FAS2820

SE Presentation

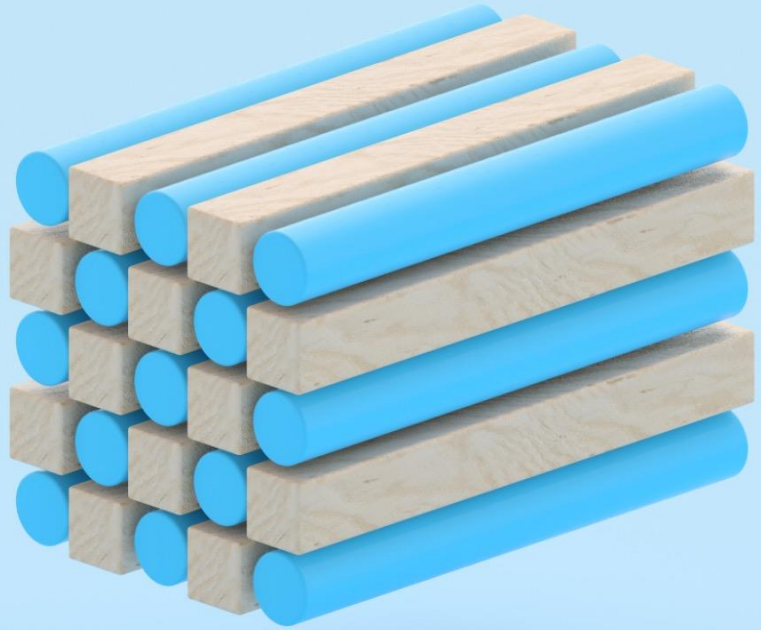

AFF and FAS PM and TME Team (ng-aff-fas-pm-tme-team)
June 2023
Version 1.1.2

Notice of NDA Content

The slides herein contain company confidential, proprietary information of NetApp and shall not be disclosed outside NetApp without a binding nondisclosure agreement in place.

These documents are for viewing purposes only and shall not be distributed.

Topics covered in order

- Architecture, positioning, and performance overview
- Specification maximums
- Technical details
  - Ports, motherboard, field replaceable parts (FRUs), and NVMEM architecture
- Detailed I/O information
  - Including supported and Best Practice configurations
- Software
- Environmental specifications
  - Power consumption, cooling, dimensions, and weight

Questions? Please email ng-aff-fas-pm-tme-team

Introducing the FAS2820

New entry FAS platform

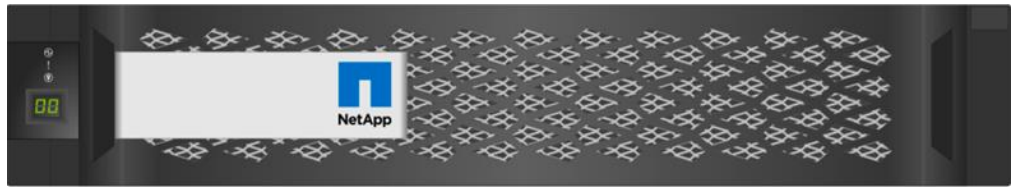

- Based on DS212C shelf
- 2U enclosure
- Two controllers in HA configuration
- 12x 3.5” form-factor drive bays

FAS2820

New entry FAS systems

- Based on the Intel IceLake-D 8-core processor
  - Provides up to 50% performance increase over the FAS2720
- Leverages existing 12Gb SAS enclosure architecture
  - Full 12Gbps SAS connectivity to internal and external drives
- Uses NVMe M.2 connection for onboard Flash Cache
- Enhanced I/O capabilities with more configuration flexibility in a 2U enclosure

Introducing the FAS2820

HA pair

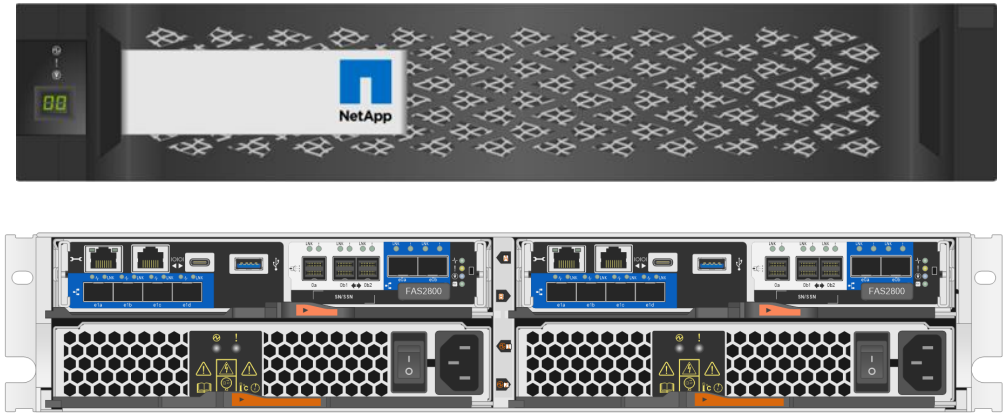

- Processor and memory (per HA pair)
  - CPU: 16 cores
  - Memory: 128GB
    - 8GB is used as NVMEM per controller
- Onboard I/O ports (per controller)
  - 2x SFP28 (25Gb Ethernet): HA and cluster interconnects
  - 2x 12Gb mini-HD SAS: External storage attachment
  - 1x Mezzanine slot for I/O expansion
  - 1x Type USB-C console port: 115,200 baud rate
  - 1x Type-A USB port: Read-only for netboot and OS updates
- Persistent write log
  - NVlog written to flash if unplanned power loss occurs

Performance

ONTAP Systems

- Refer to the NetApp Field Portal for ONTAP performance metrics across platforms
- ONTAP Hardware Performance Metrics

Systems at a Glance: FAS Series

HA pair specifications

	FAS2700	FAS2820	FAS500f	FAS8200	FAS8300	FAS8700	FAS9000
Form factor	2U[1]	2U[1]	2U[1]	3U	4U	4U	8U
CPU cores	24	16	24	32	40	64	72
Physical memory	64GB[2]	128GB[2]	128GB[2]	256GB[2]	256GB	512GB	1024GB
Max drive count	144	144	48	480	720	1,440	1,440
NVDIMM / NVRAM	8GB	8GB	16GB	16GB	32GB	32GB	64GB
Onboard / Pre-Configured Ports							
100GbE	-	-	-	-	4	4	-
25GbE	-	4	4	-	4 or 12[4]	4 or 12[4]	-
32Gb FC	-	-	-	-	-	-	-
16Gb FC	-	-	-	-	0 or 8[4]	0 or 8[4]	-
40GbE	-	-	-	-	-	-	4 or 8[5]
10GbE	4	-	-	4	-	-	-
10GBASE-T	0 or 8[3]	-	4	4	-	-	-
UTA2	0 or 8[3]	-	-	8	-	-	-
SAS	4 (12Gb)	4 (12Gb)	-	8 (12Gb)	8 (12Gb)	8 (12Gb)	8 (12Gb)
I/O Expansion Slots	-	2	4	0 or 4[3]	10	10	16
ONTAP® support	9.4RC1+	9.13.1RC1+	9.8RC1+	9.1GA+	9.7RC1+	9.7RC1+	9.1RC2+

Footnotes 1-5: See slide notes for details

Cluster node limits

ONTAP 9.13.1RC1

Homogeneous Clusters Maximum Cluster Nodes	NAS	SAN
FAS2820	24	12

Heterogeneous (Mixed) Cluster Maximum Cluster Nodes	NAS	SAN
FAS2820 AFF A900, FAS9500 AFF A250, FAS500f AFF A220, AFF C190, FAS2750, FAS2720 AFF A800, AFF A400, FAS8700, FAS8300 AFF A700, AFF A700s, FAS9000 AFF A300, FAS8200	24	12

Refer to Hardware Universe for additional information about platform mixing rules.

Aggregate and volume sizes

	FAS2820
Maximum number of volumes	1,000
Maximum aggregate size	800 TiB
Maximum volume size	100 TiB
Minimum root volume size	150 GiB

Flash limits

HA pair

	FAS2820
Minimum/maximum onboard flash	2TiB
Maximum onboard flash + Flash Pool cache	24TiB

FAS2820 controllers ship with a 1TB NVMe Flash Cache module per node.

In-band SAS management

- Starting in ONTAP 9, SAS management is now in-band over the SAS paths
  - Refer to the 12Gb SAS SE Presentation for additional information
- Supported on all platforms running ONTAP 9
  - Requires IOM12 shelf modules
- Separate ACP cabling not needed with in-band management
  - Can be left in place on upgrades to ONTAP 9
- FAS2820 controllers do not have the locked wrench port for private out-of-band ACP cables

Front view: FAS2820

Bezel Removed

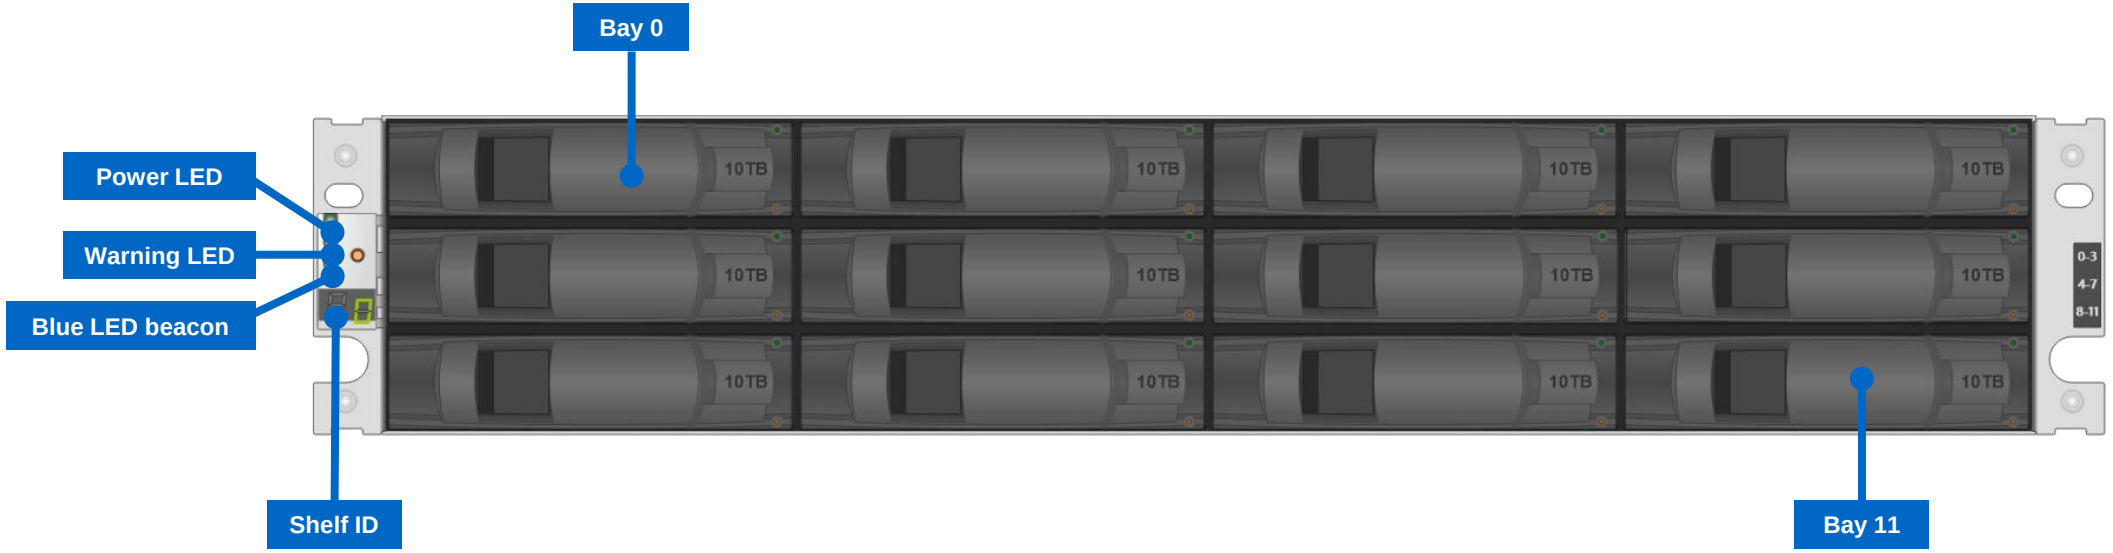

Rear view: FAS2820

Default configuration with 25Gb Ethernet mezzanine

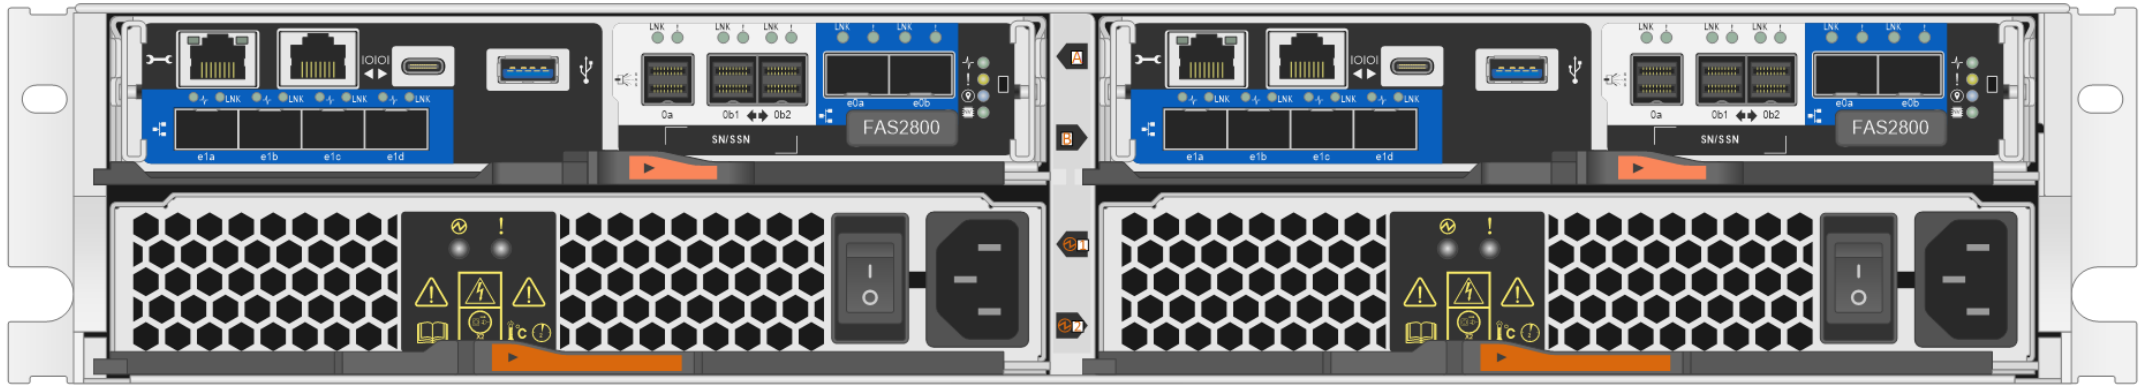

Controller I/O: FAS2820

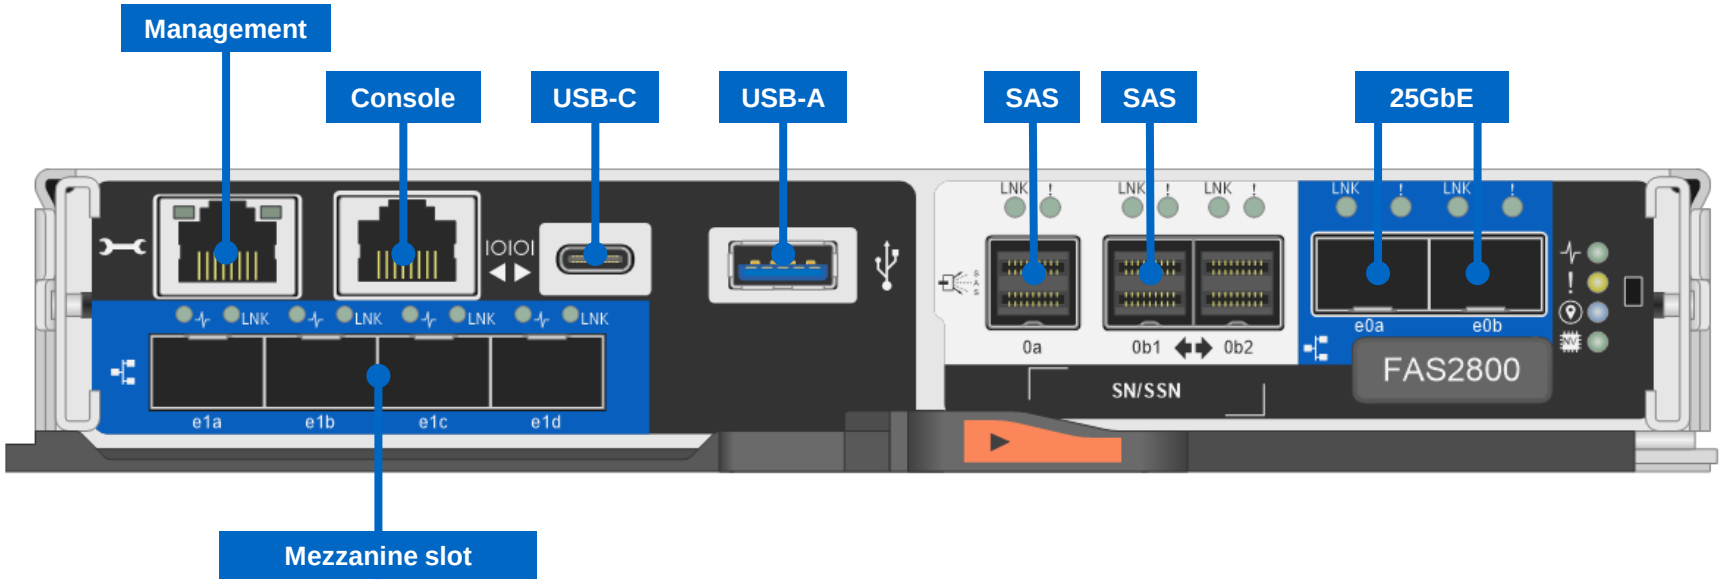

Additional details present in slide notes.

FRU map

Take a closer look

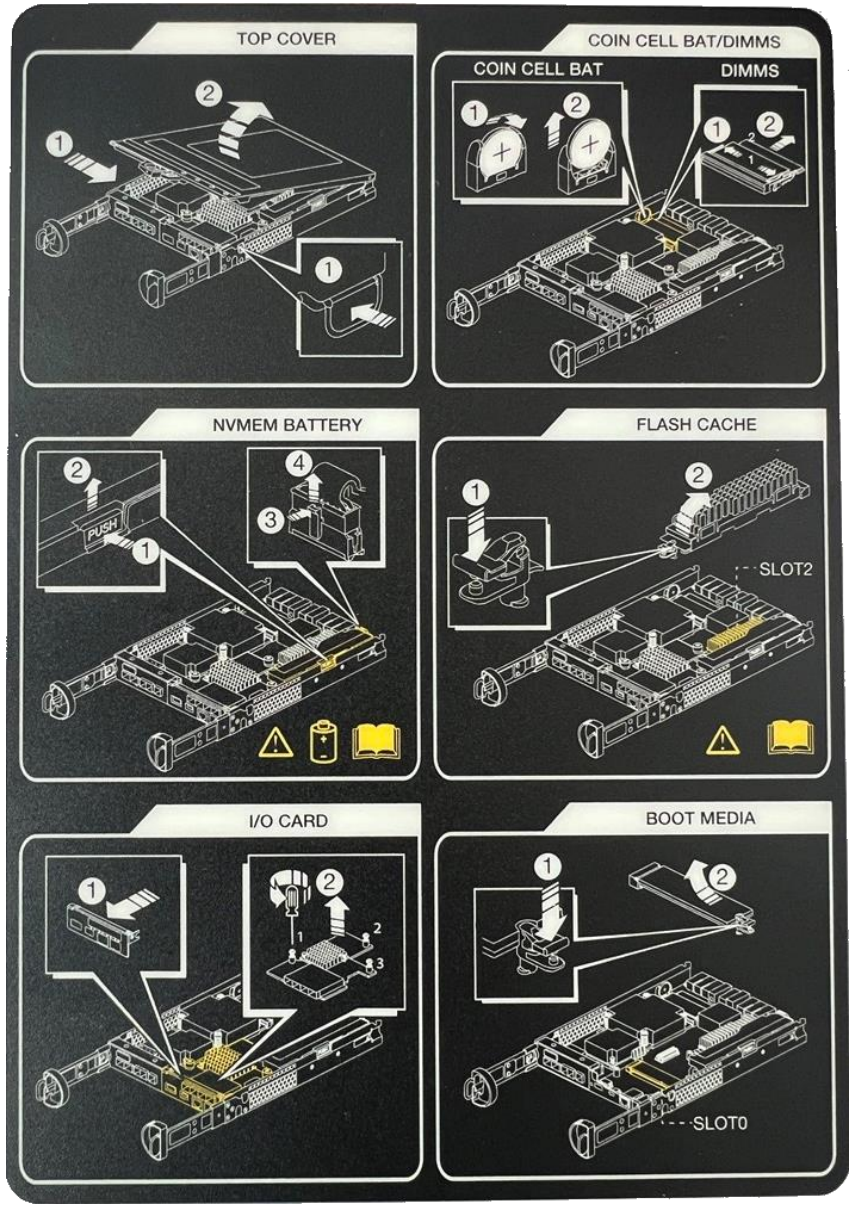

PCM block diagram

Take a closer look

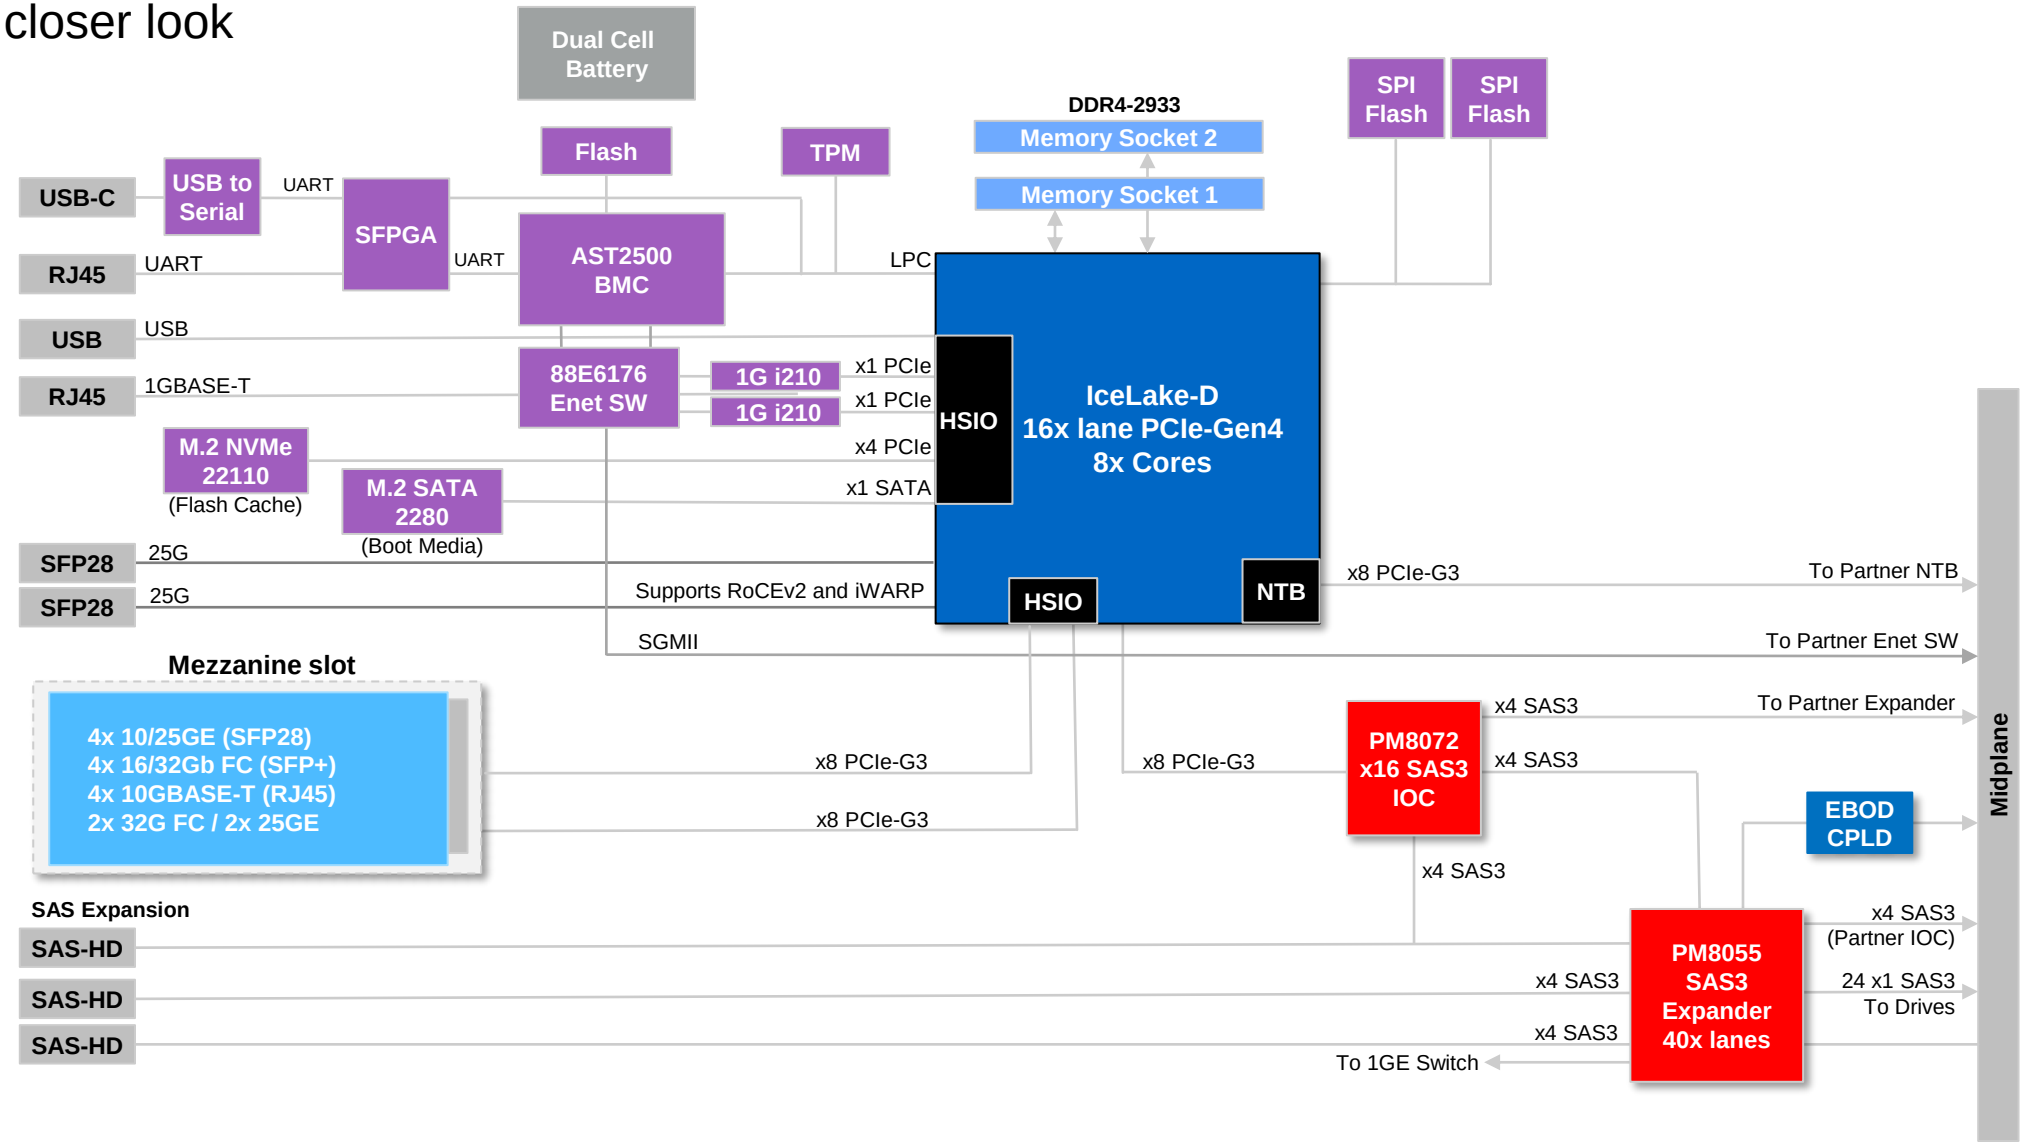

NVMEM architecture

Persistent write cache

- In the event of a dirty shutdown, the NVMEM battery maintains power to the DIMM slots so in-core writes are protected during reboot
- In the event of unexpected loss of power
  - NVMEM battery keeps motherboard powered up temporarily
  - NVMEM contents are destaged to boot device
  - Motherboard powers down after NVMEM data has been destaged to boot device
- Offers persistent write cache for write data not committed to drives
  - No need to worry about minimum 72-hour battery life to protect NVMEM data
- Destaged data stored on boot device is encrypted for added security
- A normally charged battery capable of about 25 destage events

NVMEM destage/restage

- Destaging: NVMEM battery keeps controller powered up during event

```
**********************************
* Pending cached data is detected *
* Starting the process to save to *
* media...                        *
**********************************
Destage is in progress, please wait
.....................
Destage is completed successfully
Powering off from the backup power source
```

- Restaging: NVlog replayed from flash to recover writes if needed:

```
******************************
Mon May 25 18:36:16 2023 [nv2flash.restage.progress:NOTICE]: ReStage is going to restore
non-volatile data from flash in approximately 21 second
```


Hot-swappable FRUs

- Denoted by terracotta (orange) cam lever or thumbscrew
- Controllers
  - Do not turn off PSUs to remove a controller
- PSUs
  - One PSU can provide sufficient power for the HA pair and internal drives

Cluster interconnects

- Onboard 25GbE ports provide both HA and cluster connectivity
  - Required for both 2-node switchless cluster (TNSC) and switched cluster configurations
  - Cannot be used for host network attachment
  - MetroCluster IP is not supported
- Default cluster interconnects are 25GbE ports
  - e0a and e0b
- Supported cluster switches
  - Refer to the NetApp Hardware Universe for additional details about supported cluster connections

HA Pair: FAS2820

Two-node switchless cluster (TNSC)

- FAS2820 systems require 2x 25GbE cluster interconnections
- HA & cluster interconnects
  - Connect ports e0a to e0a
  - Connect ports e0b to e0b

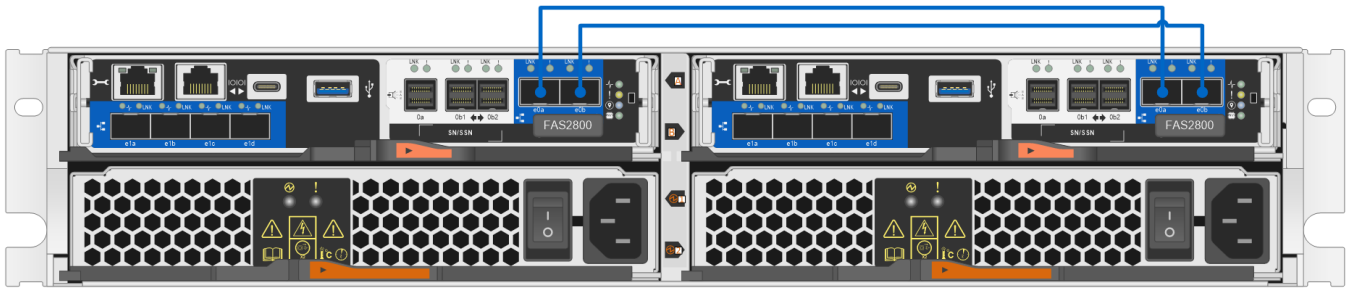

HA Pair: FAS2820

Two-node switched cluster

- FAS2820 systems require 2x 25GbE cluster interconnections
  - Connect 25GbE cables to SFP28 switch ports
    - 10GbE cables (SFP+) are also supported
  - Attach 1 connection to each controller per switch
- HA & cluster interconnects
  - Switch 1
    - FAS2820A-01: e0a
    - FAS2820A-02: e0a
  - Switch 2
    - FAS2820A-01: e0b
    - FAS2820A-02: e0b

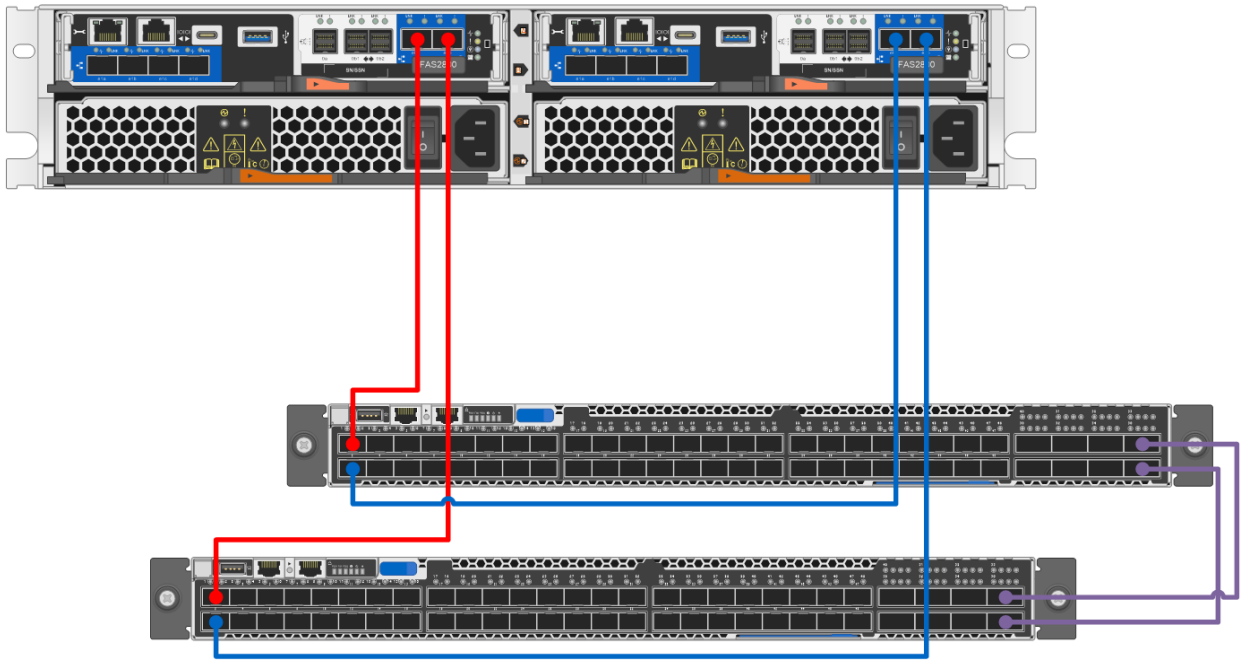

Mezzanine cards

Slot configurations

Description	NetApp P/N
4-port, 10/25Gb Ethernet, SFP28, (RoCEv2)	X1166A
4-port, 32Gb Fiber Channel, SFP+ (initiator/target)	X1167A
4-port, 32Gb Fiber Channel, SFP+ (target only)	X1168A
2-port 10/25Gb Ethernet, SFP28 + 2-port 32Gb FC, SFP+ (initiator/target)	X1169A
4-port, 10GBASE-T, RJ45	X1170A

Onboard SAS ports

- FAS2820 controllers have three onboard 12Gb SAS-3 ports
  - SAS port 0a is connected to the SAS controller
  - SAS ports 0b1 and 0b2 are connected to a SAS expander
- To implement MP-HA for internal or external storage use SAS ports 0a and 0b1
  - Use of SAS port 0b2 is not supported in ONTAP
- Inserting a cable into SAS port 0b2 will automatically disable the port
  - ONTAP will report a warning that an unsupported port is in use

MPHA cabling: FAS2820

Internal storage only

- FAS2820 HA pair
  - Connect ports 0a to 0a between controllers
  - Connect ports 0b1 to 0b1 between controllers
    - Use of port 0b2 is not supported
- Single stack configuration
- 12Gb SAS Cabling
- Refer to the 12Gb SAS SE Presentation for additional information

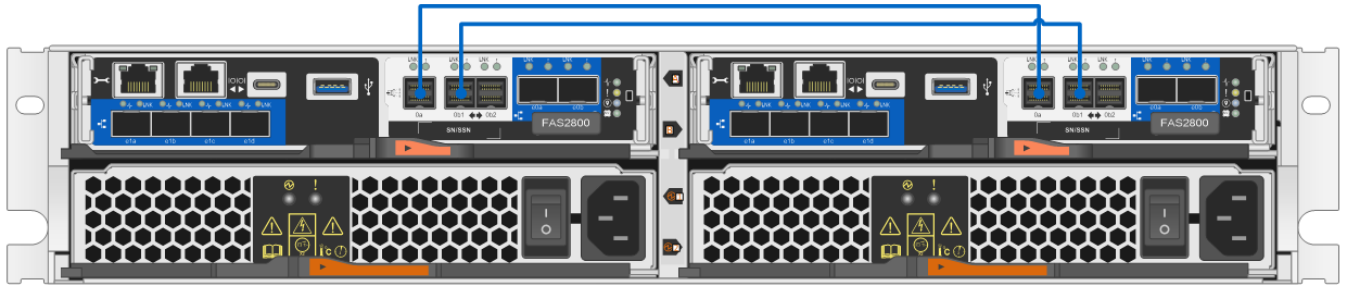

MPHA cabling: FAS2820

External storage

- FAS2820 HA pair
  - Connect ports 0a and 0b1 to external storage
    - Use of port 0b2 is not supported
- 3x DS212C or DS224C shelves
  - Single stack configuration
- 12Gb SAS Cabling
- Refer to the 12Gb SAS SE Presentation for additional information

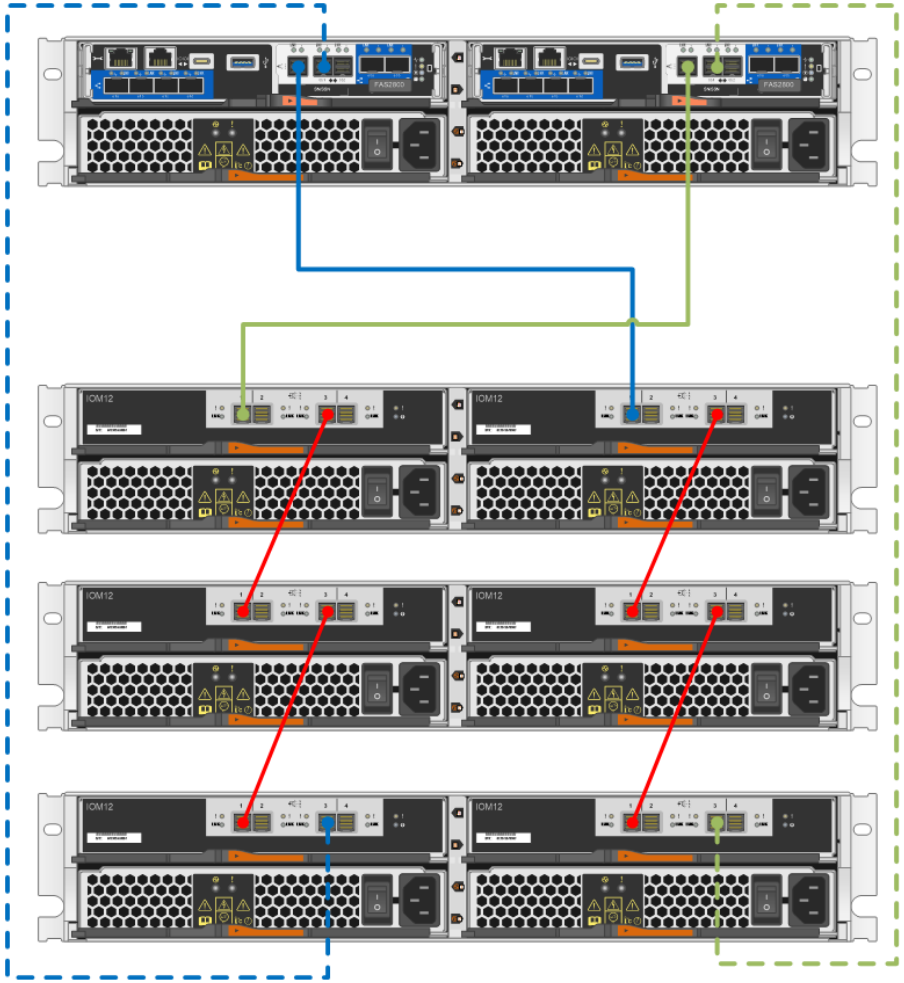

Using the onboard USB port

- USB 2.0 port enabled in ONTAP on FAS2820
- USB port can be used to
  - Boot ONTAP if the onboard boot media has failed
  - Install a release of ONTAP or a firmware service image
- The external USB port is only enabled during these activities
  - Activation is based on specific commands
- Requires USB devices have FAT32 formatting

USB port use cases

Scenario	Prerequisites	Command
Perform boot device recovery from LOADER prompt	• USB 2.0 device formatted to FAT32 with correct ONTAP image.tgz • Not hot pluggable, need to boot to LOADER prompt after USB device is inserted	• **`boot_recovery`** using the ONTAP image at the LOADER prompt • Select appropriate ONTAP image at the boot menu
Fetch ONTAP software for installation	• USB 2.0 device formatted to FAT32 with correct ONTAP image	• **`system node image update/get`** • Command has an additional option to fetch ONTAP from USB device
Fetch service image for firmware update	• USB 2.0 device formatted to FAT32 with correct service image	• **`system node firmware download`** • Command has an additional option to fetch ONTAP from USB device

Software and firmware levels

Minimum revisions

- ONTAP 9.13.1RC1
- Baseboard Management Controller (BMC) firmware level 17.1P1
- BIOS level 19.2

Baseboard Management Controller

Enhances manageability and RASM

- Baseboard Management Controller (BMC) available through Ethernet or serial console connections
  - FAS2820 leverages an IPMI-compliant BMC in lieu of a Service Processor (SP)
  - Shares the same Ethernet interface with e0M
    - Port speed will auto-negotiate to 1Gbps
    - Toggle from serial console into BMC with Ctrl-G
    - Toggle back to serial console with Ctrl-D
- BMC is used to generate system cores
  - NMI reset button is not present
- Actively manages some hardware
  - Advanced sensor management, FRU tracking, fans
  - Maintains the System Event Log (SEL) for hardware troubleshooting
- BMC firmware is updated automatically
  - BMC will automatically reboot after firmware update is completed
    - Node operations will not be impacted during BMC reboot

Baseboard Management Controller

Best practices

- Enable BMC NetApp® Active IQ®
  - Active IQ configuration is inherited from the controller ONTAP
    - Use **`system service-processor autosupport show`** to verify that **`–is-enabled`** is set to ‘true’
  - BMC utilizes SMTP for Active IQ transmission
    - Ensure that an Active IQ mailhost is configured within ONTAP
- Enable automatic BMC firmware updates
  - Use **`system service-processor show`** to verify that **`–autoupdate-enabled`** is set to ‘true’
- Configure BMC/wrench network interface on a separate management network
  - Ensures isolation from networks with heavy broadcast/multicast traffic
- Enable hardware-assisted partner takeover
  - Requires BMC IP address to reach HA partners management LIF
- Ensure that BMC address is accessible from Cluster Management LIF of the cluster
  - Aides in remote log collection, network firmware updates, etc.

FAS2820 system diagnostics

System diagnostic tool

- Diagnostic tool is used after FRU replacement
  - From the LOADER prompt run **`boot_diags`** command to execute system-level diagnostics
  - ASCII-based menu-driven interface includes a user guide
  - Replaces the former SLDIAG diagnostic tool set

Function	Description
Scan system	Collect FRUs and key component information
Test system	Test each system component
Test memory	Test all system memory DIMMs
Show VPD information	Displays VPD of all components
Show FW revision	Displays firmware revision of each programmable device
Show MAC address	Displays MAC address of each Ethernet and FC port
Show logs	Displays logged output of previously executed tests. Logs are available after scan, test system and test memory
Reboot to LOADER	

System diagnostics

SLDIAG

- SLDIAG replaces SYSDIAG
  - Both perform system-level diagnostics
- Major differences in SLDIAG
  - SLDIAG runs from maintenance mode
    - Has a CLI interface
  - SYSDIAG booted with a separate binary
    - Used menu-driven tables
- Review the SLDIAG TOI if you perform diagnostics

Power architecture

- FAS2820 systems have two power supplies (PSUs)
  - PSUs provide direct power to the chassis
  - Chassis provides indirect power to each controller and internal HDDs
- Air flows in through disks, then splits between power supplies and PCMs
- One failed PSU is allowed per chassis
  - Controller will run indefinitely with single failed PSU per chassis
  - System continues to operate with one failed PSU
  - All fans or blowers must be operational
- Blowers are powered by the backplane and will continue to run even if the PSU fails
  - Leave the failed PSU in place until replacement arrives

Environmental details

Power, thermals, weight, dimensions, acoustics, product certificates

- For the latest and most complete environmental information, refer to:

Hardware Universe

(https://hwu.netapp.com)

Certification information

EMI certifications

Type	Country	Test Standards	Notes
Safety	U.S.	60950-1	
	CAN	CAN/CSA-22-2 NO.60950-1	Equivalent to UL 60950
	EU	EN 60950-1	
Emission	U.S.	FCC part 15, Subpart B, Class A	FCC Class A minimum
	CAN	ICES-003	
	EU	EN 55022, Class A	Emission of ITE
		CISPR22	
		EN 61000-3-2	AC Line Harmonic
		EN 61000-3-3	AC Line Flicker
Immunity	U.S.	N/A	
	CAN	N/A	
	EU	EN 55024	Immunity of ITE
		EN 61000-4-2	ESD
		EN 61000-4-3	RF
		EN 61000-4-4	Fast Transient
		EN 61000-4-5	Surge
		EN 61000-4-6	Conducted
		EN 61000-4-11	Line Dip/ Interrupt
Environmental		ASHRAE A4	Compliant

Thank you