Open CloudServer JBOD specification V1.0

Authors:

Mark Shaw, Director of Hardware Engineering, Microsoft

Martin Goldstein, Principal Systems Architect, Microsoft

1 Revision History

Date	Name	Description
1/28/2014		Version 1.0

2 Scope

This document provides the technical specifications for the design of the 6G half-width JBOD blade for the Open CloudServer system.

3 Contents

4 Overview

This document describes the Open CloudServer system "just a bunch of disks" (JBOD) storage blade. The JBOD blade is a half-width form factor blade design supported in the chassis system. The design includes the metal tray parts, 6 Gigabit (6G) expander board, hard disk drive (HDD) board, and cables.

The 6G JBOD blade is a cost-effective, dense storage solution built to scale and meet storage application requirements. Like the compute blade, the 6G JBOD blade is a modular blade designed to ease serviceability, reduce system downtime, and simplify system manageability.

5 Introduction to the 6G JBOD Blade

The 6G JBOD is a 1U (1 rack unit) half-width form factor storage blade that supports 10 x 3.5" HDDs. The blade is designed so that it can be serviced independently from the chassis and the rack infrastructure, without the need to disconnect cables or change the state of any adjacent blades.

The 6G JBOD blade includes two types of boards: a single serial-attached small computer system interface (SAS) expander board and two HDDs backplane boards.

The SAS expander board provides the connection to the local 10 x HDDs and mates with the tray backplane board; this enables the power and data signals to the blade and outside of the blade. Leveraging blind-mate connectors, the SAS expander board mates to the tray backplane where the external SAS I/O ports are located; this configuration makes it possible for host nodes or other JBODs to connect to the expander board.

Figure 1 shows a view of the 6G JBOD blade (reference design shown).

Figure 1. 6G JBOD blade overview

5.1 Front View

The fronts of all blades have a similar look and a similar feature set. Each blade is secured and released with a latching handle and turn screw. Electromagnetic interference (EMI) containment is achieved with shielding materials and sized openings on the front of the blade. Each blade has two LEDs located at the front bottom right of the blade; these LEDs indicate power and attention status.

Figure 2 shows a front view of the 6G JBOD blade (reference design shown).

Figure 2. Front view of 6G JBOD blade

Between the shielding materials on the lip of the front of the blade, there is a label that helps identify the HDD location; this makes it easier to service and bring up the system.

Figure 3 shows the front lip of the 6G JBOD blade (reference design shown).

Figure 3. Front lip view of 6G JBOD blade

5.2 Rear View

The rear of the 6G JBOD mates with the rest of the system infrastructure. There are five mating components on the rear side of the blade: two FCI Electronics AirMax VSe 3x6 connectors, two mechanical guide pins, and one AirMax power connector.

The AirMax VSe 3x6 connectors carry all of the storage I/O, blade management, and ground signals (no blade power is carried through the AirMax VSe 3x6 connector). The AirMax power connector delivers the 12V power. The two mechanical guide pins enable blade blind-mating to the chassis infrastructure. The 6G JBOD blade tray lies within the 1U mechanical full-width tray, which has additional guide and stop pins that help align the blade.

There is an optional BOM-loaded LED that is visible from the rear of the blade. This LED is connected to the general purpose input output (GPIO) LED pin on the expander, which provides additional debug and serviceability features.

Figure 4 shows the rear view of 6G JBOD blade (reference design shown).

Figure 4. Rear view of 6G JBOD blade

6 Blade Design

The 6G JBOD blade is a simplified storage solution with minimal hardware components. It sits in the mechanical tray and attaches to the tray backplane board. The blade consists of a small number of components and boards that can be easily accessed and serviced once the blade is outside of the chassis.

Figure 5 shows the component layout of the 6G JBOD blade mated with the tray backplane board.

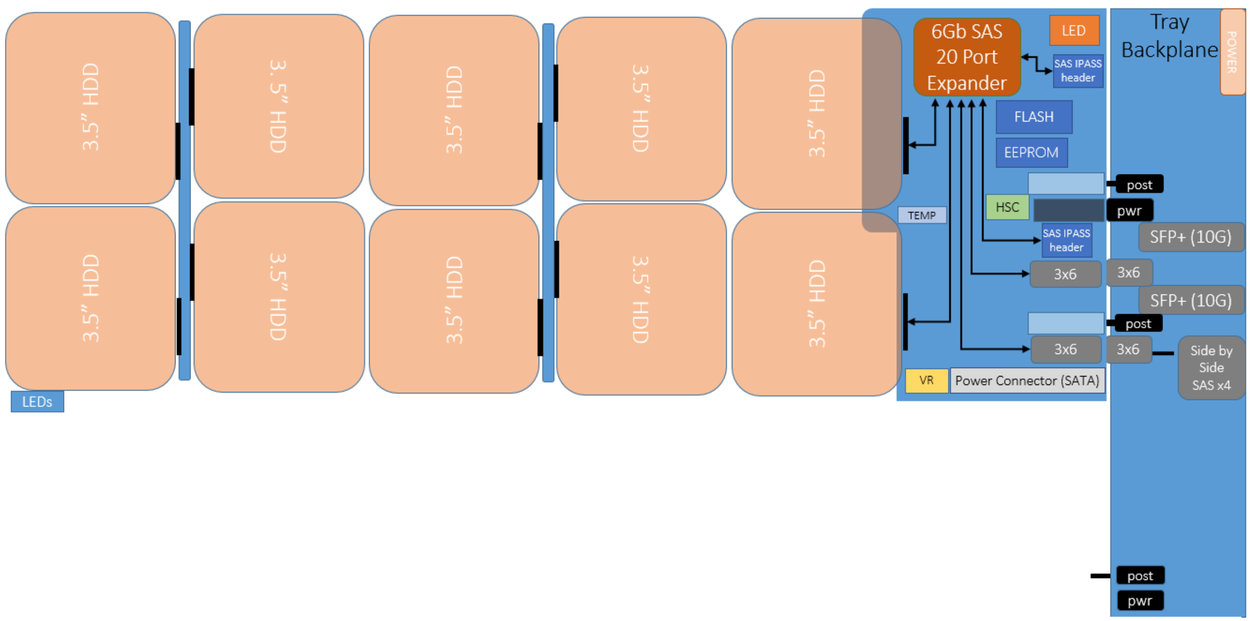

Figure 5. System component layout

6.1 Storage Solution Design

The 6G JBOD blade enables a complete storage solution within the system infrastructure. Figure 6 shows a solution-level block diagram.

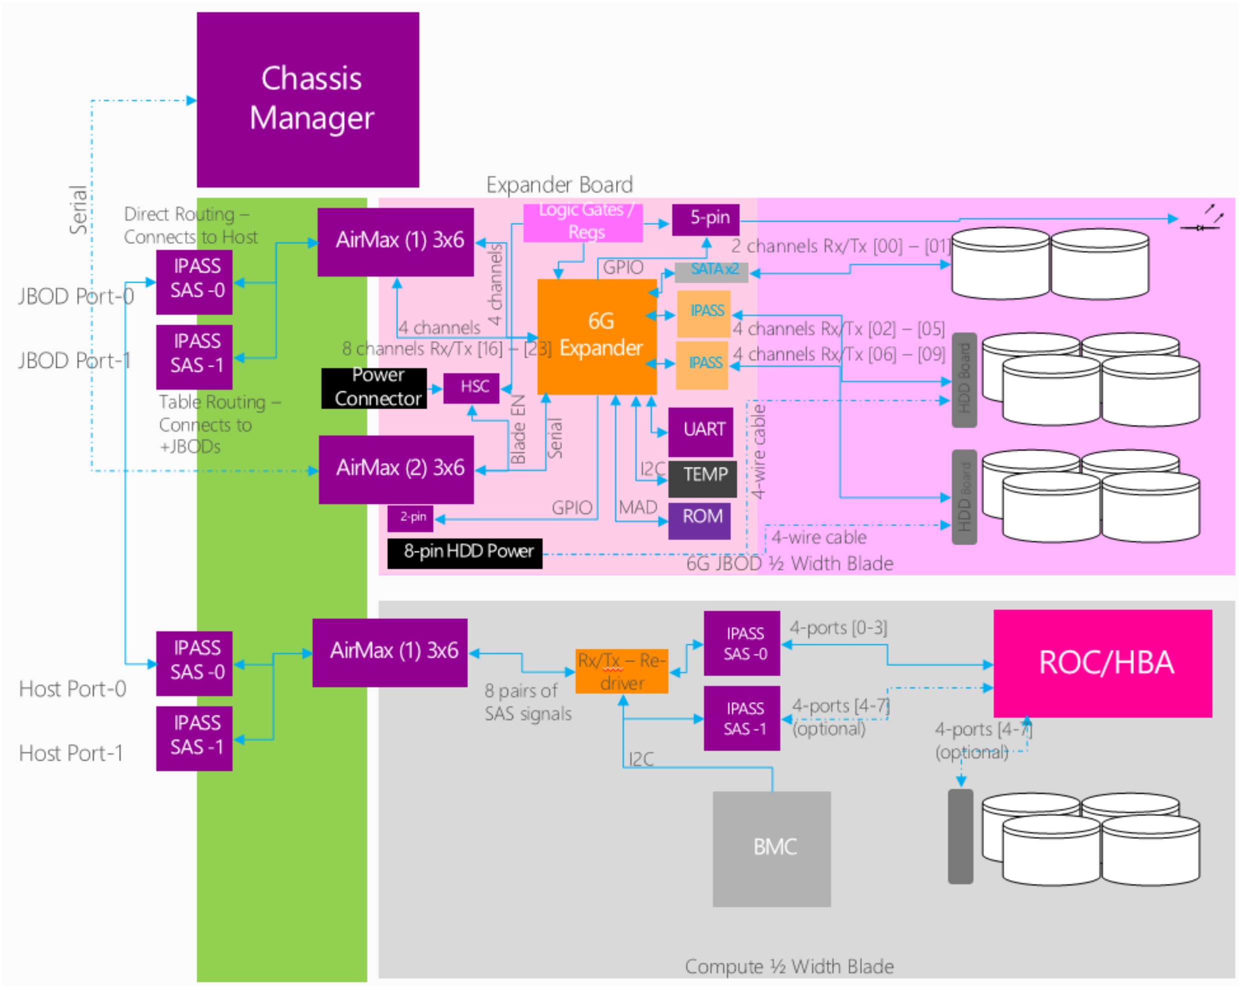

Figure 6. System block diagram

Figure 6 shows the tray backplane, the 6G JBOD expander board, the host blade (compute blade with a storage controller), and the management hardware console (chassis manager).

Both the host and JBOD blades mate with the tray backplane, which acts as a mid-plane joining the two types of blades with a single Molex iPass mini-SAS external cable. The tray backplane resides within the 1U full-width mechanical tray in the chassis, which can connect to either a compute blade or a 6G JBOD blade. Various blade combination types are supported within the system infrastructure; this makes it possible to optimize the rack-level configurations for the application.

Within the 6G JBOD blade, the HDDs are connected to serial AT attachment (SATA) backplane boards that carry the SATA signals through cables to the expander board. There are two SATA backplane boards that allow eight HDDs to be connected to the expander board by cables. Two additional drives are connected to the expander board by a SATA connector located on one of the edges.

The compute blade is located at the other end of the cable. The compute blade includes options for installing various RAID-On-Chip (ROC) and Host-Bus Adapter (HBA) controllers, enabling various storage blade feature sets and functionality. An Rx/Tx re-driver is also located inside the MCS compute blade, which helps with signal integrity on the data path.

The chassis manager, a localized management hardware console, makes out-of-band (OOB) JBOD blade management possible. Communication to the SAS expander is via a Universal Asynchronous Receiver/Transmitter (UART) interface to the chassis manager. More detail about the chassis manager and additional devices (such as the compute blade and tray backplane) can be found in the other Open CloudServer specifications.

6.2 Expander Board Design

The SAS expander board acts as the central hardware component, connecting the end-devices to the host node. The expander board is located at the rear of the 1U half-width mechanical blade. The board also distributes the power going to the HDDs and the signals going to LED panel.

Figure 7 shows a view of the expander board (reference design shown).

Figure 7. Expander board

Figure 8 shows a block diagram of the expander board.

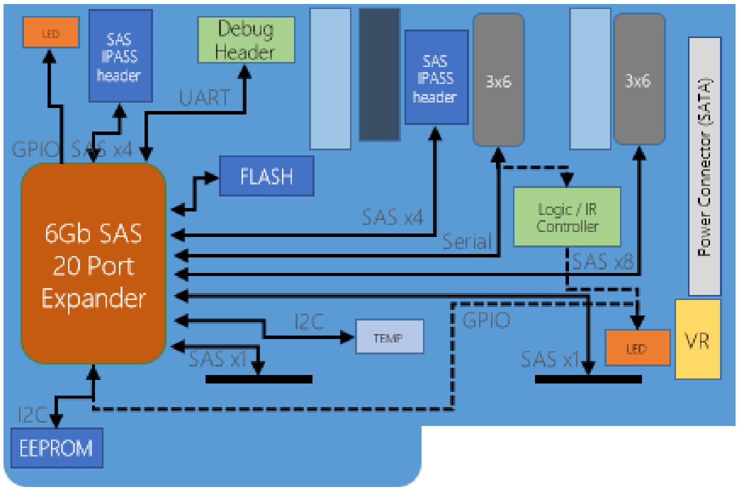

Figure 8. Blade expander board block diagram

The expander uses LSISAS2x20, a reduced packaged LSI 6G SAS expander controller. The controller carries the SAS data signals to the AirMax VSE 3x6 connectors that mate to the tray backplane board. The controller is simplified to meet the minimum feature requirements of the blade. Many of the GPIO and I2C pins are not connected. The system is managed through selected out-of-band and in-band functions. Details of the management subsystems are included in this document and in the Open CloudServer chassis management specification.

6.2.1 Expander Board Connectors

The 6G JBOD blade includes a variety of connectors for mating the blade to the tray and for connecting to HDDs, status indicators, and debug ports.

6.2.1.1 AirMax Signal and Power Connectors

The blade includes two types of FCI AirMax connectors: VSe connectors for the SAS I/O data signals and power connectors for the 12V power. Both types of connectors mate with the tray backplane.

AirMax Signal Connectors

Two 3x6 VSe connectors (FCI AirMax VSe family coplanar connectors) are used for signals only (not used for power). An example of a coplanar connector pair is shown in Figure 9. The connector layout is shown in Figure 10, with the expander board header on the left and the mating tray backplane receptacle on the right.

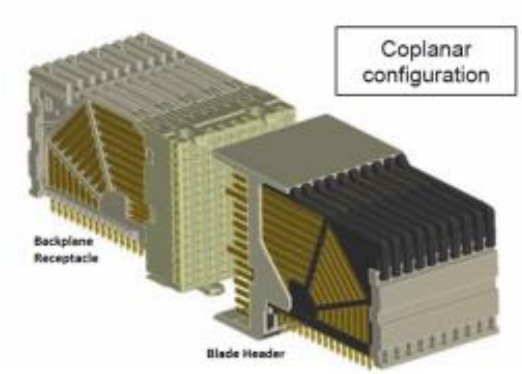

Figure 9. Blade signal connector, coplanar example

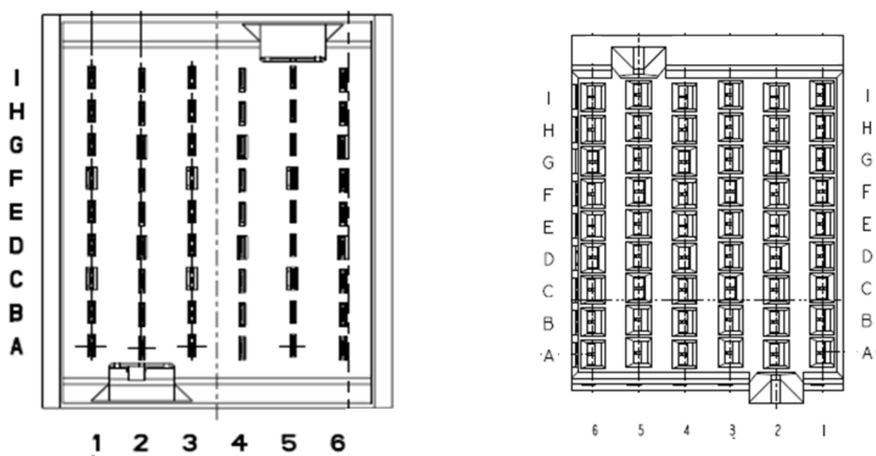

Figure 10. AirMax VSE pin arrangement, with blade header on left and backplane receptacle on right

Table 1 shows the bus type distribution for the two AirMax VSe 3x6 connectors. Table 2 and Table 3 list the connector pin definitions.

Table 1. AirMax VSe 3x6 Connector Pin Definition

Bus Type	Single Ended	Diff.	GND	Pin Count	Definition for 3 pair 6 column
AirMax VSE 3x6 (1)					
SAS - 8 Channels		16	16	48	SAS from Controller to JBOD
RSVD		2	2	6	Reserve in future.
Total Pin Count				54	Airmax 3 pair 6 column signal Connector
AirMax VSE 3x6 (2)					
Node EN #1	1			1	1 JBOD Node per sled from Mid-Plane
UART #1	2			2	1 JBOD Node per sled from Mid-Plane, RM and Blade communication bus
Mating	1			1	Connector insertion mating
Tray BP SKU ID	2			2	Tray Backplane SKU identification, 1Gb, 10Gb, SAS 6Gb, SAS 12Gb, Infinite band
RSVD		34		34	Reserve in future.
Ground			14	14	
Total Pin Count				54	Airmax 3 pair 6 column signal Connector

Table 2. AirMax VSe 3x6 (1) Connector Pin Definition

Airmax 3x6 (1) CONNECTOR PIN DEFINITION						
SAS Signals						
	1	2	3	4	5	6
A	SAS_TD_CH2P	GND	RSVD	GND	SAS_TD_CH6P	GND
B	SAS_TD_CH2N	SAS_RD_CH3P	RSVD	RSVD	SAS_TD_CH6N	SAS_RD_CH7P
C	GND	SAS_RD_CH3N	GND	RSVD	GND	SAS_RD_CH7N
D	SAS_RD_CH1P	GND	SAS_RD_CH4P	GND	SAS_RD_CH6P	GND
E	SAS_RD_CH1N	SAS_TD_CH3P	SAS_RD_CH4N	SAS_TD_CH5P	SAS_RD_CH6N	SAS_TD_CH8P
F	GND	SAS_TD_CH3N	GND	SAS_TD_CH5N	GND	SAS_TD_CH8N
G	SAS_TD_CH1P	GND	SAS_TD_CH4P	GND	SAS_TD_CH7P	GND
H	SAS_TD_CH1N	SAS_RD_CH2P	SAS_TD_CH4N	SAS_RD_CH5P	SAS_TD_CH7N	SAS_RD_CH8P
I	GND	SAS_RD_CH2N	GND	SAS_RD_CH5N	GND	SAS_RD_CH8N

Table 3. AirMax VSe 3x6 (2) Connector Pin Definition

Airmax 3x6 (2) CONNECTOR PIN DEFINITION						
Management/Mating Signals						
	1	2	3	4	5	6
A	RSVD	GND	RSVD	GND	RSVD	GND
B	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD
C	GND	RSVD	GND	RSVD	GND	RSVD
D	RSVD	GND	RSVD	GND	RSVD	GND
E	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD
F	GND	RSVD	Tray_BP_SKU_ID0	RSVD	GND	RSVD
G	RSVD	GND	NODE0_RXD	RSVD	RSVD	GND
H	RSVD	NODE0_EN_N	NODE0_TXD	RSVD	RSVD	RSVD
I	GND	RSVD	RSVD	Tray_BP_SKU_ID1	BLADE1_MATED_N	RSVD

AirMax Power Connector

The power connector, which supplies 12V to the 6G JBOD blade, is in the FCI AirMax Power family. Figure 11 shows the connector layout, and Table 4 shows the pin definition.

Note that this connector is also included in the compute blade.

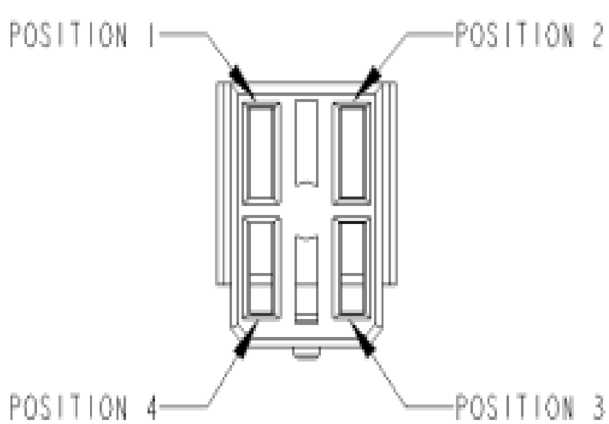

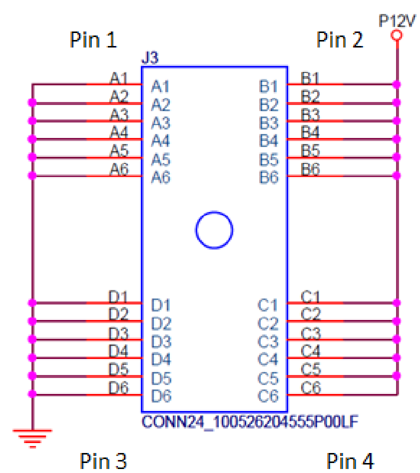

Figure 11. AirMax power receptacle pinout arrangement

Table 4. AirMax Power Receptacle Pin Definition

Pin	Signal name	Capacity
Position 1	12V return – first mate	18A
Position 2	12V supply	18A
Position 4	12V supply	18A
Position 3	12V return	18A

The maximum power capacity on the 6G JBOD blade with a single connector is 345W, 12V x 2 x 18 A x 80% de-rating. Note that depending on the implementation, these maximums might be unachievable.

6.2.1.2 Additional Connectors

There are six primary connectors on the 6G JBOD expander board, including the two described in the previous section (AirMax signal and power connectors). Table 5 provides further detail.

Table 5. 6G JBOD Expander Board Connector List

Assembly mounting	Assembly attachment	Connector description	Quantity	Manufacturing part number
Expander board	Tray backplane	Power header AirMax VS, 2x2	1	(FCI) 10028918-001LF
Expander board	Tray backplane	Signal connector—three pair, 54 contact, 2mm spacing, 17mm pitch, six column	2	(FCI) 10039851-101LF
Expander board	Tray backplane	Guide pin receptacle—10.8mm right angle, 0° key	2	(FCI) 10037912-101LF
Expander board	SAS/SATA cable	6G IPASS SFF-8087 Mini-SAS Vertical header	2	(Molex) 75784-0147
Expander board	HDD	SATA signal and power 29-pin	2	(Molex) 97945-0001
Expander board	Power cable	8-pin power vertical connector	1	N/A

The 29-pin SATA signal and power receptacle connector on the expander board provides connectivity to a single SATA HDD. There are two 29-pin SATA connectors that allow direct connection of HDDs to the expander board. The connector follows

SFF-8482 specification, enabling serial signals and power on different contacts within the same housing.

Table 6 provides the pin signal description for this 29-pin receptacle connector.

Table 6. 29-Pin Signal Connector Description

Pin	Signal
S1	Ground
S2	Transmit +
S3	Transmit -
S4	Ground
S5	Receive -
S6	Receive +
S7	Ground
S8	NC
S9	NC
S10	NC
S11	NC
S12	NC
S13	NC
S14	NC
P1	3.3V/NC
P2	3.3V/NC
P3	3.3V/NC
P4	Ground
P5	Ground
P6	Ground
P7	5V
P8	5V
P9	5V
P10	Ground
P11	NC
P12	Ground
P13	12V
P14	12V
P15	12V

Two HDDs in the JBOD blade are powered directly through the 29-pin connector; the remaining eight HDDs are powered through the 8-pin power connector, which connects to an 8-wire cable powering two HDD backplane boards.

Table 7 provides the pin description for the 8-pin vertical connector.

Table 7. 8-Pin Power Connector Description

Pin	Signal
1	12V
2	Ground
3	5V
4	Ground
5	12V
6	Ground
7	5V
8	Ground

6.2.2 Expander Board Component List

Following is list of the additional hardware components on the JBOD expander board:

- **Flash Read-Only Memory (ROM) device**

  Spansion S29JL064J

  Provides storage for the local expander firmware, configurations, and various initialization method settings.

- **Serial Electrically Erasable Programmable Read-Only Memory (EEPROM)**

  STMicroelectronics (ST) SG24C02

  Provides storage for device initialization information.

- **In-rush controller**

  Analog Devices ADM1276

  Provides in-rush current control through the 12V bus rail. More information about in-rush power control signal can be found in the compute blade specification.

- **Temperature Sensor**

  Texas Instruments (TI) TMP75

  Enables temperature monitoring. Temperature "High/Low" warning and critical thresholds can be set through the expander.

6.3 Storage HDD Backplane Board Design

The second type of board within the 6G JBOD blade is the storage HDD backplane board, a double-sided printed circuit board assembly (PCBA) that acts as a mid-plane between the 6G JBOD expander board and the HDDs. Cables are used to connect the expander and storage HDD backplane board; these cables are 3M SAS fan-out cables that have a SFF-8087 SASx4 internal header on one end, and standard SATA 7-pin single port receptacles on the other end. A single board can support 2x HDDs on each side (up to 4x HDDs per board). There are two HDD backplane boards in each 6G JBOD blade, allowing up to 8x HDDs to be supported via the HDD backplane board. Additional 2x HDDs can be supported with direct connections to the expander board, as described in the previous section.

Figure 12 and Figure 13 show views of the storage HDD backplane board (reference design shown).

Figure 12. Storage HDD backplane board, side 1

Figure 13. Storage HDD backplane board, side 2

Table 9 provides detail to the connectors on the storage HDD backplane board.

Table 9. 6G JBOD Expander Board Connector List

Assembly mounting	Assembly attachment	Connector description	Quantity	Manufacturing part number
HDD backplane	SAS/SATA cable	SATA signal 7-pin	4	(LOTES) ABA-SAT-054-P56
HDD backplane	HDD	SATA signal and power 22-pin	4	(Amphenol) SAS-F313-020-1-TR
HDD backplane	Power cable	4-pin power connector	1	N/A

Similar to the 29-pin SATA signal and power receptacle connector on the expander board, there is a 22-pin SATA signal and power receptacle connector on the HDD backplane board. This also provides connectivity to a single SATA HDD and follows SFF-8482 specification, enabling serial signals and power on different contacts within the same housing.

Table 10 provides the pin signal description of the 22-pin receptacle connector on the HDD backplane board.

Table 10. 22-Pin Signal Connector Description

Pin	Signal
S1	Ground
S2	Transmit +
S3	Transmit -
S4	Ground
S5	Receive -
S6	Receive +
S7	Ground
P1	3.3V/NC
P2	3.3V/NC
P3	3.3V/NC
P4	Ground
P5	Ground
P6	Ground
P7	5V
P8	5V

Pin	Signal
P9	5V
P10	Ground
P11	NC
P12	Ground
P13	12V
P14	12V
P15	12V

Table 11 provides the pin signal description of the 7-pin plug connector on the HDD backplane board.

Table 11. 7-Pin Signal Connector Description

Pin	Signal
S1	Ground
S2	Transmit +
S3	Transmit -
S4	Ground
S5	Receive -
S6	Receive +
S7	Ground

Table 12 provides the pin description for the 4-pin power connector.

Table 12. 4-Pin Power Connector Description

Pin	Signal
1	12V
2	Ground
3	5V
4	Ground

Table 13 provides detail of the connectors on the 3M SAS fan-out cable.

Table 13. 3M Cable Connector List

3M cable side	Cable description	Quantity	Manufacturing part number
SATA end	7-pin female SATA connector	4	(Yi-Lian Industrial) 6SAS07S-320B-D
IPASS mini-SAS end	Internal mini-SAS connector, SFF-8087	1	(3M) 78-9101-1588-6

6.4 Connectors Plating

The 6G JBOD blade is supported in environments up to 40°C with 90 percent relative humidity (RH). It is recommended that all connectors are plated with an appropriate thickness of gold plating with nickel under-plating; reliability testing and plating thickness consistency measurements are also recommended. Note that the 6G JBOD blade consists of connectors that have 15u" – 30u" Au with under-plating of up to 100u" Ni.

7 Mechanicals and Blade Physical Specification

The physical specifications for the 6G JBOD blade are detailed in the compute blade specification, including the dimensions of the space into which the blade is installed, the blade weight limits, a description of the guiding and latching features, and information about electromagnetic interference shielding. The sections that follow are a subset of the compute blade specification.

7.1 6G JBOD Blade 3D Drawings

Figure 14 shows views of the 6G JBOD blade. The drawing on the left is the bare-bone mechanical JBOD blade with no components; the drawing on the right is a mechanical JBOD blade with the expander board, HDD backplane board, and cable clips for routing.

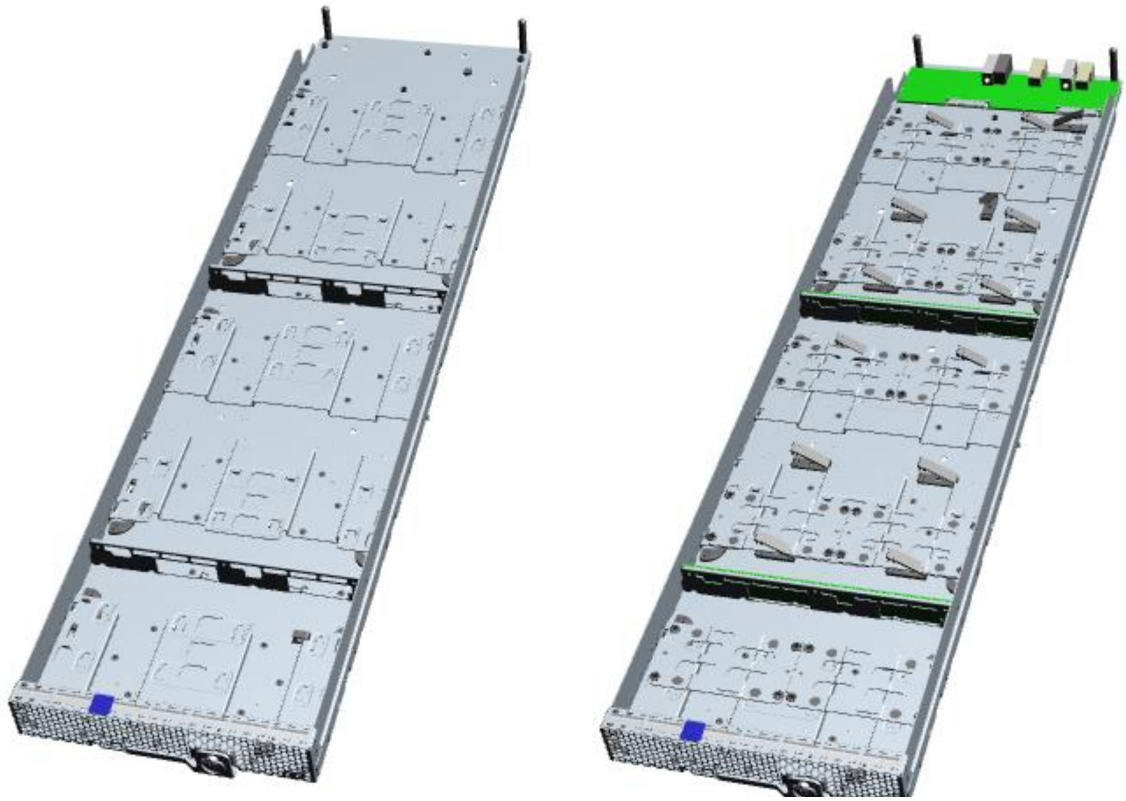

Figure 14. 6G JBOD blade 3D drawings

7.2 HDD Bracket

A single HDD bracket is required for each of the ten HDDs in the 6G JBOD blade. The HDD bracket makes it possible for 3.5" screw-mount HDDs to securely attach into the mechanical blade. The bracket must follow SFF-8300 specification requirements and meet alternative mounting locations in the SFF-8301 Rev.1.6 specification.

The HDD bracket must also include dampeners. The elasticity of the dampeners is determined based on HDD support; the recommendation is to test and support at minimum for near-line Enterprise 7200RPM SATA HDDs.

Figure 15 shows the HDD bracket (reference design shown).

Figure 15. HDD bracket

There are four large guide pins on the bracket that help with mating and securing the HDD with the blade sheet metal. The two side smaller pins are used in the HDD latching and releasing mechanism, which is described in the next section.

7.3 HDD Latch and Release Mechanism

Each of the 10x HDDs have separate latching and releasing mechanisms that allows the HDDs to be serviced and swapped quickly. There are two parts to the mechanism: the special HDD bracket (described in the previous section) and the plastic release lever.

Figure 16 shows the two pins used in the latch-and-release mechanism (reference design shown).

Figure 16. HDD bracket (2)

The HDD brackets have two dedicated pins that work with the plastic release lever. When the plastic lever is pressed on the outside of the blade, it lets the smaller pins move, allowing the HDD bracket to move away from the connectors and disconnecting the HDD from the blade.

Figure 17 shows the plastic release lever and the side of the blade where the lever can be pressed to release the pin.

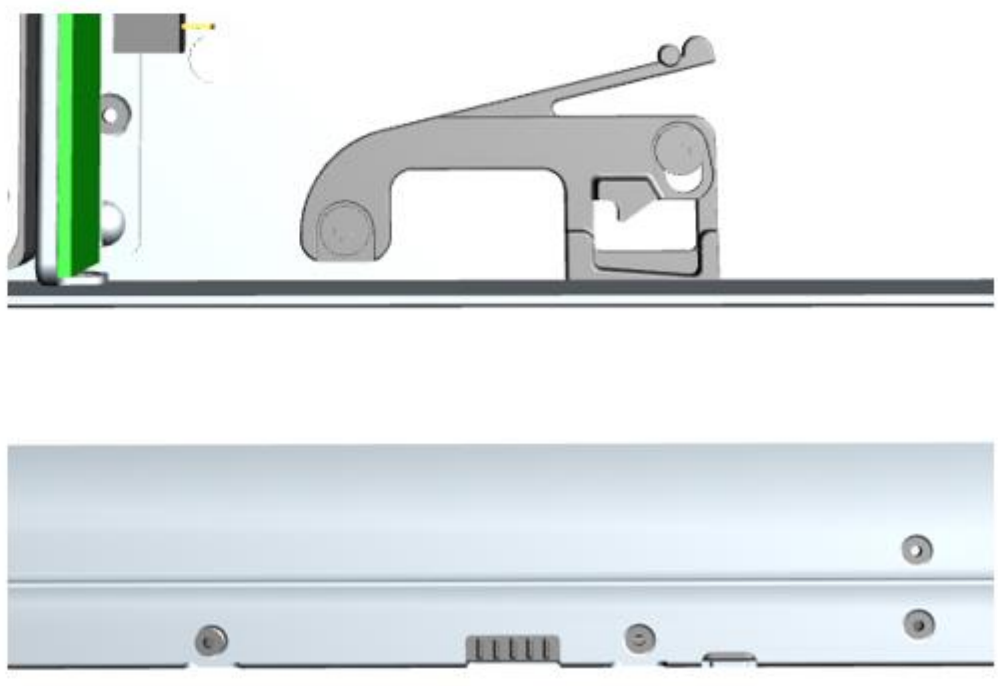

Figure 17. Plastic release lever

8 Electrical Specifications

The following sections provide a subset of information for the 6G JBOD blade electrical specification; more detail for the blades, 6G JBOD blade, and compute blade can be found in the compute blade specification.

8.1 Input Voltage, Power, Current

Table 14 lists the nominal, maximum, and minimum values for the blade input voltage. The maximum and minimum voltages include the impact of connector temperature, age, noise/ripple, and dynamic loading.

Table 14. Blade Input Voltage

Nominal voltage	Maximum voltage	Minimum voltage
12.25V DC	12.95V DC	11.65V DC

The maximum amount of power allowed per blade is defined during system power allocation. The number of blades in a chassis might be limited by the capacity of the AC power cord or by the cooling capacity of the deployment. Table 15 lists the input power allocation for a low-power blade.

Table 15. Input Power Allocation

	Nominal voltage	System power allocation (in W)	Maximum current (in A)
Low-power blade (1x power connector)	12.25VDC	300W	27.3A

The blade provides in-rush current control through the 12V bus rail; return-side in-rush control is not used. The in-rush current rises linearly from 0A to the load current over a 5 millisecond (ms) period. (Note that this time period must be no longer than 200ms.)

8.2 Current Interrupt Protection and Power, Voltage, and Current Monitoring

The 6G JBOD blade provides a cost-effective way to measure and report blade voltage and current consumption, and to make instance reporting available to the system. The blade includes power consumption measurements at the in-rush controller, and provides a way to interrupt current flow within 1 microsecond (μs) of exceeding three times the maximum current load.

9 Management Subsystem

The JBOD blade enables both in-band and out-of-band management. Like the compute blade, the 6G JBOD blade has two LEDs and an optional BOM-load LED for debugging.

More detail about the management subsystem can be found in the compute blade specification.

9.1 In-Band Manageability

The in-band functionality is implemented within the LSI expander (LSISAS2x20). Full management capability with the expander is enabled via the storage enclosure processor (SEP) software, which provides the SCSI Enclosure Service (SES) that monitors status and enables control of the HDDs. Various other management functions are enabled by SES, including HDD status, temperature monitoring, and DC voltage monitoring. The in-band functionality is supported through LSI utilities and software; management functions are fully supported though the LSI manageability implementation.

9.2 Out-of-Band Manageability/Blade API

In addition to self-monitoring and in-band management, the 6G JBOD blade permits out-of-band management through the management hardware console known as the chassis manager. The backplane communicates with the SAS expander through a UART interface.

The SAS expander firmware supports VT100 console output to enable communication with the backplane. The expander firmware also monitors the byte stream received on the UART interface for Intelligent Platform Management Interface (IPMI) packets.

The filtering for IPMI packets works as follows:

1. The received bytes are monitored for 0xA020. 0xA0 is the start byte for the IPMI packet, and 0x20 is the baseboard management controller (BMC) responder address. In this case, the expander plays the role of the BMC.
2. If 0xA020 is received, monitor for 0xA5 to signify the end of the IPMI packet. If 0xA5 is received, the packet is processed as an IPMI command. If 0xA5 is not received within 128 bytes, the whole byte stream is processed as a VT100 console input.
3. If 0xA020 is not received before the end of the packet or a carriage return, the byte stream is processed as a VT100 console input.

Following is the list of commands supported within the 6G JBOD blade:

- Get Device ID
- Get System GUID
- Get Channel Authentication Capabilities
- Read FRU Data
- Write FRU Data
- Get Disk Status
- Get Disk Info
- Get Chassis Status

Through the backplane, the 6G JBOD blade also provides a basic command-line interface (CLI) through the VT100 console as an alternative for users to monitor, configure, or update the JBOD blade. Features and functionality are based on the default LSI CLI modules from the LSI Bobcat software development kit (SDK).

Following is a subset of the list of functions supported:

- Show version
- Get/set SAS address
- Show disk status
- Show/clear all PHYS error count
- Read EEPROM data
- Update expander FW
- Reset expander

10 Device Support and Scaling

The sections that follow provide information about device support and scaling.

10.1 Device Support

The blade is intended to support single port storage devices only. The path from the expander to each SATA connector on the HDD board uses cable through one SAS path. At minimum, the 6G JBOD blade needs to support 6Gb/s near-line Enterprise 7200RPM SATA HDDs. The current design does not support different HDD types within each JBOD blade.

10.2 Cascading

It is possible to scale the storage capacity by daisy-chaining the 6G JBOD blades (cascading JBODs). Based on the selected LSI expander (LSISAS2x20), the expander table routing enables connections with as many as 1024 SAS addresses. It is recommended, however, that the maximum recommended cascading JBODs in series per x4 port is four layers; additional layers can result in increased latency and performance degradation to the application. Capacity of a storage solution with four layers is eight JBODs and 80 HDDs.

Figure 18 shows a storage topology with 4x 6G JBOD blades daisy-chained off of a single x4 port on the storage head node, compute blade.

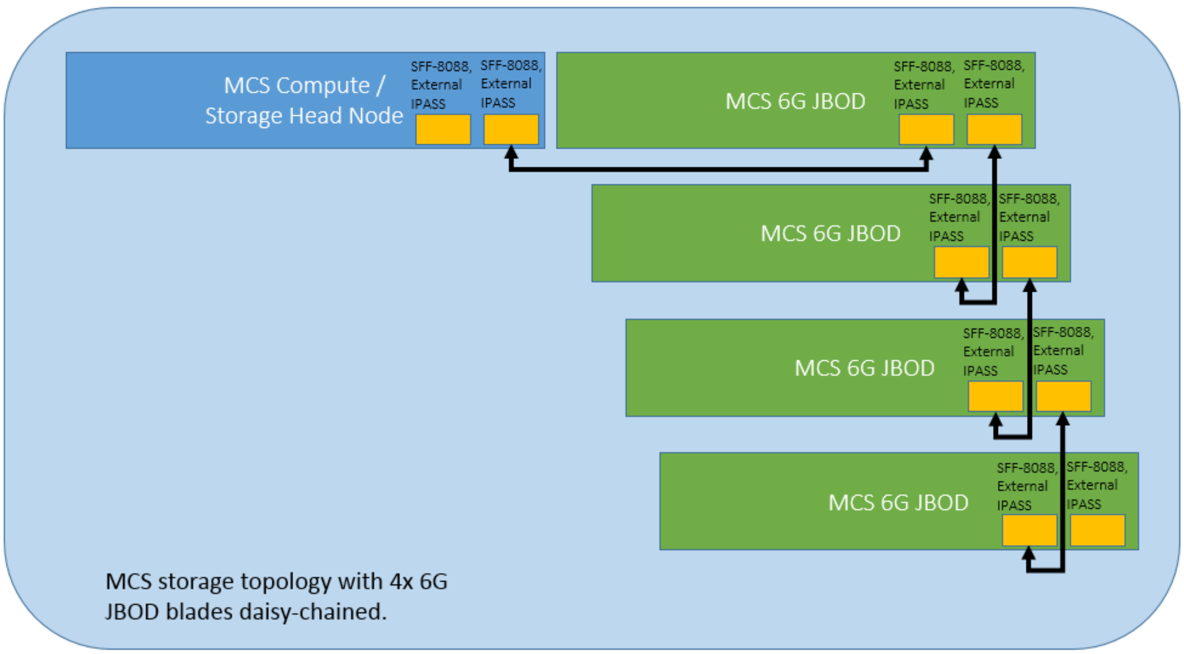

Figure 18. Cascading JBODs

11 Appendix: Commonly Used Acronyms

This section provides definitions of acronyms used in the system specifications.

ACPI – advanced configuration and power interface

AHCI – advanced host controller interface

AHJ – authority having jurisdiction

ANSI – American National Standards Institute

API – application programming interface

ASHRAE – American Society of Heating, Refrigerating and Air Conditioning Engineers

ASIC – application-specific integrated circuit

BCD – binary-coded decimal

BIOS – basic input/output system

BMC – baseboard management controller

CFM – cubic feet per minute (measure of volume flow rate)

CLI – command line interface

CM – backplane

CMOS – complementary metal–oxide–semiconductor

COLO – co-location

CTS – clear to send

DCMI – Data Center Manageability Interface

DDR3 – double data rate type 3

DHCP – dynamic host configuration protocol

DIMM – dual inline memory module

DPC - DIMMs per memory channel

DRAM – dynamic random access memory

DSR – data set ready

DTR – data terminal ready

ECC – error-correcting code

EEPROM - electrically erasable programmable read-only memory

EIA – Electronic Industries Alliance

EMC – electromagnetic compatibility

EMI – electromagnetic interference

FRU – field replaceable unit

FTP – file transfer protocol

GPIO – general purpose input output

GUID – globally unique identifier

HBA – host-bus-adapter controller

HBI – high business intelligence

HCK – Windows Hardware Certification Kit

HMD – hardware monitoring device

HT – hyperthreading

I^2C – inter-integrated circuit

IBC – international building code

IDE – integrated development environment

IEC - International Electrotechnical Commission

IOC – I/O controller

IPMI – intelligent platform management interface

IPsec – IP security

ITPAC – IT pre-assembled components

JBOD – "just a bunch of disks"

KCS – keyboard controller style

L2 – layer 2

LAN – local area network

LFF – large form factor

LPC – low pin count

LS – least significant

LUN – logical unit number

MAC – media access control

MDC – modular data center containers

MLC – multi-level call

MTBF – mean time between failures

MUX – multiplexer

NIC – network interface card

NUMA – non-uniform memory access

OOB – out of band

OSHA - Occupational Safety & Health Administration

OTS – off the shelf

PCB – printed circuit board

PCIe – peripheral component interconnect express

PCH – platform control hub

PDB – power distribution backplane

PDU – power distribution unit

Ph-ph – phase to phase

Ph-N – phase to neutral

PNP – plug and play

POST – power-on self-test

PSU – power supply unit

PWM – pulse-width modulation

PXE – preboot execution environment

QDR – quad data rate

QFN – quad flat package no-lead

QPI – Intel QuickPath Interconnect

QSFP – Quad small form-factor pluggable

RAID – redundant array of independent disks

REST - representational state transfer

RM – Rack Manager

RMA – remote management agent

ROC – RAID-on-chip controller

RSS – receive-side scaling

RTS – request to send

RU – rack unit

RxD – received data

SAS – serial-attached small computer system interface (SCSI)

SATA – serial AT attachment

SCK – serial clock

SCSI – small computer system interface

SDA – serial data signal

SDR – sensor data record

SEP – storage enclosure processor

SES – SCSI enclosure service

SFF – small form factor

SFP – small form-factor pluggable

SMBUS – systems management bus

SMBIOS – systems management BIOS

SOL – serial over LAN

SPI – serial peripheral interface

SSD – solid-state drive

TB – tray backplane

TDP – thermal design power

TB – tray backplane

TOR – top of rack

TPM – trusted platform module

TxD – transmit data

U – rack unit

UART – universal asynchronous receiver/transmitter

UEFI – unified extensible firmware interface

UL – Underwriters Laboratories

UPS – uninterrupted power supply

Vpp – voltage peak to peak

WMI –Windows Management Interface