SSI Ethernet Midplane Design Guide

November 2010

Revision 1.0.2

Disclaimer:
THIS DRAFT DESIGN GUIDE IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. WITHOUT LIMITATION, THE PROMOTERS (Intel Corporation, Dell Computer Corporation, Silicon Graphics, Inc., and International Business Machines Corporation) DISCLAIM ALL LIABILITY FOR COST OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS OF USE, LOSS OF DATA OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT, INDIRECT, OR SPECIAL DAMAGES, WHETHER UNDER CONTRACT, TORT, WARRANTY OR OTHERWISE, ARISING IN ANY WAY OUT OF USE OR RELIANCE UPON THIS DRAFT SPECIFICATION OR ANY INFORMATION HEREIN.

The Promoters disclaim all liability, including liability for infringement of any proprietary rights, relating to use of information in this Draft Specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

This specification and the information herein is the confidential and trade secret information of the Promoters. Use, reproduction and disclosure of this Draft Specification and the information herein is subject to the Terms and Conditions set forth on the website on which the download of this Draft Specification was made available.

Revision 1.0.2

Contents

Figures

Tables

Revision History

The following table lists the revision schedule based on revision number and development stage of the product.

Revision	Project Document State	Date
1.0.0	Initial public release	9/18/2009
1.0.1	Bug fixes: 268, 273	11/10/2010
1.0.2	Added notation for previous bug fixes: 321, 381	11/12/2010

Notes:

- Not all revisions may be published.

This page intentionally blank.

1 Midplane Overview

1.1 Introduction

The midplane (Figure 1-1) is a central feature of a blade system. It provides the interconnect topology between the compute blades, switches, management, power subsystems, and other key building blocks of the system.

The term “mid” implies there are elements on both sides of the midplane; however, this is only one possible implementation. There may be implementations where elements plug in to one side. (Such an implementation is commonly referred to as a backplane.) Throughout this document the term midplane will be used, but the concepts may equally apply to a backplane implementation.

Figure 1-1: Midplane Connectivity Diagram

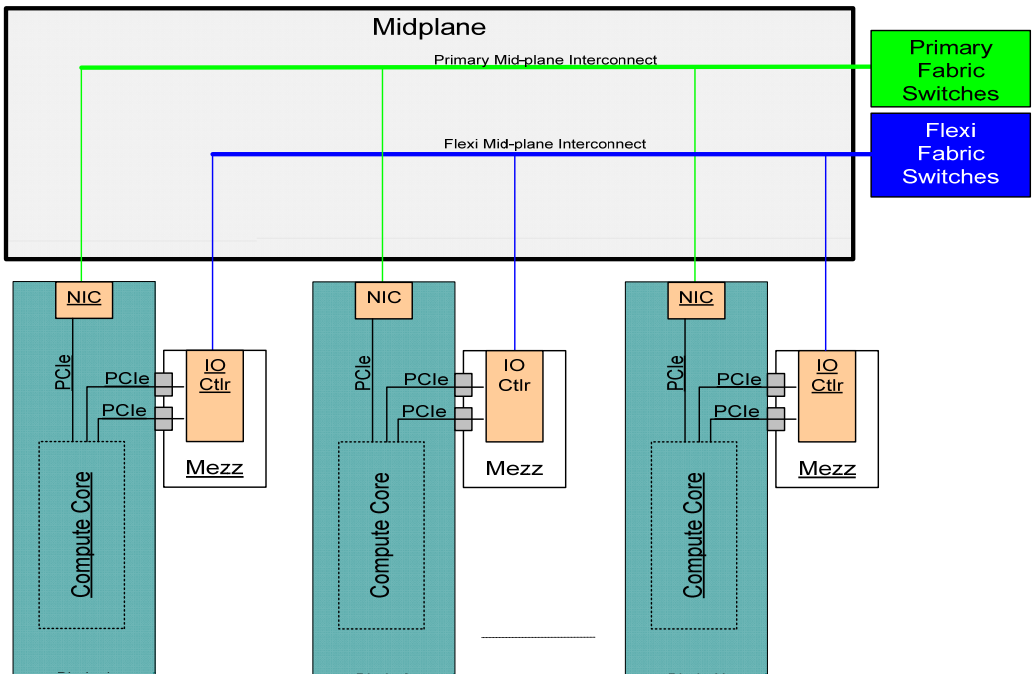

1.2 Purpose

This document is intended to give guidance for the printed circuit board (PCB) design of a midplane that supports SSI compute blades and switches, with emphasis on the electrical design of the high-speed fabrics. These fabrics include the following, specified in the Backplane Ethernet IEEE Std 803.3ap™-2007:

- 1000 BASE KX PMD
- 10GBASE-KX4 PMD
- 10GBASE-KR PMD

(Gigabit, 4-lane 10 Gigabit, and 1-lane Serial 10 Gigabit)

This guide is not limited to any specific implementation, and it does not guarantee that a specific implementation will or will not function with all SSI building blocks. It is up to the designer to verify the electrical characteristics and function of the design.

1.3 Reference Documents

- *SSI Compute Blade Specification*
- *SSI Compute Blade Mezzanine Specification*
  SSI Chassis Management Module (CMM) Specification
  SSI Ethernet Midplane Electric Specification
- *SSI Switch Base Specification (Base Specification for Switch Module Subsystmes)*
- *SSI Switch SERDES Specification (Base Specification and Design Guide for SERDES High-speed Electrical Signaling)*
- *SSI Switch VPD Specification (Base Specification for Vital Product Data (VPD))*
  IEEE Std 803.3ap™- 2007
  "Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment 4: Ethernet Operation over Electrical Backplanes". Copyright © 2007 by the Institute of Electrical and Electronics Engineers, Inc. All rights reserved. Published 22 May 2007. Printed in the United States of America. IEEE and 802 are registered trademarks in the U.S. Patent & Trademark Office, owned by the Institute of Electrical and Electronics Engineers, Incorporated.
- IPMI – Intelligent Platform Management Interface Specification, v2.0 rev 1.0E3, February 16, 2006, Copyright © 2004, 2005, 2006 Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., All rights reserved.
- IPMI – Platform Management FRU Information Storage Definition, V1.0, Document revision 1.1, September 27, 1999 Copyright © 1998, 1999 Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., All rights reserved.

- IPMI – Intelligent Platform Management Bus Communications Protocol Specification, V1.0, Document revision 1.0, November 15, 1999 Copyright © 1998, 1999 Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., All rights reserved.
- SD card specifications - http://www.sdcard.org/home/ [273]

1.4 Terms and Abbreviations

Table 1-1 lists terms and acronyms used in specific ways throughout this specification.

Table 1-1: Terms and Abbreviations

Term	Definition
ASHRAE	American Society of Heating, Refrigerating, and Air Conditioning Engineers.
Base Management Interface (BMI)	This is the IPMB-based management interface used by the Chassis Manager to communicate with the blade management controllers.
blade	This is a resource module that plugs into the blade chassis. A blade can provide many different types of resources to the chassis, including compute functions, storage capabilities, additional I/O interfaces and switching capabilities, and special purpose capabilities. A blade can be a single-wide module (assumed) or a double-wide module, occupying two adjacent slots in the chassis.
blade server	A system comprising a chassis, chassis resources (power, cooling, Chassis Manager), compute blades, and communication (switch) blades. The chassis may contain additional modules, such as storage.
bottom	When used in reference to a board, the end that would be on the bottom in a vertically oriented chassis.
CFM	Cubic Feet per Minute. A measure of volumetric airflow. One CFM is equivalent to 472 cubic centimeters per second.
chassis	The mechanical enclosure that consists of the mid-plane, front boards, cooling devices, power supplies, etc. The chassis provides the interface to boards, and it consists of the guide rails, alignment, handle interface, face plate mounting hardware, and mid-plane interface.
chassis ground	A safety ground and earth return that is connected to the chassis metal and available to all PBAs.
Chassis Management Module (CMM)	Dedicated intelligent chassis module that hosts the Chassis Manager functionality.
Chassis Manager (CM)	Set of logical functions for hardware management of the chassis. This may be implemented by one or more dedicated Chassis Management Modules or by one or more blade management controllers and/or payload processors.
cold start	Cold start is the time when blades receive the payload power for the first time.
component side 1	When used in reference to a PBA, the side on which the tallest electronic components would be mounted.

Term	Definition
component side 2	When used in reference to a PBA, the side normally reserved for making solder connections with through-hole components on Component Side 1, but on which low-height electronic components may also be mounted.
creepage	Surface distance required between two electrical components.
face plate	The front-most element of a PBA, perpendicular to the PBA, that serves to mount connectors, indicators, controls, and mezzanines.
guide rail	Provides for the front board guidance feature in a slot.
handle	An item or part used to insert or extract blades into and out of chassis.
Intelligent Platform Management Bus (IPMB)	IPMB is an I2C-based bus that provides a standardized interconnection between managed modules within a chassis. ftp://download.intel.com/design/servers/ipmi/ipmb1010ltd.pdf
Intelligent Platform Management Interface (IPMI)	IPMI v2.0 R1.0 specification defines a standardized, abstracted interface to the platform management subsystem of a computer system. ftp://download.intel.com/design/servers/ipmi/IPMIv2_0rev1_0.pdf
interconnect channel	An interconnect channel comprises two pairs of differential signals. One pair of differential signals for transmit and another pair of differential signals for receive.
LFM	Linear Feet per Minute. A measure of air velocity. One LFM is equivalent to 0.508 centimeters per second.
logic ground	Chassis-wide electrical net used on blades and mid-planes as a reference and return path for logic-level signals that are carried between boards.
managed module	Any component of the system that is addressable for management purposes via the specified management interconnect and protocol. A managed module is interfaced directly to the chassis BMI.
Management Controller	This is an intelligent, embedded microcontroller that provides management functionality for a blade or other chassis module.
may	Indicates flexibility of choice with no implied preference.
mezzanine	The mezzanine is a PBA that installs on a blade PBA horizontally. It provides additional functionality on the blade PBA and provides electrical interface between the blade PBA and the mid-plane PBA. Both the blade PBA and mezzanine PBA are contained inside the blade module.
mid-plane	Equivalent to a system backplane. This is a PBA that provides the common electrical interface for each blade in the chassis and on both the front and back of the PBA.
module	A physically separate chassis component which may be independently replaceable (e.g., a blade or cooling module) or attached to some other component (e.g., a mezzanine board).
open blade	A blade that conforms to the requirements defined by the Open Blade standard set of specifications.
out-of-band (OOB)	Communication between blades that does not need the host or payload to be powered on.
payload	The hardware on a blade that implements the main mission function of the blade. On a compute blade, this includes the main processors, memory, and I/O interfaces. The payload is powered separately from the blade management subsystem. Payload power is controlled by the blade management controller.

Term	Definition
PBA	Printed board assembly. A printed circuit board that has all electronic components attached to it.
PCB	Printed circuit board without components attached.
peak power	The maximum power a blade can draw for a very short period of time during a hot insertion, hot removal, or a cold start.
pitch line	Horizontal pitch line between slots.
shall	Indicates a mandatory requirement. Designers must implement such mandatory requirements to ensure interchangeability and to claim conformance with this specification. The use of **shall not** (in bold) indicates an action or implementation that is prohibited.
should	Indicates flexibility of choice with a strongly preferred implementation. The use of **should not** (in bold) indicates flexibility of choice with a strong preference that the choice or implementation be avoided.
slot	A slot defines the position of one blade in a chassis.
top	When used in reference to a blade, the end which would be on top in a vertically oriented chassis.
U	Unit of vertical height defined in IEC 60297-1 rack, shelf, and chassis height increments. 1U=44.45 mm.
WDT	Watchdog timer.

2 Midplane Mechanical Guidelines

This chapter describes the midplane mechanical guidelines.

2.1 PCB Guidelines

This section details the PCB guidelines.

2.1.1 Thickness

The midplane thickness is primarily driven by layer count, via stub length, and mechanical rigidity. A thickness of 125 mils is recommended. With this thickness, 16 layers are possible, and if careful, routing techniques via stub lengths can be minimized. If the PCB thickness is greater than 0.125" then advanced PCB technology, such as via back drilling or HDI-PCB* may be necessary to minimize via stub length.

2.1.2 Thermal Consideration

Each compute blade and switch requires airflow to cool components. Typically this air flow is provided by fans or blowers located in the rear of the chassis. Openings in the midplane must be made to allow sufficient airflow to and across the compute blades and switches. Factors to consider are placement of fans, and air flow directions for the compute blades, switches, and other components in the system. For airflow requirements of the compute blades and switches, see their respective specifications.

2.2 Connectors

This section shows the relative mechanical location of connectors on the midplane for the compute blade and switches (Figure 2-1 through Figure 2-4). It only shows the connector placement locations for a given compute blade or switch module, not the relative location from blade to blade or blade to switch. The mechanical placement for all compute blades and switches is up to the system designer. Table 2-1 lists the connectors specified for each module. For pin-out detail, see each individual building block specification.

Note: Pin-outs in each specification indicate signal direction with respect to that module. For example, a compute blade's RX is its receiver signal. The midplane **should** connect it to the appropriate TX of a switch.

* High Density Interconnect Printed Circuit Boards utilize blind via and micro-vias.

Table 2-1: Connector Reference Table

[Items in this table were updated by bug 268.]

Building block	Connector	Description	Manufacturer	Part Number
Compute Blade	96-pin primary fabric	4-pair vertical, 2 mm, 96-pin receptacle	FCI*	10084605-101LF
	96-pin optional connector	4-pair vertical, 2 mm, 96-pin receptacle	FCI	10084605-101LF
	Guide	10.8 mm guide pin	FCI	10084610-101LF
	Power	Power 2x2 vertical receptacle	FCI	10084602-4554P00LF
Mezzanine	96-pin Flexi-Fabric (Mezzanine)	4-pair vertical, 2 mm, 96-pin receptacle	FCI	10084605-101LF
Chassis Management Module	120-pin signal interface (2 instances)	4-pair vertical, 2 mm, 120-pin receptacle	FCI	10084601-101LF
	Guide	10.8 mm guide pin	FCI	10084610-101LF
Low-speed Switch Module (1x)	VHDM 1 of 2	6 row x 20 BP assembly	Molex*	74074-9987
			Amphenol*	498-5010-022
	VHDM 2 of 2	6 row x 20 BP assembly	Molex	74074-9976
			Amphenol	498-5110-022
	Power	6 row BP power assembly, 2 circuits	Molex	74029-6998
			Amphenol	437-6050-000
High-speed Switch Module (4x)	Guide pin right	Mating guide pin	Molex	75834-5205
			Amphenol	325-4020-025
	GbX end wall right	2 pair x 40 row signal conn	Molex	75834-5805
			Amphenol	326-4320-025
	Guide pin left	Mating guide pin	Molex	75834-3205
			Amphenol	325-4120-025
	GbX end wall left	2 pair x 40 row signal conn	Molex	75834-3805
			Amphenol	326-4420-025

Compute Blade Connector Location

Figure 2-1 illustrates the locations of compute blade connectors on the midplane.

Figure 2-1: Compute Blade Connector Layout on Midplane

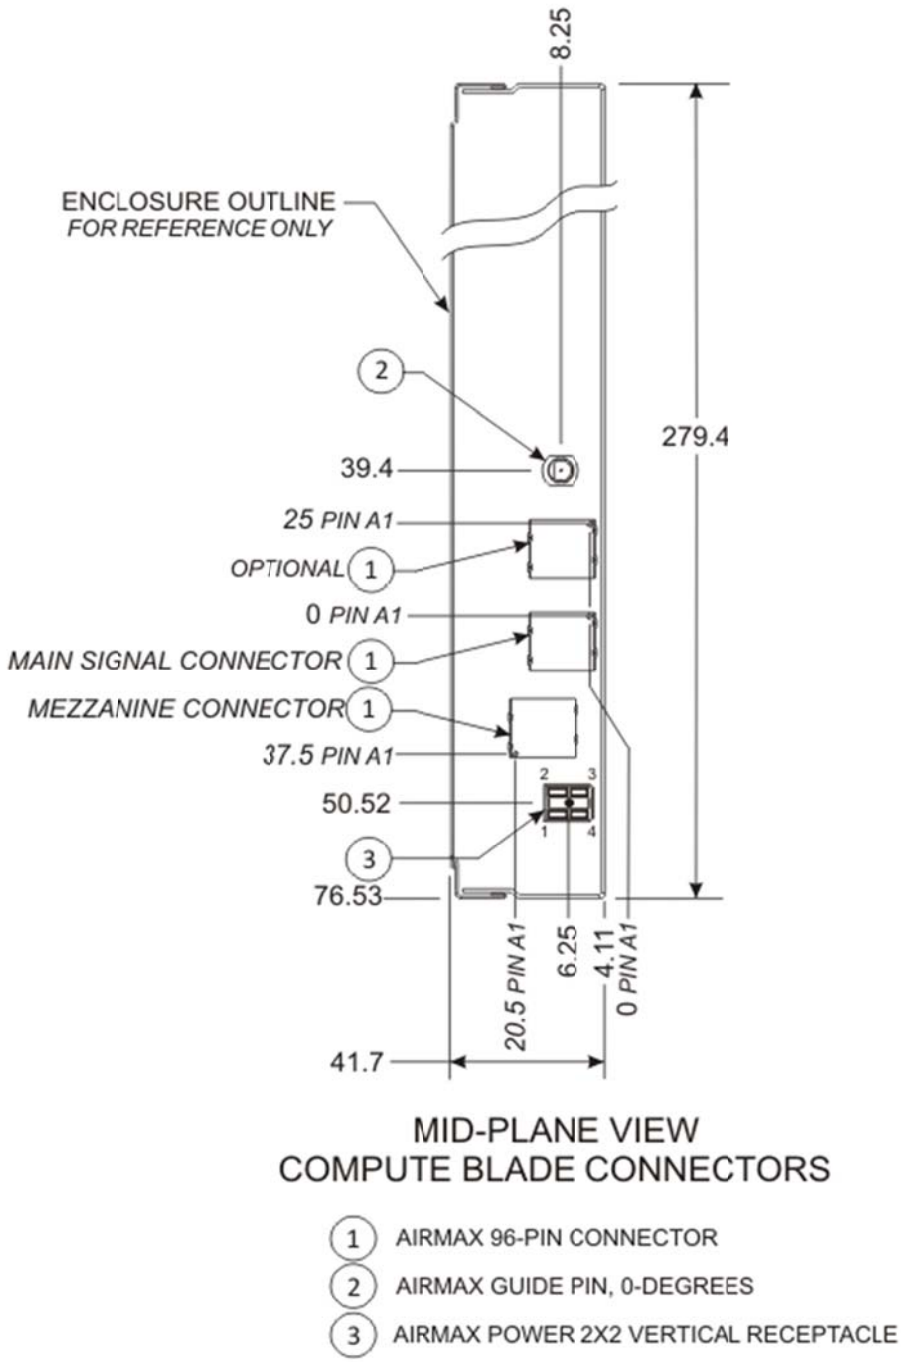

2.2.1 Chassis Management Module Connector Location

Figure 2-2 illustrates the locations of CMM connectors on the midplane [381].

Figure 2-2: CMM Connector Layout on Midplane

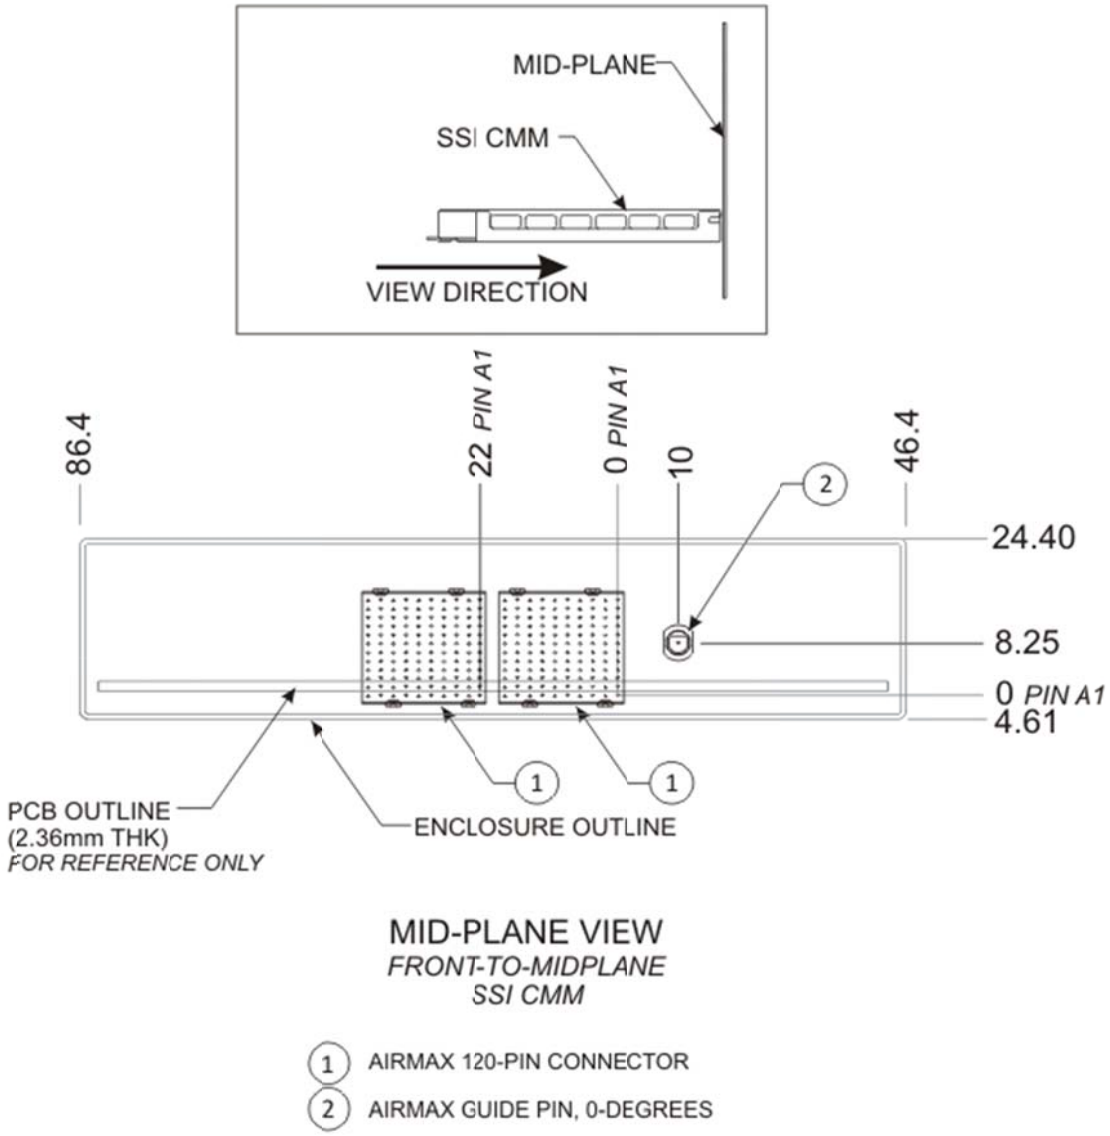

2.2.2 Switch Connector Location

2.2.2.1 SSI 1X Switch Module (LSSM)

Figure 2-3: 1X (LSSM) Connector Layout on Midplane

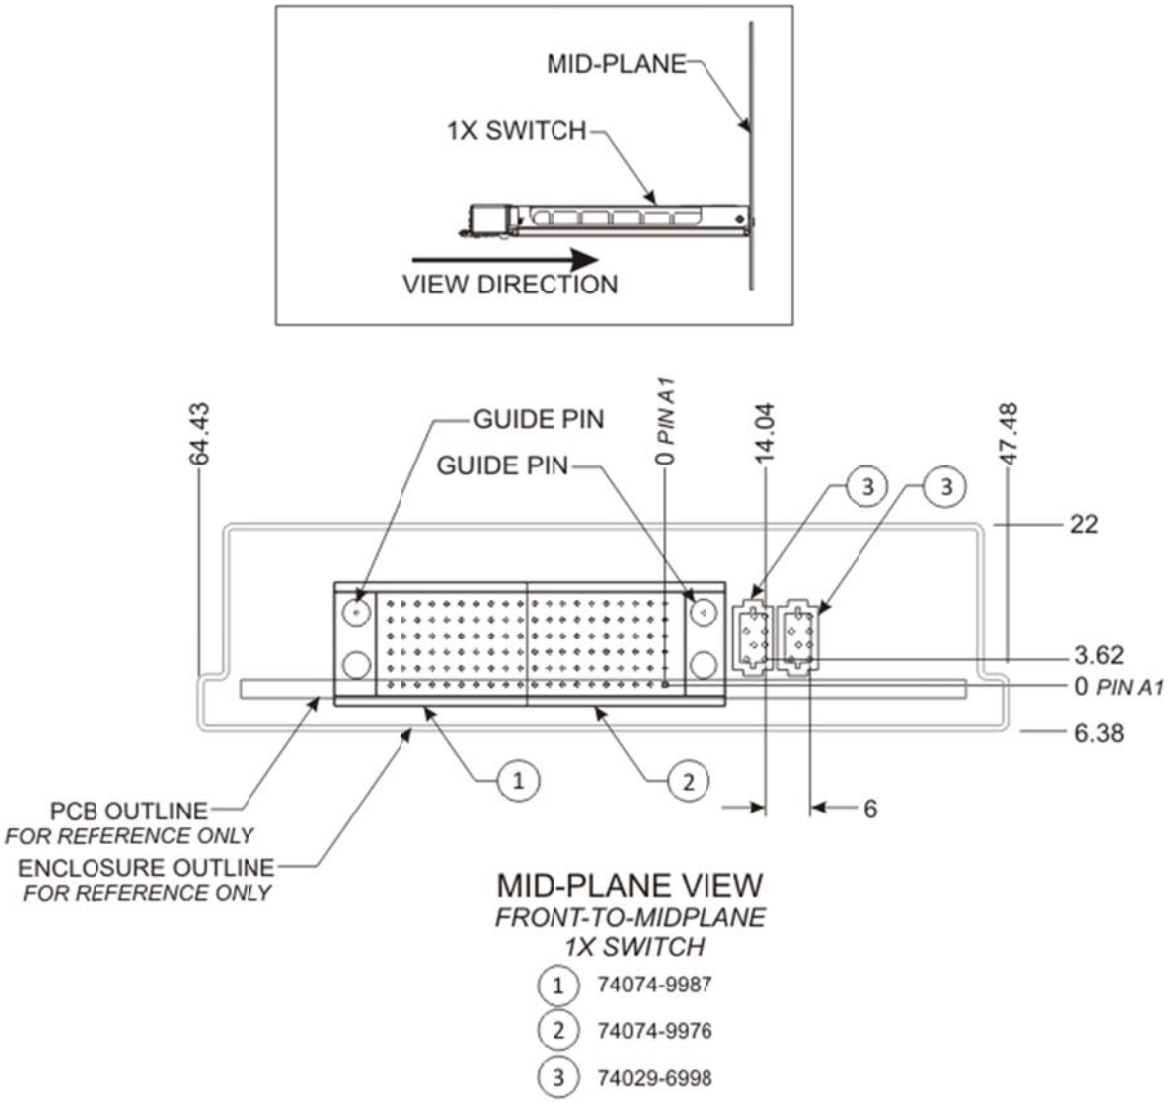

2.2.2.2 SSI 4X Switch Module (HSSM)

Figure 2-4: 4X (HSSM) Connector Layout on Midplane

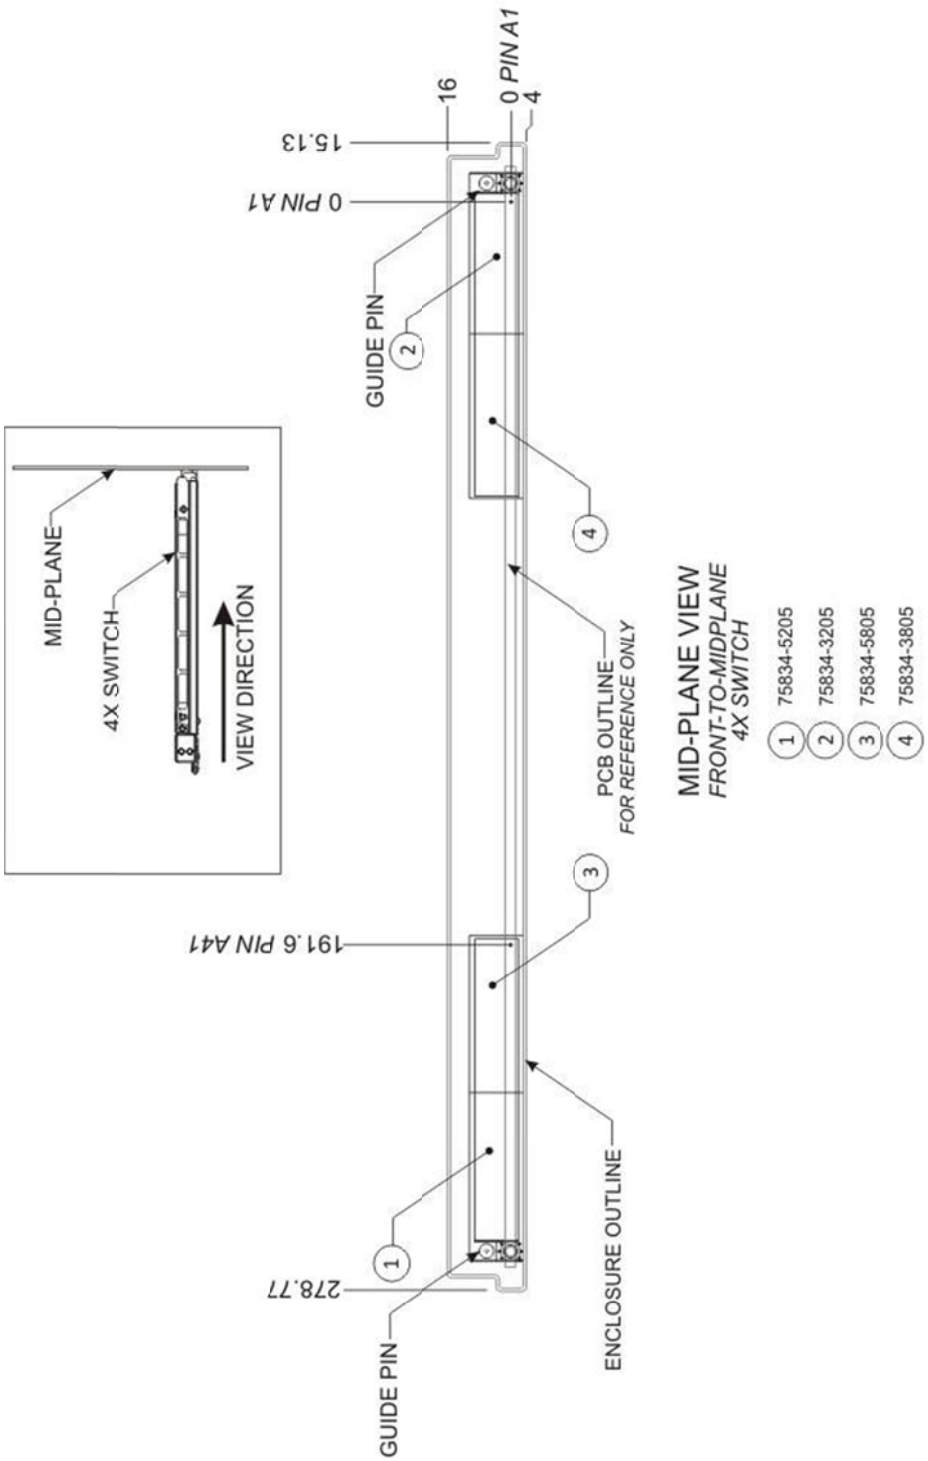

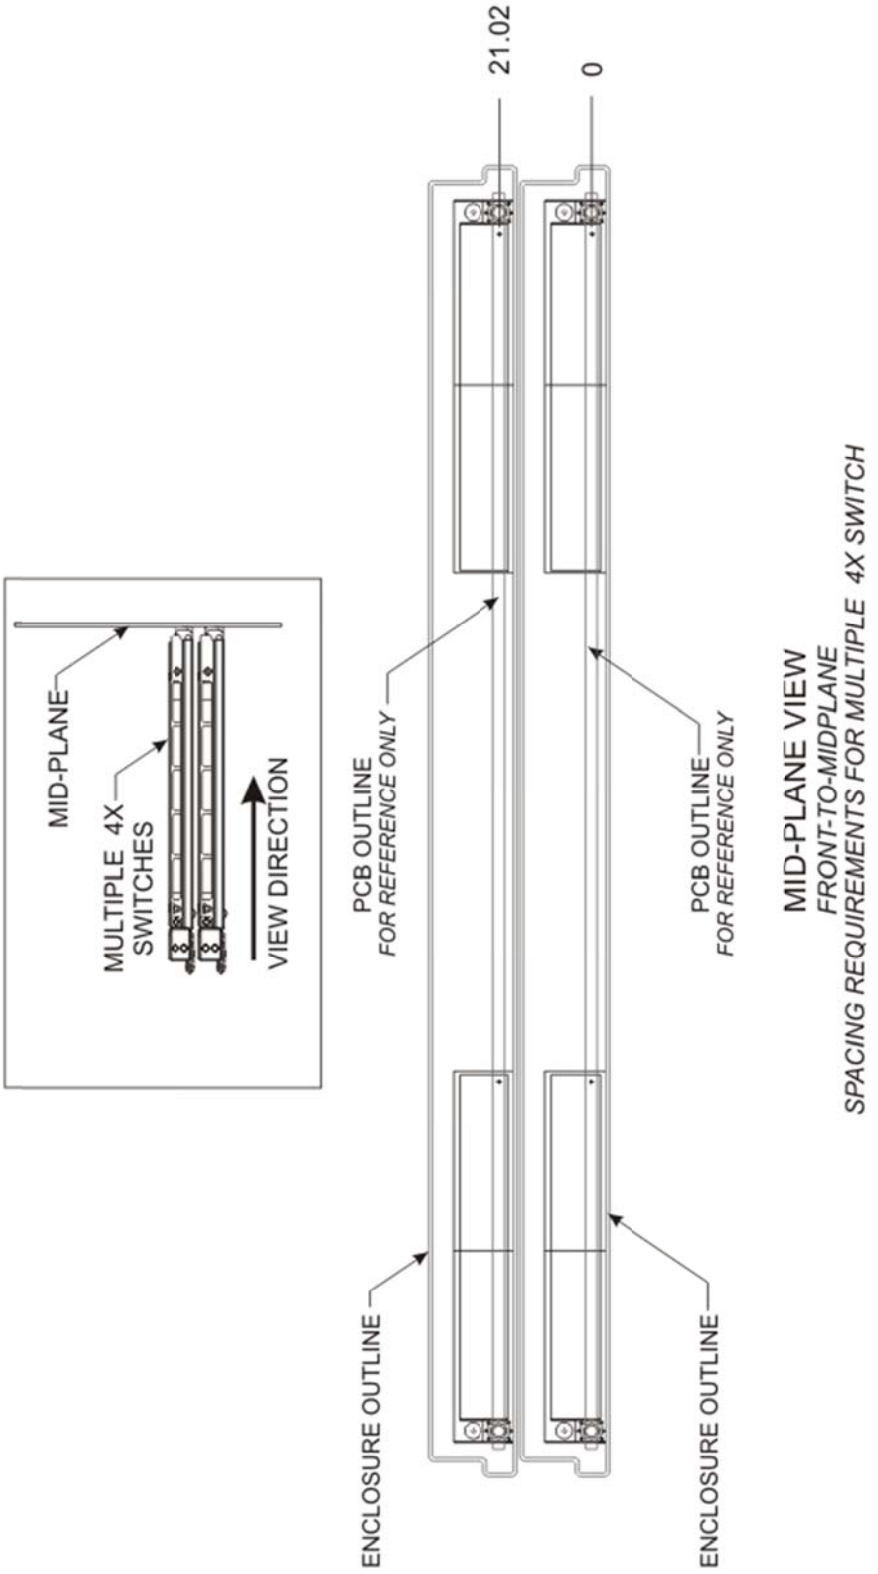

MID-PLANE VIEW
FRONT-TO-MIDPLANE
SPACING REQUIREMENTS FOR MULTIPLE 4X SWITCH

3 Electrical Guidelines

3.1 Stackup

A 16-layer stackup example, with trace width, spacing, and minimum isolation distance between the differential pairs, is shown in Table 3-1.

Table 3-1: 16-layer Stackup Example

	total thickness *125.2* 16 layers	single-ended		DIFFERENTIAL			
Layer Name	Thickness	isolation spacing	width	isolation spacing	width	pair space	
Solder Mask	0.5						
Top	2						
prepreg, 2-ply	6						
Plane2	1.3						
core	7	30	8.5	30	7	8	mils
Sig3	1.3		42.5		85.0		Target Z0
prepreg	7						
Plane4	1.3						
core	7	30	8.5	30	7	8	mils
Sig5	1.3		42.5		85.0		Target Z0
prepreg	7						
Plane6	1.3						
core	7	30	8.5	30	7	8	mils
Sig7	1.3		42.5		85.0		Target Z0
prepreg	7						
Plane8	1.3						
core	6						
Plane9	1.3						
prepreg	7						
Sig10	1.3		42.5		85.0		Target Z0
core	7	30	8.5	30	7	8	mils
Plane11	1.3						
prepreg	7						
Sig12	1.3		42.5		85.0		Target Z0
core	7	30	8.5	30	7	8	mils
Plane13	1.3						
prepreg	7						
Sig14	1.3		42.5		85.0		Target Z0
core	7	30	8.5	30	7	8	mils
Plane15	1.3						
prepreg, 2-ply	6						
Botm	2						
Solder Mask	0.5						

3.1.1 Power Distribution

The 12 V power **should** be delivered through a low-impedance path. In the stackup example, it's routed on the top and bottom layers utilizing 2 oz. copper.

In the stackup example, a solid ground plane is on layer 8. Power rails other than 12 V could be routed on layer 9. If this is done, layer 10 **should not** be used for high-speed differential pairs.

3.2 Ethernet Fabric Signal Routing

All high-speed signals are required to meet the *SSI Midplane Electrical Specification*. This design guide provides guidance in design techniques for Ethernet technology in meeting the *SSI Midplane Electrical Specification*.

Three different PCB technology cost points are considered to determine maximum trace lengths achievable. These are shown in Table 3-2.

Table 3-2: PCB Technology Options

	PCB Material	Via Stub Length	Differential Z0 tolerance
Low-cost (LC)	FR4 Tan d < 0.018	S-VIA, VIA stub < 48 mils	10%
Medium-cost (MC)	FR408 Tan d < 0.014	S-VIA VIA stub < 48 mils	10%
High-cost (HC)	Nelco 4000 13-SI Tan d < 0.008	S-VIA, VIA stub < 48 mils	8%

Maximum high-speed trace length on the midplane varies, given the switch connector's performance characteristics combined with various cost points of PCB technologies and length budgets on the compute blade and switches. Table 3-3 lists some midplane route lengths with these constraints in mind [321].

Table 3-3: Midplane Route Length

[This table was updated by bug 321.]

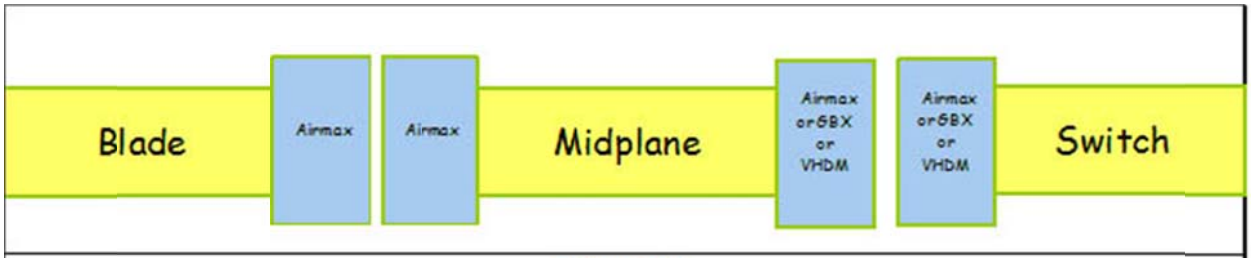

Channel										
Blade			Midplane					Switch		
Channel Technology	Impedance Target	Connector Type		Design Point LC	Design Point MC	Design Point HC	Impedance Target	Connector Type	Impedance Target	Channel Technology
KX	85	Airmax	Max Len	16.5	19.5	22.5	85	VHDM	100	KX
			Rev	0.8	0.8	0.8				
KX4	85	Airmax	Max Len	10	11.5	13	85	GBX	100	KX4
			Rev	0.8	0.8	0.8				
KR	85	Airmax	Max Len	9	11	13	85	GBX	100	KR
			Rev	0.8	0.8	0.8				

Rev Definitions
- 0.0 Engineering targets to meet timings and noise margins
- 0.5 Simulations complete pending Blade, Switch or Midplane Specification changes
- 0.8 Simulations complete without outstanding specification changes
- 1.0 Simulations correlated with hardware
- 2.0 Demonstrated in Compatibility and Interoperability Lab

Table 3-4 shows mandatory PCB layout rules for all materials.

Table 3-4: Routing Techniques

		Requirement	NOTES
Layer assignment	See Figure 3-1	Mandatory for S-via boards	Other via stub reduction methods can replace this
Trace width target	> 7 mil	Mandatory	
Isolation space between differential pairs (coupling)	> 4 times the dielectric height	Mandatory	
Corners	Mitered corners are recommended	Not mandatory	
Min. length	2 inches	Mandatory	
Max. length	See Table 3-3	Mandatory	
Length matching	2 mils between lines	Mandatory	
Connector pair matching	Compensate for the delay mismatch in connector pairs	Mandatory	See connector specification from connector manufacturers.
Pair trace length match	< 5 mils	Mandatory	
Impedance	85 ohm target	Recommended	See Table 3-2 for tolerances
Construction	Ground referenced symmetric strip line	Mandatory	Micro strip is not acceptable
Copper	Reverse treat copper	Not mandatory for KX Mandatory for KR and IB QDR	Intent is to reduce loss
Via anti pad to via hole	> 3 mils hole diameter is desirable	Not mandatory	Intent is to minimize via stub impairment
Via stub	< 35 mils	Mandatory and restricted	See Table 3-2 Maximum high-speed trace length on the midplane varies, given the switch connector's performance characteristics combined with various cost points of PCB technologies and length budgets on the compute blade and switches. Table 3-3 lists some midplane route lengths with these constraints in mind [321].

			for cost options
Non-functional via pads	Unused or non-functional via pads should be removed	Mandatory	This is usually an option in layout tools.
PCB material	See Table 3-2 for cost options	Mandatory	
Board thickness	< 0.125 mils	Not mandatory	Thicker boards require additional via mitigation than suggested in Table 3-2 Maximum high-speed trace length on the midplane varies, given the switch connector's performance characteristics combined with various cost points of PCB technologies and length budgets on the compute blade and switches. Table 3-3 lists some midplane route lengths with these constraints in mind [321].

3.3 General Signal Routing Guidelines

This section describes the general signal routing guidelines.

3.3.1 S-VIA Routing technique

This section describes a method to reduce via stub known as S-Via Routing, as shown in Figure 3-1. This method is used when the compute blade and switches are on opposite sides of the midplane.

Figure 3-1: S-Via Routing and Layer Assignment

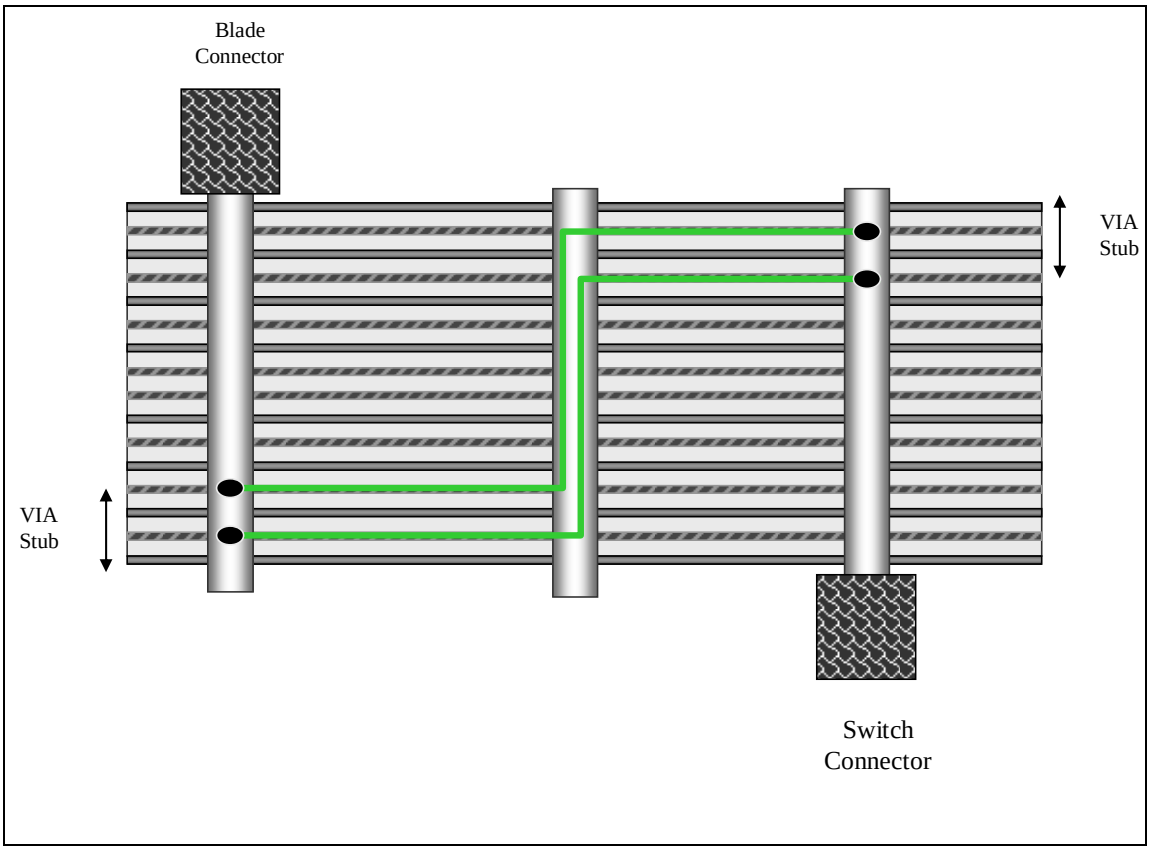

Utilizing the stackup shown in Table 3-1, depending on PCB stackup, this method limits via stub length to meeting the requirement in Table 3-2.

3.3.2 Length Matching and Connector Break Out

Serpentine routing, as indicated in Figure 3-2, **shall** be utilized to match intra-pair length mismatches. Also, the connector break out is illustrated.

Figure 3-2: Serpentine Trace Length Compensation

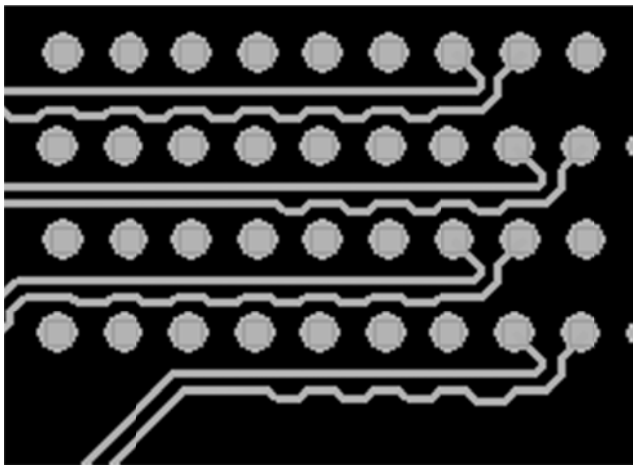

Serpentine routing is a good way to compensate corner length mismatch by using back-to-back turns as shown in Figure 3-3.

Figure 3-3: Back-to-Back Corner Compensation

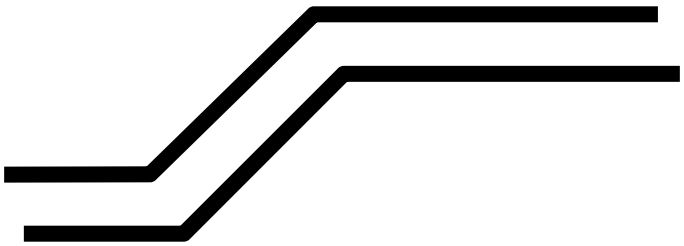

3.3.3 Test Points

No additional vias **shall** be added for test points.

3.3.4 Void, Splits, and Proximal Metal (Mandatory)

Caution: It is required to NOT cross splits or voids in adjacent reference planes.

Traces that run parallel to a split are not permitted. Either side of a differential trace **shall** be greater than 4 times the dielectric height linear distance away from any metal or void in adjacent reference planes.

Exception: An exception is the connector break out.

3.3.5 Connectors

The midplane connectors are specified in the mechanical section. In general, the frequency rating on a connector is not sufficient for connector selection.

Midplane routing **shall** compensate for the delay mismatch in connector pairs.

3.3.6 Dielectric Weave Compensation

It is required that traces do not run parallel to the dielectric FR4 weave. The recommended compensation method is to rotate the artwork 12 to 15 degrees as shown in Figure 3-4. Artwork rotation is the recommended method to mitigate effects of fiber weave. Refer to IPC 2008 APEX publication “Image Rotation to Mitigate the Fiber Weave Effect — Its Impact on PCB Manufacturing.”

Figure 3-4: Dielectric Weave Compensation

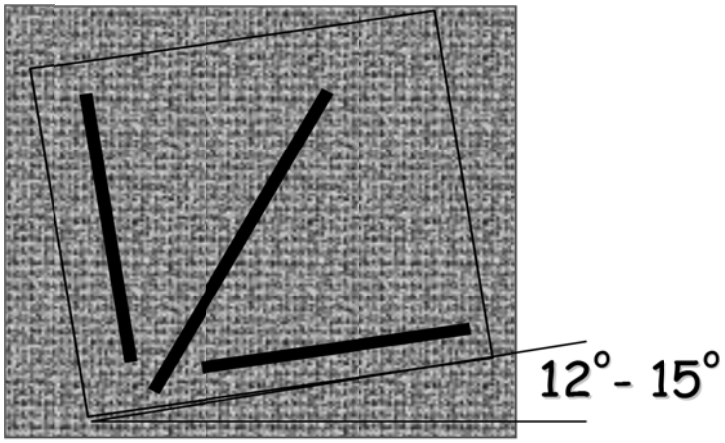

3.3.7 Trace Spacing

Maintain constant differential trace spacing as shown Figure 3-5.

Figure 3-5: Trace Spacing

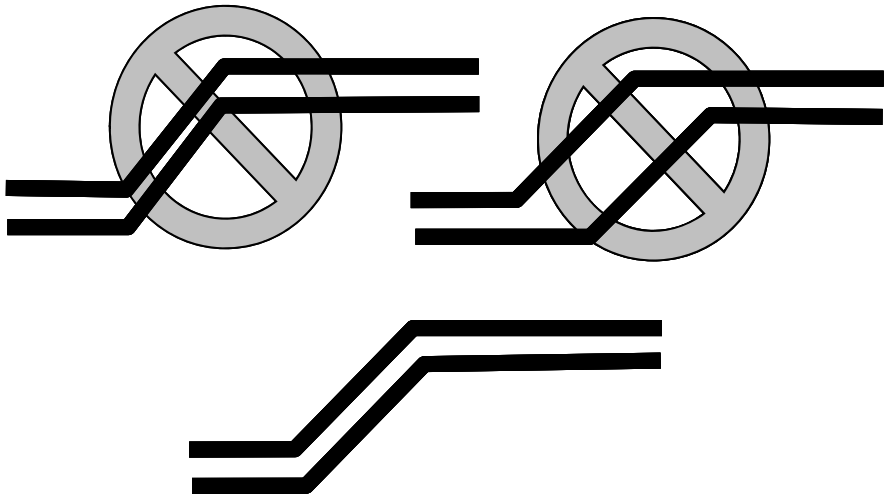

3.3.8 Vias

Traces **shall not** have anti pad overlap as shown in Figure 3-6.

Figure 3-6: Anti-via Overlap

There **shall** be no signal vias placed between a differential via pair as shown in Figure 3-7.

Figure 3-7: No Signal Vias Between Differential Vias

A differential via design is recommended for connector and s-via areas, as shown in Figure 3-8. In general, holes within a pair placement **should** be as close as manufacturing constraints permit, and hole size **should** be as small as manufacturing requirements allow. In an area of the board that is unconstrained by other nearby traces, the anti via **should** be three times the diameter of the hole. In other areas, the anti pad can be as large as manufacturing and nearby routing constraints allow. Conversely, the pad **should** be as small as possible. In general, anti pads are larger than pads.

Figure 3-8: Differential Vias

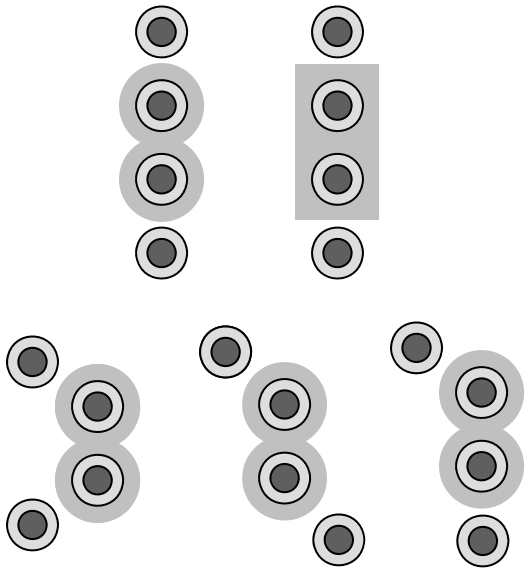

An example of balanced differential via routing is shown in Figure 3-9. An alternative three-via routing pattern is shown in Figure 3-10 with a ground via between the differential pair.

Figure 3-9: Differential Via Routing

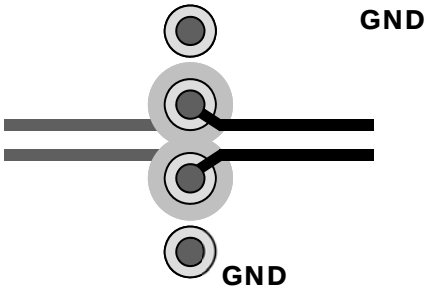

Figure 3-10: Alternative Differential Via Routing

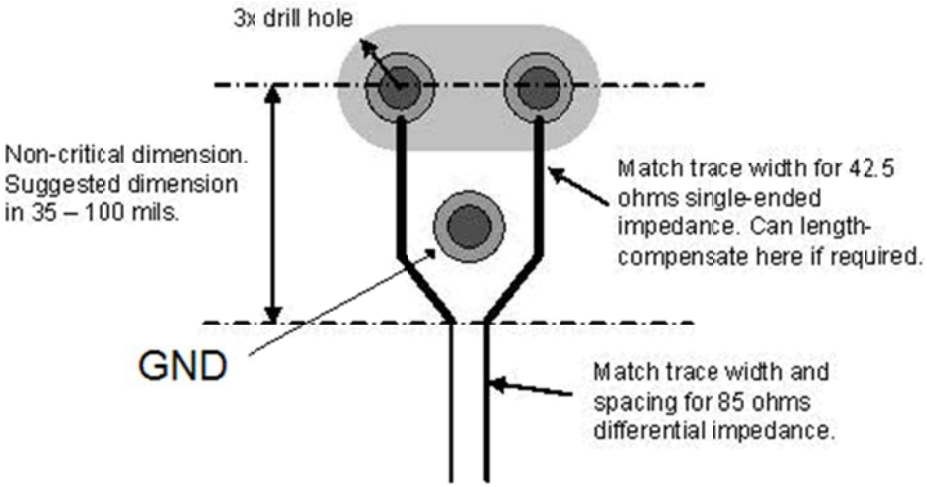

3.3.9 Guard traces

Guard traces are not recommended as shown in Figure 3-11.

Figure 3-11: No Guard Traces

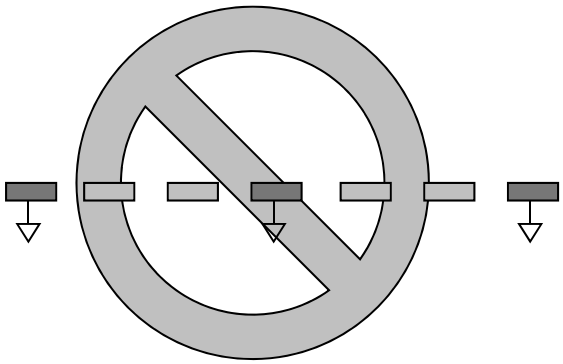

3.3.10 Signal Referencing

All SERDES traces **shall** be referenced to ground, as shown in Figure 3-12. SERDES traces **shall not** be referenced to power. If striplines are employed, use only symmetric striplines, not asymmetric striplines.

Figure 3-12: Signal Referencing

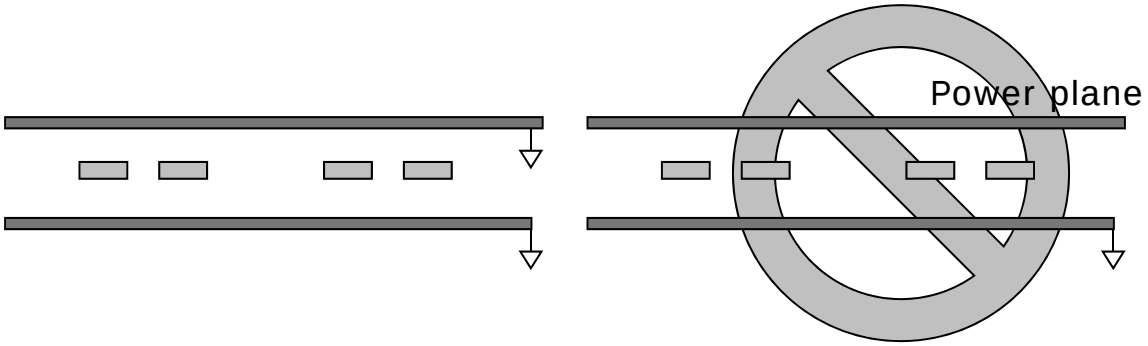

3.3.11 Corners

Use Figure 3-13 as a guide to trace corners routing.

Figure 3-13: Mitered Corners

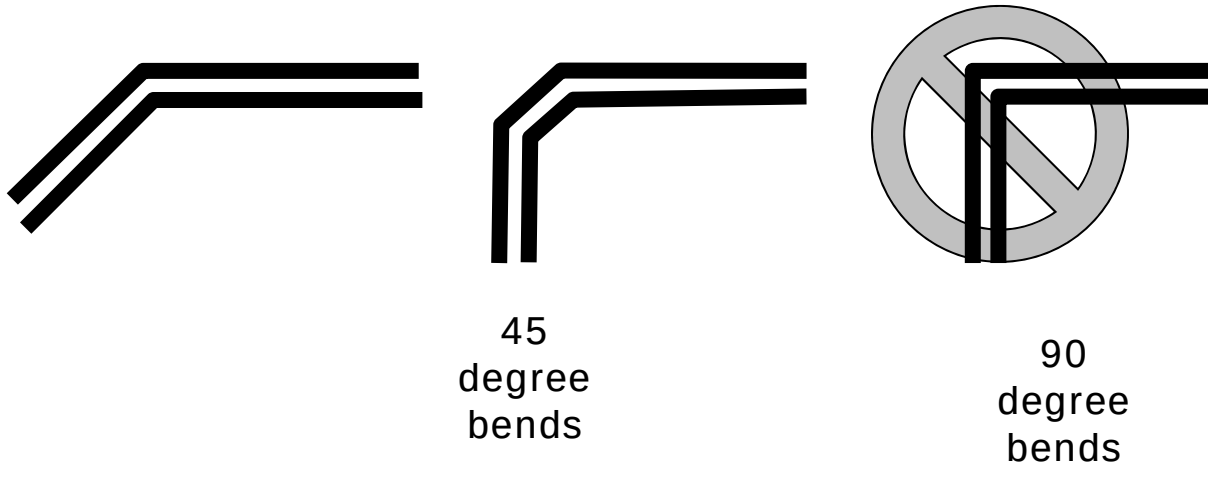

3.3.12 Topologies

All high-speed fabrics **shall** utilize single-layer routing unless s-vias are employed. In that case, two-layer routing is permitted. However, only one via **shall** be used.

3.3.13 Length Compensation

Trace length mismatch through any connector going to the midplane **should** be compensated on the midplane itself, not on the compute blade or switch modules. Length compensation **shall** assume a mated connector. Refer to the connector manufacturer specifications for specific routing length.

3.3.14 Reverse Copper Treatment

At higher frequencies (2GHz and higher) skin depth of PCB traces approach that of the copper foil roughness. Reverse copper treatment minimizes trace roughness and is required for traces used for high frequency signaling.

3.4 Low-speed Signals

This section describes low-speed signals.

3.4.1 BMI Topologies

BMI interconnects utilize I2C technology. It is recommended that each compute blade and switch have dedicated I2C busses, such that there is point-to-point routing on the midplane. See the *SSI Chassis Management Module (CMM) Specification* for more information concerning available I2C busses.

Each I2C bus **shall** be pulled up to 5V via a 2.2k ohm resistor on the midplane. This requirement for 5V on the midplane may necessitate a 5V regulator on the midplane. I2C busses going to the switches **shall** be pulled up to 5.0V via a 2.2k ohm resister on the midplane.

3.4.2 MM_SELECT_A/B_N Signals

The CMM drives a MM_SELECT_A/B_N signal to the switches for I2C master selection. These signals require a pull up resistor on the midplane to 5V. They also require an inverter on the midplane to change the logic state going into the switches from active low to active high. See Figure 3-14 for details.

Careful loading analysis **should** be made by the designer to ensure that the MM_SELECT_A/B_N signals are not over loaded.

Figure 3-14: MM_SELECT_A/B Circuit

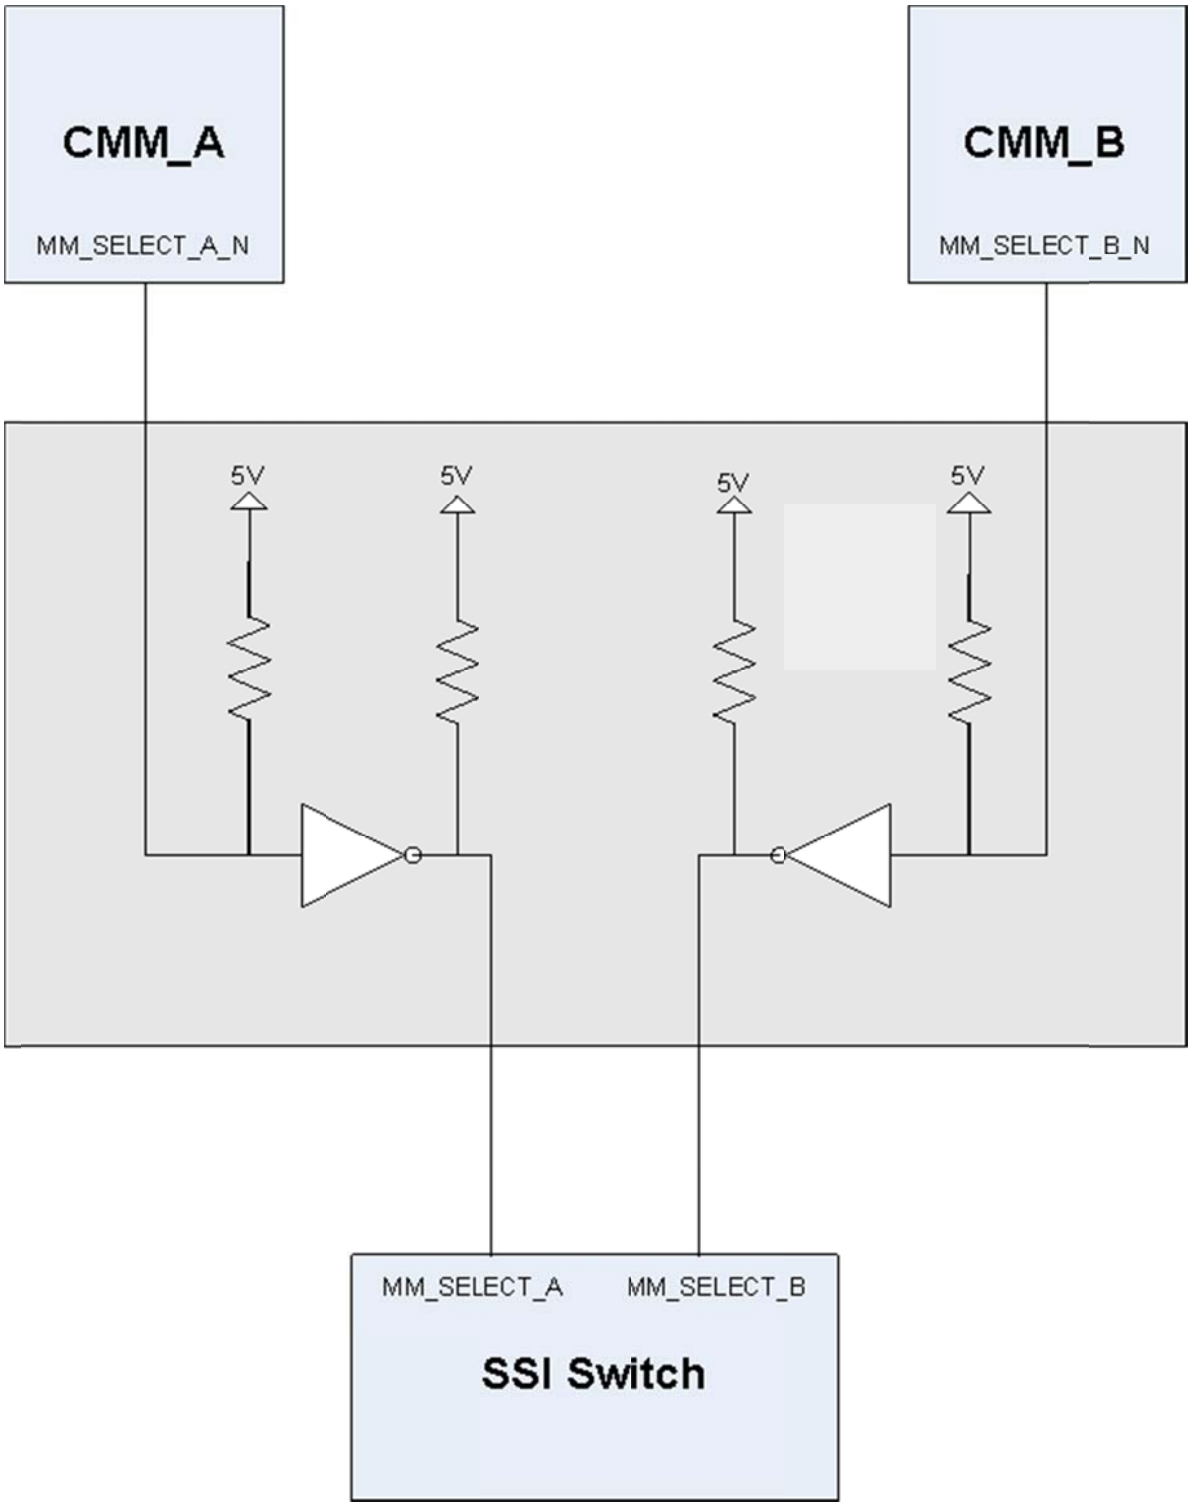

4 Other Considerations

4.1 Architecture

When architecting a backplane, consideration must be given to what type of switch a given fabric will support. There are different switch form factors, which are not interchangeable. For example, the LSSM supports only x1 fabric, and it is a different form factor than the HSSM, which supports x4 fabric.

A given Compute Blade implementation may only implement a x1 Ethernet fabric, either KX or KR. If it is plugged into a fabric with an HSSM Switch Module, three channels per blade will go unused. On the other hand if a Compute Blade implements a x4 fabric and is plugged into a fabric with an LSSM Switch Module, only a single lane from the Compute Blade will be utilized.

Another consideration is how a technology mismatch between a Compute Blade and Switch Module will auto-negotiate. The auto-negotiation feature of the 802.3ap specification will negotiate to the highest common performance link possible as shown in Table 4-1.

Table 4-1: Primary Link Negotiation

Blade NIC Technology	Switch Technology	Negotiated Link
KX	KX	KX
	KX4	KX
	KR	KX
KX4	KX	KX
	KX4	KX4
	KR	KX
KR	KX	KX
	KX4	KX
	KR	KR

4.2 Airflow

For a given system design, careful analysis of airflow is required. Openings in the midplane will be required to allow proper flow to cool the various modules. Furthermore, airflow requirements of the various modules are not uniform. For example, the compute blade requires front to back airflow where as the switch module allows for several options.

Openings made in a mid-plane blocks routing channels, so a trade off between routing requirements and airflow will need to be made by the system designer.

4.3 CMM SD Storage Card

The CMM specification provides a mechanism for a centralized SD storage device on the midplane of a chassis. This can optionally be used to keep the CMM software on the chassis itself. This could be used to store databases only or the entire CMM software stack. When designing a midplane, support **should** always be made for this form of storage by providing an SD mounting device and properly routing it to the CMM slot.

The recommended implementation is a SD storage device per CMM. Alternatively, a single SD storage device can be shared between redundant CMMs utilizing a multiplexer and the signals CMM_SELECT_A_N and CMM_SELECT_B_N.

To ensure a reliable communication to the SD card, it is recommended that a means is implemented to reset the SD interface using the FM_SDCD_FLSH_PW-RST signal as shown in Figure 4-1.

A pnp transistor is required for the FM_SDCD_FLSH_PW-RST signal to operate properly in the operational state.

A 100k to 150k ohm resister from the SD DAT3 signal is needed to ensure software can detect the presence of an SD storage card.

To ensure signal integrity, the SD Card connector **shall** be placed within 5” of the CMM connector on the mid-plane.

Refer to the following for more information:

- SD card specifications - http://www.sdcard.org/home/ [273]

Figure 4-1: Example CMM SD Interface Circuit

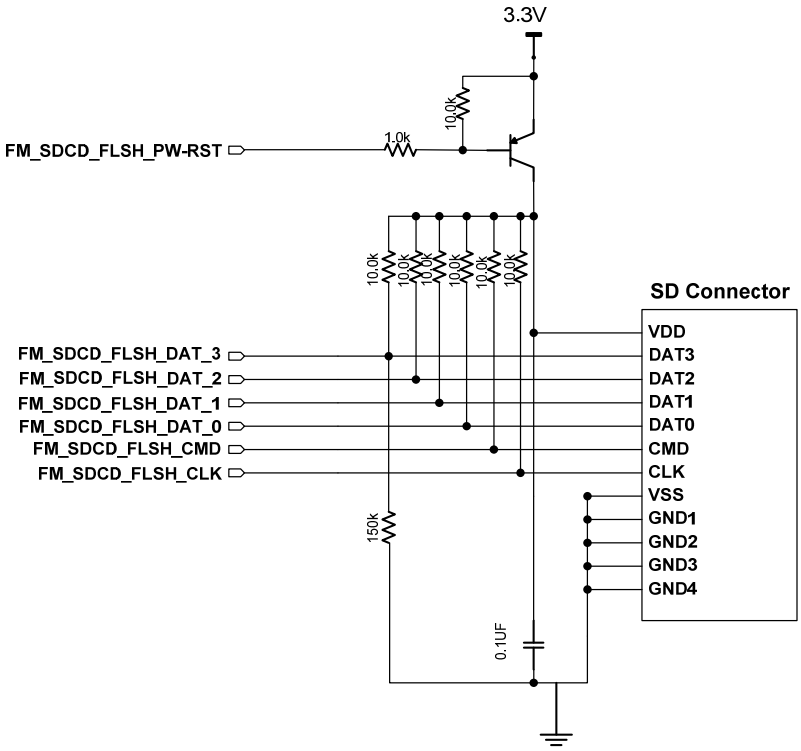

4.4 Midplane IDROM Device

The midplane **shall** provide a 5V tolerant I2C Serial EEPROM of type 24C64 (or compatible) that is capable of 100 kHz operation at a minimum, and **shall** be inter-operable with 400 KHz devices. This device will store chassis configuration information and **shall** be connected to CMM SMB_SPARE I2C Bus.

See the SSI Chassis Management Module (CMM) Specification for details.