:04 EST [11c2-fas3240-01: download.request:notice]: Operator
requested download initiated
download: Downloading boot device (Service Area)
...
.
.
Filer*> Mon NovMon Nov 21 12:51:05 EST [11c2-fas3240-01:
download.requestDone:notice]: Operator requested download completed
21 12:51:05 EST [11c2-fas3240-01: download.requestDone:notice]:
Operator requested download completed
```

- b. Confirm the service image is installed on the boot media:

```

File> version -b
/cfcard/x86_64/freebsd/image1/kernel: OS 8.0.1P5
/cfcard/x86_64/freebsd/image2/kernel: OS 8.0.1P5
/cfcard/x86_64/diag/diag.krn: 5.4.7
/cfcard/x86_64/firmware/excelsio/firmware.img: Firmware 1.9.0
/cfcard/x86_64/firmware/DrWho/firmware.img: Firmware 2.5.0
/cfcard/x86_64/firmware/SB_XV/firmware.img: Firmware 4.4.0
/cfcard/x86_64/firmware/SB_XVI/firmware.img: Firmware 5.1.1
/cfcard/x86_64/firmware/SB_XVIII/firmware.img: Firmware 7.0.1
/cfcard/common/firmware/zdi/zdi_fw.zpk: PAM II Firmware 1.10 (Build
0x201012200653)
/cfcard/common/firmware/zdi/zdi_fw.zpk: X1936A FPGA Configuration
PROM 1.0 (Build 0x200706131558)

```

5. Perform a cluster failover and giveback to load the new service image into memory:

- a. Perform cluster failover from the partner node

```
Filer> cf takeover
```

- b. Observe the firmware update during the reboot

System rebooting...

Phoenix TrustedCore(tm) Server
Copyright 1985-2006 Phoenix Technologies Ltd.
All Rights Reserved
BIOS version: 5.1.0
Portions Copyright (c) 2007-2009 NetApp, Inc. All Rights Reserved

```
CPU = 1 Processors Detected, Cores per Processor = 4
Intel(R) Xeon(R) CPU           L5410  @ 2.33GHz
Testing RAM
512MB RAM tested
8192MB RAM installed
6144 KB L2 Cache
System BIOS shadowed
USB 2.0: MICRON eUSB DISK
BIOS is scanning PCI Option ROMs, this will take few seconds...
.....
```

Boot Loader version 3.0
Copyright (C) 2000-2003 Broadcom Corporation.
Portions Copyright (C) 2002-2008 NetApp

CPU Type: Intel(R) Xeon(R) CPU L5410 @ 2.33GHz

The boot media contains a newer firmware image (5.1.1).

Warning: Automatic firmware update will start in 5 seconds.
To abort, press Ctrl-C during this period.

!!!DO NOT POWER OFF SYSTEM!!!
Firmware is being updated.

```
New BIOS Version: 5.1.1
New Loader Version: 3.2
Saving Primary Image to Backup flash device
Programming -----+.+-----
----- done. 2097152 bytes written
Updating Primary Boot Flash
Programming .+.+.+.+.+.+.+.+.+.+-----+.+ done. 1310720
bytes written
```

Firmware update complete. System rebooting...

```
Phoenix TrustedCore(tm) Server
Copyright 1985-2006 Phoenix Technologies Ltd.
All Rights Reserved
BIOS version: 5.1.1
Portions Copyright (c) 2007-2011 NetApp, Inc. All Rights Reserved
```

```
CPU = 1 Processors Detected, Cores per Processor = 4
Intel(R) Xeon(R) CPU          L5410  @ 2.33GHz
Testing RAM
512MB RAM tested
8192MB RAM installed
6144 KB L2 Cache
System BIOS shadowed
USB 2.0: MICRON eUSB DISK
BIOS is scanning PCI Option ROMs, this will take few seconds...
.....
```

Boot Loader version 3.2

```
Copyright (C) 2000-2003 Broadcom Corporation.
Portions Copyright (C) 2002-2008 NetApp
```

```
CPU Type: Intel(R) Xeon(R) CPU          L5410  @ 2.33GHz
```

c. Perform cluster giveback from the partner node

```
Filer> cf giveback
```

6. Once the system is normalized confirm the new BIOS version is on the boot media and is now loaded on the system:

a. For the boot media:

```
Filer> version -b
/cfcard/x86_64/freebsd/image1/kernel: OS 8.0.1P5
/cfcard/x86_64/freebsd/image2/kernel: OS 8.0.1P5
/cfcard/x86_64/diag/diag.krn: 5.4.7
/cfcard/x86_64/firmware/excelsio/firmware.img: Firmware 1.9.0
/cfcard/x86_64/firmware/DrWho/firmware.img: Firmware 2.5.0
/cfcard/x86_64/firmware/SB_XV/firmware.img: Firmware 4.4.0
/cfcard/x86_64/firmware/SB_XVI/firmware.img: Firmware 5.1.1
/cfcard/x86_64/firmware/SB_XVIII/firmware.img: Firmware 7.0.1
/cfcard/common/firmware/zdi/zdi_fw.zpk: PAM II Firmware 1.10 (Build
0x201012200653)
/cfcard/common/firmware/zdi/zdi_fw.zpk: X1936A FPGA Configuration
PROM 1.0 (Build 0x200706131558)
```

b. For the currently loaded BIOS version:

```
Filer> sysconfig -a
NetApp Release 8.0.1P5 7-Mode: Thu May 19 07:59:27 PDT 2011
System ID: 1573798315 (11c2-fas3240-01); partner ID:
1573798381 (11c2-fas3240-02)
System Serial Number: 700000458174 (11c2-fas3240-01)
System Rev: F0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Full Connectivity
slot 0: System Board 2.3 GHz (System Board XVI F0)
    Model Name: FAS3240
    Part Number: 111-00693
    Revision: F0
    Serial Number: 5000564915
    BIOS version: 5.1.1
    Loader version: 3.2
    Processors: 4
    Processor ID: 0x1067a
    Microcode Version: 0xa0b
    Processor type: Intel(R) Xeon(R) CPU
L5410 @ 2.33GHz
    Memory Size: 8192 MB
    Memory Attributes: Bank Interleaving
    Hoisting
    Chipkill ECC
    NVMEM Size: 1024 MB of Main Memory Used
    CMOS RAM Status: OK
```

11.2 Service Processor Firmware upgrade procedure

For reference purposes the FAS3200 service processor firmware update procedure is shown below. The example demonstrates an upgrade from 1.3.3 to 1.3.3P2.

The complete procedure is available on NetApp's support site at the following location:

http://support.netapp.com/NOW/download/tools/serviceimage/30602240_ontap/

1. Download the firmware archive (zip) file from the above location
2. Rename the firmware archive SP_FW.zip
3. Save the SP_FW.zip file to a web server to be used as a distribution server for this upgrade
4. SSH to the FAS3200 series system and confirm the SP firmware version currently running. In this example note that version 1.3.3 is installed and is therefore a candidate to upgrade to version 1.3.3P2. If the SP has an existing network configuration, make a note of the details to confirm the IP address, netmask and gateway values are correct post upgrade.

```
11c2-fas3240-01> sp status
      Service Processor           Status: Online
      Firmware Version:         1.3.3
      Mgmt MAC Address:         00:A0:98:14:15:1E
      Ethernet Link:            up, 100Mb, half duplex, auto-neg
complete
      Using DHCP:                no
      IPv4 configuration:
      IP Address:                10.40.53.76
      Netmask:                   255.255.255.192
      Gateway:                   10.40.53.65
      IPv6 configuration:        Disabled
```

5. Check the status of the FCMT0 flag by following the procedure in section [11.6](#).
6. Confirm if a pre-existing service processor firmware image is in the local software repository (/etc/software) and delete it if one exists (this is preferred to using the software install -f approach). Confirm the firmware image has been removed:

```
11c2-fas3240-01> software list
814P4_q_image.tgz
SP_FW.zip
814P5_q_image.tgz
814P6_q_image.tgz
814P1_q_image.tgz
11c2-fas3240-01> software delete SP_FW.zip

11c2-fas3240-01> software list
ontap_zoneinfo.zip
814P4_q_image.tgz
814P5_q_image.tgz
814P6_q_image.tgz
814P1_q_image.tgz
```

7. Download the latest firmware image from your distribution server used in step 3. In this example a lab distribution server 10.40.42.115 is used. Note the command is all one line but is word wrapped here for presentation purposes.

```
11c2-fas3240-01> software install
http://10.40.42.115/repo/netapp/firmware/sp/3200/v1.3.3P2/SP_FW.zip
software: You can cancel this operation by hitting Ctrl-C in the next 6
seconds.
software: Depending on system load, it may take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to SP_FW.zip
software: 100% file read from location.
software: /etc/software/SP_FW.zip has been copied.
software: installing software, this could take a few minutes...
software: installation of SP_FW.zip completed.
Mon May  4 05:11:37 EDT [11c2-fas3240-01:cmds.software.installDone:info]:
Software: Installation of SP_FW.zip was completed.
```

8. Initiate the firmware update process with the 'sp update' command.
Before performing this step ensure your SSH session is established through an interface other than the Service Processor. In tests this stage took between 25 and 30 minutes.

```
11c2-fas3240-01> sp update
Updating the SP firmware.
Please ensure you have run 'software install' command with correct
parameters for 'sp update' prior to this step.
DO NOT reset this system during this process.
Mon May  4 05:12:52 EDT [11c2-fas3240-01:sp.orftp.failed:warning]: SP
communication error, receiver unable to create file.
New SP version : 1.3.3P2
Current SP version : 1.3.3
Sending files to SP
.....
.....
.....
.....
Installing package on SP
.....
SP: Firmware updated successfully!
You need to reboot the SP for the new image to take effect
NOTE : If your console connection is through the SP,it will be disconnected
```

9. Confirm the SP reboot request by entering 'Y'. Note rebooting the SP does not impact the FAS3200 series systems ability to serve data. An SP reboot takes about 5 minutes on average. The messages displayed while the SP is being updated (shown below) are normal and expected.

```
Do you want to reboot the SP now? (y/n) :? y
Mon May  4 05:25:47 EDT [acp-vfiler@11c2-fas3240-01:acp.locked.wrench.port.up:info]: The on-board locked wrench port is up.
Mon May  4 05:25:47 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.up:info]: IOXM switch port 1 is up.
Mon May  4 05:25:47 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.up:info]: IOXM switch port 2 is up.
Mon May  4 05:25:47 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.up:info]: IOXM switch port 3 is up.
Mon May  4 05:25:47 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.up:info]: IOXM switch port 4 is up.
```

Console connection through the SP will now be dropped

```
11c2-fas3240-01> Mon May  4 05:26:46 EDT [11c2-fas3240-01:netif.linkDown:info]: Ethernet e0M: Link down, check cable.
Mon May  4 05:26:46 EDT [11c2-fas3240-01:netif.linkDown:info]: Ethernet e0P: Link down, check cable.
Mon May  4 05:26:47 EDT [11c2-fas3240-01:netif.linkUp:info]: Ethernet e0M: Link up.
Mon May  4 05:26:47 EDT [11c2-fas3240-01:netif.linkUp:info]: Ethernet e0P: Link up.
11c2-fas3240-01> Mon May  4 05:28:07 EDT [acp-vfiler@11c2-fas3240-01:acp.locked.wrench.port.up:info]: The on-board locked wrench port is up.
Mon May  4 05:28:07 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.down:info]: IOXM switch port 1 is down.
Mon May  4 05:28:07 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.down:info]: IOXM switch port 2 is down.
Mon May  4 05:28:07 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.down:info]: IOXM switch port 3 is down.
Mon May  4 05:28:07 EDT [acp-vfiler@11c2-fas3240-01:acp.ioxm.port.down:info]: IOXM switch port 4 is down.
```

10. Once the SP has completed its reboot confirm the new firmware version and check that the network settings are consistent with those collected in step 4. Use the 'SP status' command to complete this task.

```
11c2-fas3240-01> sp status
      Service Processor           Status: Online
      Firmware Version:         1.3.3P2
      Mgmt MAC Address:         00:A0:98:14:15:1E
      Ethernet Link:            up, 100Mb, half duplex, auto-neg
complete
      Using DHCP:                no
      IPv4 configuration:
      IP Address:                10.40.53.76
      Netmask:                  255.255.255.192
      Gateway:                  10.40.53.65
      IPv6 configuration:        Disabled
```

11. Recheck the status of the FCMT0 flag by following the procedure in section [11.6](#).

11.3 Service Processor Firmware downgrade procedure

The downgrade procedure has been removed from this document as of Service Processor firmware release 1.3.2P1 as a downgrade of SP firmware should no longer be required as part of normal operational processes.

11.4 Procedure to validate NV-MEM battery firmware version prevents bug 536445

This procedure is no longer required for the FAS3200 series storage systems when Service Processor firmware version 1.3.2P1 or later is installed on the system.

11.5 NV-MEM battery firmware upgrade procedure (bug 536445)

This procedure is no longer required for the FAS3200 series storage systems when Service Processor firmware version 1.3.2P1 or later is installed on the system. This is because the battery firmware is pre-bundled as part of the SP firmware package.

11.6 Procedure to validate the status of the Service Processor FCMT0 flag

The FCMT0 flag state is related to bug [536445](#) but should also be checked as part of the pre and post Service Processor firmware upgrade tasks. This procedure is documented as part of [knowledge base article 2016592](#).

1. Enter the SP shell from the console

```
11c2-fas3240-01> CTRL+G
```

2. Login to the Service Processor and run the following command (the output below is line wrapped):

```
SP 11c2-fas3240-01> system sensors
```

Sensor Name	Current	Unit	Status	LCR	LNC
UNC	UCR				
Bat_1.8V	1.806	Volts	ok	1.613	1.625
1.974	2.000				
Bat_8.0V	8.100	Volts	ok	na	na
8.900	9.000				
Bat_Curr	0.000	Amps	ok	na	na
2.200	2.300				
Bat_Capacity	4.096	Amps * hour	ok	na	na
na	na				
Bat_Run_Time	114.000	hour	ok	72.000	80.000
na	na				
Bat_Temp	27.000	degrees C	ok	-40.000	1.000
70.000	75.000				
Charger_Curr	0.000	Amps	ok	na	na
2.200	2.300				
Charger_Cycles	0.000	cycles	ok	na	na
250.000	251.000				
Charger_Volt	8.200	Volts	ok	na	na
8.900	9.000				

3. Check if “Bat_Run_Time” is below 80-hr Warning or below 72-hr Critical.
4. Confirm that “Bat_Curr” = 0 mA and “Charger_Volt” = 0 mV.

For systems below 72-hr Critical Low threshold – LOW_BATT, take immediate action. The shutdown sequence has begun. Follow the procedure in section [11.7](#).

For systems below 80-hr but > 72-hr “run time,” there is some time prior to initiating the shutdown sequence. Under normal conditions, a typical battery can take up to 14 days to decrease from 80 to 72-hr “run time”. Follow the procedure in section [11.7](#).

If the system sensors output does not meet any of the criteria list above proceed to step 5 (below).

5. Press CTRL-D to exit the SP shell (if SSH was used to access the SP, connect to the system console port).

11.7 Procedure to reset FCMT0 flag

The FCMT0 flag reset requirement is related to bug [536445](#).

Before attempting this procedure you must have completed the procedure to verify whether the FCMT0 flag issue is present on the storage system. As documented in section [11.6](#).

You are only required to follow the process documented below if the system sensors output on the service processor meets the criteria defined in step 4 of section [11.6](#).

**This procedure requires a storage failover (cf takeover) and an RFC (Request for Change) / green zone maintenance window is required.
Console connectivity is required.**

Also note that if the Bat_Run_Time value of the system sensors command is less than 72 hr “runtime”, the system will not boot Data ONTAP at step 4 of the following process until the battery has charged to a minimum of 72 hr “runtime” is reached. This is usually achieved in a few minutes. If the system is at less than 72hr “runtime” be prepared for the system to be in a degraded (failed over state) for an extended period of time (possibly as long as 4 hours worse case).

1. Assuming an HA configuration, run the following command to initiate a cf takeover.

- a. From the partner controller ensure autogiveback is not enabled (that is, ON).

```
11c2-fas3240-02> options cf.giveback.auto.enable
```

If enabled:

```
11c2-fas3240-02> options cf.giveback.auto.enable off
```

- b. From the partner controller issue the following command.

```
11c2-fas3240-02> cf takeover
```

If the AUTOBOOT environment variable was set to TRUE, the local node boots to the Waiting for giveback...(Press Ctrl-C to abort wait) prompt.

Press **CTRL+c** to abort waiting and halt the node. Enter **y** when the system prompts, Do you wish to halt this node rather than wait [y/n]?

2. Enter Maintenance Mode. While booting Data ONTAP, when the CTL-C message appears, press **<CTL+C>** to interrupt the boot process.

```
LOADER> boot_ontap
```

Select 5 as the Maintenance Mode boot method.

3. Clear the “no charger” condition by running the following command:

```
*> halt
```

4. Recover the system by running the following command:

```
LOADER> boot_ontap
```

For systems with less than 72 hr “runtime”, the system will not boot Data ONTAP until the battery has charged to a minimum of 72 hr “runtime” is reached. This is usually achieved in a few minutes.

5. If a takeover was performed per step 1, the controller will boot to waiting for giveback.

From the partner controller, run the following command:

```
11c2-fas3240-02> cf giveback
```

6. Enter the SP shell from the console

```
11c2-fas3240-01> CTRL+G
```

7. Confirm the battery status by running the system sensors command.
8. Verify that the “Charger_Volt” is 8200 mV and there is “Bat_Curr” >0 mA (if the battery needs to be charged). If the battery is fully charged, Bat_Curr will be 0mA and “Charger_Volt” will be 8.200 V.

Recheck “Bat_Run_Time”: This value should be increasing. For systems that do not reach the 72 hr “runtime” threshold, the battery should recover and the warning messages will then discontinue.

9. Press **CTRL-D** to exit the SP shell (if SSH was used to access the SP, connect to the system console port).

Restore the auto giveback option to the original state:

```
11c2-fas3240-02> options cf.giveback.auto.enable <on | off>
```

11.8 Specification comparison between 3210, 3220, 3240, 3250 and 3270 models (for ONTAP prior to 8.1.4)

Feature	3210	3220	3240	3250	3270
CPU					
Vendor	Intel		Intel		Intel
Codename	Wolfedale		Harpertown		Wolfedale
Branding	Xeon		Xeon		Xeon
Model	E5220		L5410		E5240
Cores	2		4		2
Clock rate (core)	2.33 GHz		2.33 GHz		3.0 GHz
Cache	6MB (L2)		2 x 6MB (L2)		6MB (L2)
Total CPU count ¹⁴	2		2		4
Total core count ¹⁵	4		8		8
Controller architecture	64-bit		64-bit		64-bit
RAM	5GB		8GB		20GB
NV-MEM	512MB		1GB		2GB
Number Native FC ports (speed)	2 (4GB SFP)		2 (4GB SFP)		2 (4GB SFP)
Number Native 1GB Ethernet ports	2		2		2
Number Native 10GB Ethernet ports	0		0		0
Expansion slots	2 PCIe		2 PCIe (without IOXM) 6 PCIe (with IOXM)		2 PCIe (without IOXM) 6 PCIe (with IOXM)
FlashCache (maximum)					1TB
Maximum raw disk capacity (TB) (HA) (ONTAP 7.3/ ONTAP 8.0.2 (7-mode))	240/720		600/ 1800		960/ 2880
Maximum number of SAS/BSAS (SATA) drives (HA)	240		600		960
Maximum raw 64-bit aggregate size (TB) ONTAP 8.0.2 (7-mode) only)	50		50		70

¹⁴ Each figure is a per Cluster amount. These figures should be halved for a standalone configuration.

¹⁵ Each figure is a per Cluster amount. These figures should be halved for a standalone configuration.

11.9 Root Volume Recommendation for ONTAP 7.3.6/ ONTAP 8.0.2

Hardware model	NetApp Recommended Root Vol Size (ONTAP 7.3.6/ ONTAP 8.0.2 (7-mode))	Citi Recommended Root Vol Size (ONTAP 7.3.6/ ONTAP 8.0.2 (7-mode))	Number of disks (inc parity) (ONTAP 7.3.6/ ONTAP 8.0.2 (7-mode))					
			300 GB	450 GB	600 GB	900 GB	1 TB	2 TB
3210	17GB/ 151GB	34GB/ 302GB	3/ 4	3/ 3	3/ 3	3/ 3	3/ 3	3/ 3
3220	NA	NA	-	-	-	-	-	-
3240	22GB/ 205GB	44GB/ 410GB	3/ 4	3/ 3	3/ 3	3/ 3	3/ 3	3/ 3
3250	NA	NA	-	-	-	-	-	-
3270	44GB/ 250GB	88GB/ 500GB	3/ 4	3/ 4	3/ 3	3/ 3	3/ 3	3/ 3

11.10 Maximum Volume (Single) for ONTAP 7.3.6/ ONTAP 8.0.2

Hardware model	Maximum number of FlexVols per storage controller	Maximum number of FlexVols per HA-Pair
3210	200	400
3220	NA	NA
3240	500	1000
3250	NA	NA
3270	500	1000

11.11 Concurrent replication process maximums for ONTAP 7.3.6/ ONTAP 8.0.2

Memory management without nearstore license	3210	3220	3240	3250	3270
Snapvault streams with ONTAP 7.3.6/ ONTAP 8.0.2	64	NA	64	NA	64
Source volume snapmirror streams with ONTAP 7.3.6/ ONTAP 8.0.2	50	NA	50	NA	50
Destination volume snapmirror streams with ONTAP 7.3.6/ ONTAP 8.0.2	50	NA	50	NA	50
Memory management with nearstore license	3210	3220	3240	3250	3270
Snapvault streams with ONTAP 7.3.6/ ONTAP 8.0.2	120	NA	320	NA	320
Source volume snapmirror streams with ONTAP 7.3.6/ ONTAP 8.0.2	50	NA	50	NA	50
Destination volume snapmirror streams with ONTAP 7.3.6/ ONTAP 8.0.2	100	NA	100	NA	100

12 Document Version Control

This section contains a detailed list of the updates made during each document release. Section 1 is purely an overview.

Version	Changes
0.1	Initial draft – general peer review
0.2	- Corrected minor typos - Added clarification regarding the 2.5 SAS stack maximum size with the FAS3240
1.0	- Removed draft document designation - Added total core count row to Table 3 - Transposed core count and cpu count numbers on FAS3270 in Table 3 to correct typo - Added names of CTI regional reviewers to Table 1 - Added section 7.8.5 detailing maximum number of volumes supported per Filer model
1.1	- Updated minimum ONTAP release to 8.0.1P5 - Updated Citi minimum ONTAP release to 8.0.1P5 in Table 2
1.2	Updated minimum recommended Service Processor firmware from 1.2 to 1.2.2 in Table 6
1.3	- Updated minimum recommended System Board firmware/ Service Image Bios from 5.1.0 to 5.1.1 in Table 6 - Added section 11 detailing the Service Image upgrade process
1.4	Changed the recommended service processor level from latest recommend to version 1.2.2 as the latest current version (1.2.3) is only supported by NetApp with a minimum of ONTAP 8.0.2P3 or later
2.0	- Updated Table 2 to include Citi recommended ONTAP 7.x release and Citi recommended ONTAP 8.x / 7-mode release - Reworded section Error! Reference source not found. regarding ONTAP 7.x support of the FAS3200 series in line with the hardware support statement enclosed in the ONTAP 7.3.6 CATE Release Notes - Updated service processor firmware version support in Table 6 - Updated reference location for System Configuration Guides in section 7.6 - Updated concurrent replication process counts in Table 15
2.1	Reworded section 6.1
2.2	- Removed draft document designation - Changed the recommended port for NIC2 in 10 GbE configurations to port B in Table 12
2.3	Added bug 536445 to section 8
2.4	- Updated service processor firmware version support in Table 6 for minimum and recommended levels to version 1.3 - Updated bug 536445 in section 8 now strategic fix is available - Added section 11.2 to describe SP firmware update procedure and section 11.3 to describe steps involving in upgrading NV-MEM battery firmware to correct for bug 536445
2.5	- Reverted service processor firmware version support in Table 6 for minimum and recommended levels to version 1.2.3 due to bugs 617701 and 634521 - Change the SP firmware upgrade example in section 11.2 of 1.3 upgrade from 1.2.3, to 1.2.3 upgrade from 1.2.2 to avoid confusion and prevent accidental deployment of SP FW version 1.3 - Added new section 11.3 detailing SP firmware downgrade procedure as corrective action to mitigate risk of bugs 617701 and 634521 with firmware version 1.3 - Added new section 11.4 detailing steps to ensure NV-MEM battery firmware is at correct level to prevent bug 536445 - Added sections 11.6 and 11.7 relating to checking and resetting the FCMTO flag - Added bugs 617701 and 634521 to section 8 - Updated URL for system configuration guides to point to the new interactive Hardware Universe in section 10.

Version	Changes
2.6	- Added Sev 1 bug 519766 to section 8 - Changed hyperlinks from now.netapp.com to support.netapp.com
2.7	- Updated section 11.2 with SP F/W v1.2.3 to v1.3.2 steps - Updated expected results for test 2.3.2 in section 13.2 - Added test 2.5 to section 13.2 - Updated SP firmware recommended version in Table 6
2.7.1	- Removed procedure from section 11.4 as it's no longer required post SP F/W 1.3.2 - Removed procedure from section 11.5 as it's no longer required post SP F/W 1.3.2 - Renamed title of section 11.6 to more appropriate description - Removed SP F/W downgrade process from section 11.3 as this procedure will no longer be common practice post F/W version 1.3.2 - Added pre and post SP F/W upgrade FCMT0 flag status check steps to section 11.2
2.7.2	- Updated SP F/W version to 1.3.2P1 - Added section 8.1 describing issues fixed with latest firmware releases - Updated test 2.5 to section 13.2
2.7.3	- Removed draft document designation - Removed tactical fix instructions from bugs 536445, 617701, 634521 - Updated the strategic fix for bugs 536445, 617701, 634521 to SP 1.3.2P1 or later
2.7.4	Corrected typo on number of 1-gig port s for FAS3170 in Table 5
2.8	- Added FAS3220 and FAS3250 details - Added Section 11.8 – 11.11
2.9	RFT release
3.0	Removed draft document designation and released to production
3.0.1	- Added draft documentation designation - Updated SP F/W version to 1.3.3 - Added section 8.1 describing issues fixed with latest firmware releases - Updated test 2.5 to section 13.2
3.1	- Removed draft document designation - Updated ONTAP version details from 8.1.3 to 8.1.4P1 - Updated Service Processor image file names and versions in Table 6
3.1.1	- Updated document with Service Processor update from v1.3.3 to v1.3.3P2 - Removed ONTAP7.3.X/ 8.0.X support from Section 6.1 Operating System Dependencies and updated Recommended version to ONTAP8.1.4P6 - Included Bug 822180 in Section 8 Known Issues with strategic fix after updating SP firmware to v1.3.3P2 - Updated Section 8.1 Issues fixed with this release for Bugs 862792, 853063, 853081, 853143, 853410, 853418 - Updated Section 11.2 with SP F/W v1.3.3 to v1.3.3P2 - Updated table for DS4243 shelf to include 3TB disks
3.1.2	- Updated Section 8.1 Issues fixed with Bug 822180 and 536445 - Removed the above mentioned bugs from Section 8 as they are fixed in this release - Release for Testing
3.2	Release for Production

13 Test Results

13.1 Physical Tests

This section describes the steps used in the physical testing a FAS32x0.

This test plan assumes the CPOC lab team has already installed the latest ONTAP 8.0.1 patch release on the cluster

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.0	FAS3200 series physical tests					
1.1	Alternate hardware configuration - [A] – [clustered single chassis]					
1.1.1	Base build	Configure system network interfaces	Filer has ping-able IP address and C\$ CIFS share can be mounted from client	As expected	Pass	
1.1.2	Environment	Capture output of 'environment status chassis'	Command provides output – ensure 'Partner Ctrl Pre ' is listed as present	As expected	Pass	Is there another way to determine the hardware configuration other than this command
1.1.3	System config	Capture output of 'sysconfig -a'	IOXM should not be shown in the list	As expected	Pass	
1.1.4	Operations Manager reporting	Register 32x0A in Operations Manager 4.0.1 and determine if any information is provided to determine the model and configuration (A in this case)	Filer added to Operations Manager 4.0.1	As expected	Pass	Determine if there is a mechanism in OM to distinguish between each model and its configuration

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.2	Alternate hardware configuration - [AE] – [clustered dual chassis with IOXM installed]					
1.2.1	Base build	Configure system network interfaces.	Filer has ping-able IP address and C\$ CIFS share can be mounted from client	As expected	Pass	
1.2.2	Environment	Capture output of 'environment status chassis'	Command provides output – ensure 'Partner Ctrl Pre ' is listed as not present	As expected	Pass	Is there another way to determine the hardware configuration other than this command
1.2.3	System config	Capture output of 'sysconfig –a'	IOXM should be shown in the list	As expected	Pass	
1.2.4	Operations Manager reporting	Register 32x0AE in Operations Manager 4.0.1 and determine if any information is provided to determine the model and configuration (A in this case)	Filer added to Operations Manager 4.0.1	As expected	Pass	Determine if there is a mechanism in OM to distinguish between each model and its configuration
1.3	Alternate hardware configuration comparison [A vs AE]					
1.3.1	Compare configurations	Compare output from tests 1.1 and 1.2 and determine a short list of indicators / commands that can be used to confirm the presence of each configuration	Different configurations can be confirmed with one or two simple commands	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.4	DS4243 and DS2246 disk enclosure tests					
1.4.1	Validate stacks	Ensure each shelf type is cabled in a separate SAS stack	Cabled in separate stacks	As expected	Pass	
1.4.2	Validate LANE speed	Sysconfig -v, check phy state: for each SAS port	3.0 gb/s for DS4243 6.0 gb/s for DS2246	As expected	Pass	
1.4.3	Validate disks	Sysconfig -v and sysconfig -r	Ensure all disks are reported	As expected	Pass	
1.4.4	Create aggregate on DS4243	Aggr create aggr1 -t raid_dp 16@size	Aggregate created	As expected	Pass	
1.4.5	Create volume on DS4243	Vol create vol1 -s volume aggr1 50g	Volume created	As expected	Pass	
1.4.6	Create CIFS share on DS4243	Cifs shares -add vol1\$ /vol/vol1	Share created	As expected	Pass	
1.4.7	Mount share from Windows client	Net use * \\filer_ip\vol\$ /u:filer_ipadministrator	Share mapped	As expected	Pass	
1.4.8	Create, modify and delete files	Create, modify and delete files	Normal CIFS behavior	As expected	Pass	
1.4.9	Create aggregate on DS2246	Aggr create aggr2 -t raid_dp 16@size	Aggregate created	As expected	Pass	
1.4.10	Create volume on DS2246	Vol create vol2 -s volume aggr2 50g	Volume created	As expected	Pass	
1.4.11	Create CIFS share on DS2246	Cifs shares -add vol2\$ /vol/vol2	Share created	As expected	Pass	
1.4.12	Mount share from Windows client	Net use * \\filer_ip\vol\$ /u:filer_ipadministrator	Share mapped	As expected	Pass	
1.4.13	Create, modify and delete files	Create, modify and delete files	Normal CIFS behavior	As expected	Pass	
1.4.14	Remove DS4243 SAS cable	Remove SAS cable from one port on SAS HBA for DS4243 on node A	Vol 1 still accessible	As expected	Pass	
1.4.15	Confirm degraded state	Filer> storage show disk -p	Only one active path	As expected	Pass	
1.4.16	Reinstate SAS cable	Filer> storage show disk -p	Both paths active	As expected	Pass	
1.4.17	Reboot system	Reboot -f	Filer rebooted	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.4.18	Remove other DS4243 SAS cable	Remove other SAS cable from one port on SAS HBA for DS4243 on node A	Vol 1 still accessible	As expected	Pass	
1.4.19	Confirm degraded state	Filer> storage show disk -p	Only one active path	As expected	Pass	
1.4.20	Reinstate SAS cable	Filer> storage show disk -p	Both paths active	As expected	Pass	
1.4.21	Reboot system	Reboot -f	Filer rebooted	As expected	Pass	
1.4.22	Remove DS4243 SAS cable from Node B	Remove SAS cable from one port on SAS HBA for DS4243 on node B	Vol 1 still accessible	As expected	Pass	
1.4.23	Confirm degraded state	Filer> storage show disk -p	Only one active path	As expected	Pass	
1.4.24	Reinstate SAS cable	Filer> storage show disk -p	Both paths active	As expected	Pass	
1.4.25	Reboot system	Reboot -f	Filer rebooted	As expected	Pass	
1.4.26	Remove other DS4243 SAS cable from Node B	Remove other SAS cable from one port on SAS HBA for DS4243 on node B	Vol 1 still accessible	As expected	Pass	
1.4.27	Confirm degraded state	Filer> storage show disk -p	Only one active path	As expected	Pass	
1.4.28	Reinstate SAS cable	Filer> storage show disk -p	Both paths active	As expected	Pass	
1.4.29	Reboot system	Reboot -f	Filer rebooted	As expected	Pass	
1.4.30	Remove DS2246 SAS cable	Remove SAS cable from one port on SAS HBA for DS2246 on node B	Vol 2 still accessible	As expected	Pass	
1.4.31	Confirm degraded state	Filer> storage show disk -p	Only one active path	As expected	Pass	
1.4.32	Reinstate SAS cable	Filer> storage show disk -p	Both paths active	As expected	Pass	
1.4.33	Reboot system	Reboot -f	Filer rebooted	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.4.34	Remove other DS2246 SAS cable	Remove other SAS cable from one port on SAS HBA for DS2246 on node B	Vol 2 still accessible	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.35	Confirm degraded state	Filer> storage show disk –p	Only one active path	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.36	Reinstate SAS cable	Filer> storage show disk –p	Both paths active	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.37	Reboot system	Reboot -f	Filer rebooted	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.38	Remove DS2246 SAS cable from Node A	Remove SAS cable from one port on SAS HBA for DS2246 on node A	Vol 2 still accessible	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.39	Confirm degraded state	Filer> storage show disk –p	Only one active path	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.40	Reinstate SAS cable	Filer> storage show disk –p	Both paths active	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.41	Reboot system	Reboot -f	Filer rebooted	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.4.42	Remove other DS2246 SAS cable from Node A	Remove other SAS cable from one port on SAS HBA for DS2246 on node A	Vol 2 still accessible	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.43	Confirm degraded state	Filer> storage show disk –p	Only one active path	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.44	Reinstate SAS cable	Filer> storage show disk –p	Both paths active	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.4.45	Reboot system	Reboot -f	Filer rebooted	Tests 1.4.1 – 1.4.33 are consistent with disk shelf results – skipping the remainder of these SAS tests	Pass	
1.5	Alternate Control Path (ACP)					
1.5.1	Show ACP state	Filer> storage show acp	All paths active	As expected	Pass	
1.5.2	Show ACP polling information	Filer> acpadmin list_all	All devices contactable	As expected	Pass	
1.5.3	Offline e0P on node A	Filer A> ifconfig e0p down	Port off-lined	As expected	Pass	
1.5.4	Show ACP state on node A	Filer A > storage show acp	Some connectivity degraded	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.5.5	Show ACP state on node B	Filer B > storage show acp	Some connectivity degraded	As expected	Pass	
1.5.6	Validate active data paths	Filer A> storage show disk -p	Both paths active	As expected	Pass	
1.5.7	Online e0P on node A	Filer A> ifconfig e0p up	Port online	As expected	Pass	
1.5.8	Show ACP state on node A	Filer A > storage show acp	All devices contactable	As expected	Pass	
1.5.9	Show ACP state on node B	Filer B > storage show acp	All devices contactable	As expected	Pass	
1.5.10	Offline e0P on node B	Filer B> ifconfig e0p down	Port off-lined	As expected	Pass	
1.5.11	Show ACP state on node A	Filer A > storage show acp	Some connectivity degraded	As expected	Pass	
1.5.12	Show ACP state on node B	Filer B > storage show acp	Some connectivity degraded	As expected	Pass	
1.5.13	Validate active data paths	Filer B> storage show disk -p	Both paths active	As expected	Pass	
1.5.14	Online e0P on node B	Filer A> ifconfig e0p up	Port online	As expected	Pass	
1.5.15	Show ACP state on node A	Filer A > storage show acp	All devices contactable	As expected	Pass	
1.5.16	Show ACP state on node B	Filer B > storage show acp	All devices contactable	As expected	Pass	
1.6	Flashcache (PAM II)					
1.6.1	Show FlashCache cards are installed	Filer> sysconfig -v	Card(s) listed	As expected	Pass	
1.6.2	Enable FlashCache	Filer> options flexscale.enable on	Feature enabled	As expected	Pass	
1.6.3	Generate IO	Run sio from window clients 50:50 ratio	Load running	As expected	Pass	
1.6.4	Validate caching	Filer> stats show ext_cache	Cache_objects count > 0	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.7	Power supply redundancy [Single chassis HA-Pair configuration - A]					
1.7.1	Show status	Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as OK	As expected	Pass	
1.7.2	Degraded power PSU1	Remove power cable from PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as failed	As expected	Pass	
1.7.3	Re-instate power	Re-insert power cable to PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as OK	As expected	Pass	
1.7.4	Turn off PSU1	Turn off PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 ON listed as OFF	As expected	Pass	
1.7.5	Re-instate power	Turn on PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 ON listed as ON	As expected	Pass	
1.7.6	Remove PSU1	Turn off PSU1 and then remove from chassis Filer A> environment status chassis Filer B> environment status chassis	PSU1 Present listed as Not Present	As expected	Pass	
1.7.7	Re-instate power	Re-insert PSU1 to chassis and then turn back on Filer A> environment status chassis Filer B> environment status chassis	PSU1 Present listed as Present	As expected	Pass	
1.7.8	Reset test environment	Filer A> cf disable Filer A> reboot -f Filer B> reboot -f Filer A> cf enable	Both nodes reboot successfully and cluster returned to normal operation	As expected	Pass	
1.7.10	Show status	Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as OK	As expected	Pass	
1.7.11	Degraded power PSU2	Remove power cable from PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as failed	As expected	Pass	
1.7.12	Re-instate power	Re-insert power cable to PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as OK	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.7.13	Turn off PSU2	Turn off PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 ON listed as OFF	As expected	Pass	
1.7.14	Re-instate power	Turn on PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 ON listed as ON	As expected	Pass	
1.7.15	Remove PSU2	Turn off PSU2 and then remove from chassis Filer A> environment status chassis Filer B> environment status chassis	PSU2 Present listed as Not Present	As expected	Pass	
1.7.16	Re-instate power	Re-insert PSU2 to chassis and then turn back on Filer A> environment status chassis Filer B> environment status chassis	PSU2 Present listed as Present	As expected	Pass	
1.7.17	Reset test environment	Filer A> cf disable Filer A> reboot -f Filer B> reboot -f Filer A> cf enable	Both nodes reboot successfully and cluster returned to normal operation	As expected	Pass	
1.8	Power supply redundancy [Dual chassis HA-Pair configuration - AE]					
1.8.1	Show status	Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as OK	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.2	Degraded power PSU1	Remove power cable from PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as failed	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.3	Re-instate power	Re-insert power cable to PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as OK	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.4	Turn off PSU1	Turn off PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 ON listed as OFF	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.5	Re-instate power	Turn on PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 ON listed as ON	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.6	Remove PSU1	Turn off PSU1 and then remove from chassis Filer A> environment status chassis Filer B> environment status chassis	PSU1 Present listed as Not Present	As expected	Pass	For chassis A (namely the chassis containing node A)

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.8.7	Re-instate power	Re-insert PSU1 to chassis and then turn back on Filer A> environment status chassis Filer B> environment status chassis	PSU1 Present listed as Present	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.8	Reset test environment	Filer A> cf disable Filer A> reboot -f Filer B> reboot -f Filer A> cf enable	Both nodes reboot successfully and cluster returned to normal operation	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.10	Show status	Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as OK	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.11	Degraded power PSU2	Remove power cable from PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as failed	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.12	Re-instate power	Re-insert power cable to PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as OK	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.13	Turn off PSU2	Turn off PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 ON listed as OFF	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.14	Re-instate power	Turn on PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 ON listed as ON	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.15	Remove PSU2	Turn off PSU2 and then remove from chassis Filer A> environment status chassis Filer B> environment status chassis	PSU2 Present listed as Not Present	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.16	Re-instate power	Re-insert PSU2 to chassis and then turn back on Filer A> environment status chassis Filer B> environment status chassis	PSU2 Present listed as Present	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.17	Reset test environment	Filer A> cf disable Filer A> reboot -f Filer B> reboot -f Filer A> cf enable	Both nodes reboot successfully and cluster returned to normal operation	As expected	Pass	For chassis A (namely the chassis containing node A)
1.8.18	Degraded power PSU1 and PSU2	Remove power cable from PSU1 and PSU2 Filer A> environment status chassis Filer B> environment status chassis	Cluster failover	As expected	Pass	For chassis A (namely the chassis containing node A)

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.8.20	Show status	Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as OK	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.21	Degraded power PSU1	Remove power cable from PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as failed	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.22	Re-instate power	Re-insert power cable to PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 input listed Input OK as OK	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.23	Turn off PSU1	Turn off PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 ON listed as OFF	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.24	Re-instate power	Turn on PSU1 Filer A> environment status chassis Filer B> environment status chassis	PSU1 ON listed as ON	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.25	Remove PSU1	Turn off PSU1 and then remove from chassis Filer A> environment status chassis Filer B> environment status chassis	PSU1 Present listed as Not Present	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.26	Re-instate power	Re-insert PSU1 to chassis and then turn back on Filer A> environment status chassis Filer B> environment status chassis	PSU1 Present listed as Present	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.27	Reset test environment	Filer A> cf disable Filer A> reboot -f Filer B> reboot -f Filer A> cf enable	Both nodes reboot successfully and cluster returned to normal operation	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.28	Degraded power PSU2	Remove power cable from PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as failed	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.29	Re-instate power	Re-insert power cable to PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 input listed Input OK as OK	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.30	Turn off PSU2	Turn off PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 ON listed as OFF	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.31	Re-instate power	Turn on PSU2 Filer A> environment status chassis Filer B> environment status chassis	PSU2 ON listed as ON	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.8.32	Remove PSU2	Turn off PSU2 and then remove from chassis Filer A> environment status chassis Filer B> environment status chassis	PSU2 Present listed as Not Present	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.33	Re-instate power	Re-insert PSU2 to chassis and then turn back on Filer A> environment status chassis Filer B> environment status chassis	PSU2 Present listed as Present	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.34	Reset test environment	Filer A> cf disable Filer A> reboot -f Filer B> reboot -f Filer A> cf enable	Both nodes reboot successfully and cluster returned to normal operation	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.8.35	Degraded power PSU1 and PSU2	Remove power cable from PSU1 and PSU2 Filer A> environment status chassis Filer B> environment status chassis	Cluster failover	Skipped testing as functionality has been tested in 1.8.1 - 1.8.19	Pass	For chassis B (namely the chassis containing node B)
1.9	Controller redundancy [Single chassis HA-Pair configuration - A]					
1.9.1	Cluster failover A	Filer B> cf status Filer B> cf takeover Filer B> cf status	Takeover successful Filer A vol1\$ CIFS share continues to serve data	As expected	Pass	
1.9.2	Cluster reset	Filer B> cf status Filer B> cf giveback Filer B> cf status	Failback successful	As expected	Pass	
1.9.3	Cluster failover A	Filer A> cf status Filer A> cf takeover Filer A> cf status	Takeover successful Filer B vol2\$ CIFS share continues to serve data	As expected	Pass	
1.9.4	Cluster reset	Filer A> cf status Filer A> cf giveback Filer A> cf status	Failback successful	As expected	Pass	
1.10	Controller redundancy [Dual chassis HA-Pair configuration - AE]					
1.10.1	Cluster failover A	Filer B> cf status Filer B> cf takeover Filer B> cf status	Takeover successful Filer A vol1\$ CIFS share continues to serve data	As expected	Pass	
1.10.2	Cluster reset	Filer B> cf status Filer B> cf giveback Filer B> cf status	Failback successful	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.10.3	Cluster failover A	Filer A> cf status Filer A> cf takeover Filer A> cf status	Takeover successful Filer B vol2\$ CIFS share continues to serve data	As expected	Pass	
1.10.4	Cluster reset	Filer A> cf status Filer A> cf giveback Filer A> cf status	Failback successful	As expected	Pass	
1.10.5	Cross connect cluster interconnect	Connect port c0a on Node A with c0b on Node B and c0b on Node A with c0a on Node B	ONTAP should warn a cluster cable interconnect cross-over	As expected	Pass	
1.10.6	Inline connect cluster interconnect	Connect port c0a on Node A with c0a on Node B and c0b on Node A with c0b on Node B	CF status is normal	As expected	Pass	
1.11	Network interface redundancy					
1.11.1	Single mode vif	Using 2 of the 4 network interfaces configure a single mode vif on node A	VIF created	As expected	Pass	
1.11.2	VIF favor	Favor each interface within the VIF on node A	CIFS access to vol1\$ share is maintained	As expected	Pass	
1.11.3	Unplug active port	Remove cable from active VIF interface	Connectivity fails over to redundant port, CIFS access to vol1\$ share is maintained	As expected	Pass	
1.11.4	Re-instate cable	Re-instate the cable from the previous test point	Connectivity remains on the active (favored VIF)	As expected	Pass	
1.11.5	VIF favor	Favor the interface in the VIF to the initial port	Active interface failed back	As expected	Pass	
1.11.6	Unplug redundant port	Remove cable from passive VIF interface	No change in connectivity	As expected	Pass	
1.11.7	Re-instate cable	Re-instate the cable from the previous test point	Connectivity remains on the active (favored VIF)	As expected	Pass	
1.11.9	Multi-mode vif	Using 4 of the 4 network interfaces configure a single mode VIF comprised of two multimode VIFs on node A	VIF created	As expected	Pass	
1.11.10	VIF favor	Favor each multi-mode VIF within the single mode VIF on node A	CIFS access to vol1\$ share is maintained	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
1.11.11	Unplug active port	Remove cable from one active VIF interface	VIF stays up in a degraded state, CIFS access to vol1\$ share is maintained	As expected	Pass	
1.11.12	Unplug active port	Remove second cable from one active VIF interface	Connectivity fails over to redundant port, CIFS access to vol1\$ share is maintained	As expected	Pass	
1.11.13	Re-instate cables	Re-instate the cable from the previous test points	Connectivity remains on the active (favored VIF)	As expected	Pass	
1.11.14	VIF favor	Favor the interface in the VIF to the initial port	Active interface failed back	As expected	Pass	
1.11.15	Unplug redundant port	Remove first cable from passive VIF interface	No change in connectivity	As expected	Pass	
1.11.16	Unplug redundant port	Remove second cable from passive VIF interface	No change in connectivity	As expected	Pass	
1.11.17	Re-instate cables	Re-instate the cable from the previous test points	Connectivity remains on the active (favored VIF)	As expected	Pass	

13.2 Research Tests

This section describes the testing criteria in the research phase of this project

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
2.0	FAS3200 series research tests					
2.1	Field replaceable units – hot swap or non-disruptive replacement (non-disruptive implies cluster failover required)					
2.1.1	FRU – boot device	Review doc 215-052028_A0 as to the procedure for replacing the boot device	Replacement requires a cluster failover	As expected	Pass	
2.1.2	FRU – memory DIMM	Review doc 215-05029_A0 as to the procedure for replacing a DIMM	Replacement requires a cluster failover	As expected	Pass	
2.1.3	FRU – Fan module	Review doc 215-05024_A0 as to the procedure for replacing the fan module	Fan module is a hot swappable component	As expected	Pass	
2.1.4	FRU – Power supply unit	Review doc 215-05030_A0 as to the procedure for replacing a power supply	Power supply is a hot swappable component	As expected	Pass	
2.1.5	FRU – NVMEM battery	Review doc 215-05031_A0 as to the procedure for replacing the NV MEM battery	Replacement requires a cluster failover	As expected	Pass	
2.1.6	FRU – NVMEM DIMMS	Review doc 215-05031_A0 as to the procedure for replacing the NV MEM dimms	Replacement requires a cluster failover	As expected	Pass	
2.1.7	FRU – PCIe cards	Review doc 215-05027_A0 as to the procedure for replacing PCIe cards	Replacement requires a cluster failover	As expected	Pass	
2.1.8	FRU – IO Expansion module	Review doc 215-05032_A0 as to the procedure for replacing the IOXM	Replacement requires a cluster failover	As expected	Pass	
2.1.9	FRU – FAS3200 chassis	Review doc 215-05026_A0 as to the procedure for replacing the chassis	Depends on configuration of system – shared chassis vs dedicated chassis	As expected	Pass	Chassis is passive and replacement of entire chassis is expected to be highly unlikely
2.1.10	FRU – FAS3200 storage controller	Review doc 215-05025_A0 as to the procedure for replacing the storage controller	Replacement requires a cluster failover	As expected	Pass	
2.1.11	FRU – Real Time Clock (RTC) battery	Review doc 215-05033_A0 as to the procedure for replacing the RTC battery	Replacement requires a cluster failover	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
2.2	Field replaceable units – accessibility					
2.2.1	FRU access	Review procedures from tests in section 2.1	FRU components are readily accessible	As expected	Pass	
2.2.2	FRU special tools	Review procedures from tests in section 2.1	Nothing more than a Philips screwdriver is required	As expected	Pass	
2.3	Recommended firmware					
2.3.1	System board	Review system image information on the NetApp support site	Version 5.0.1	As expected	Pass	
2.3.2	Service processor	Review system image information on the NetApp support site	Version 1.3.2	As expected	Pass	
2.4	System configuration					
2.4.1	PCIe slot assignments	Review System configuration guides on the NetApp support site	Installation order documented in system configuration guide	As expected	Pass	
2.4.2	Sizing	Review System configuration guides on the NetApp support site	System capacities documented in system configuration guide	As expected	Pass	
2.5	Service Processor firmware upgrade to version 1.3.3P2					
2.5.1	Install	Software install http://10.40.42.115/netapp/firmware/sp/3200/v1.3.3P2/SP_FW.zip	Firmware installs without error	As expected	Pass	
2.5.2	Update	Sp update	Firmware is uploaded to and installed on the SP without error	As expected	Pass	
2.5.3	Reboot	Reboot the SP post update	SP reboots with no impact to ONTAP's data serving abilities, and the network configuration of the SP remains the same prior to upgrade/ reboot	As expected	Pass	

Test #	Test Conditions	Action / Test Procedure	Expected Results / Acceptance Criteria	Actual Results / Observed Results	Pass / Fail	Notes
2.5.4	Confirmation	Confirm version 1.3.3P2 is installed with an 'SP status' command in ONTAP and a 'events newest 20' command on the Service Processor	SP status shows firmware version 1.3.3P2 installed and the event log shows 1.3.3P2 loaded ([SP.notice]: Running primary version 1.3.3P2)	As expected	Pass	
2.5.5	Stability	Review SP event log, ONTAP messages log and compare to logs prior to upgrade for a period of two week either side of the upgrade	No new error messages appear in the logs	As expected	Pass	
2.5.6	Stability	Review SP and ONTAP uptime ensuring no unexpected reboots or e0M and SP network interface issues occurred in the period up to two weeks after the firmware upgrade	No unexpected reboots occurred and no degradation in network interface availability or performance was observed	As expected	Pass	