COM-HPC® Carrier Design Guide

Guidelines for Designing COM-HPC® Carrier Boards

February 16, 2024

Rev. 2.2

This Design Guide is not a specification. It provides COM-HPC® Carrier implementation information but does not replace the PICMG COM-HPC® specification.

The full COM-HPC® specification is needed in conjunction with this Design Guide for signal descriptions, signal integrity information and loss budgets, Module and Carrier connector pinouts, PCB mechanical details and more.

Table of Contents

Index of Tables

Table of Figures

1.   Preface

1.1.   About This Document

This document provides information for designing project specific Carrier Boards for systems using COM-HPC Modules. This document is a design guide and not a specification document. It should be used by together with the **COM-HPC Base Specification**, with other industry specifications (listed in Section 1.9. below), with silicon and component vendor's documentation and with your COM-HPC Module vendor's product documentation.

The PICMG **COM Express Carrier Board Design Guide** is also a very useful additional source of information. The COM-HPC and COM Express pin names are not the same, but it is not hard to correlate them. The COM Express design guide document is available for free download on the public PICMG website (www.picmg.org). No membership is required to down load the design guides.

Revision 2.2 of this document includes numerous notes that clarify how COM-HPC Client examples need to be modified slightly for use in COM-HPC Mini designs. The differences between the COM-HPC Client and the Mini are well described in the COM-HPC Module Base Specification Rev. 1.20, primarily in Section 10 of that document.

The notes that clarify the COM-HPC Mini differences are presented in this bold italic blue font to make the notes stand out.

1.2.   Intended Audience

This design guide is intended for electronics engineers and PCB layout engineers designing Carrier Boards for PICMG COM-HPC Modules. It may also be useful to COM-HPC Module designers for them to better understand how COM-HPC Modules are used, and to understand how some of the design rules (trace length recommendations, trace length matching recommendations etc.) are shared between Module and Carrier designs.

1.3.   No Special Word Usage

Unlike a PICMG specification, which assigns special meanings to certain words such as "shall", "should" and "may", there is no such usage in this document. That is because this document is not a specification; it is a non-normative design guide.

1.4.   No Statements of Compliance

As this document is not a specification but a set of guidelines, there should not be any statements of compliance made with reference to this document.

1.5.   Correctness Disclaimer

The schematic examples given in this document are believed to be correct but no guarantee is given.

1.6.   Name and Logo Usage

The PCI Industrial Computer Manufacturers Group policies regarding the use of its logos and trademarks are as follows:

Permission to use the PICMG® organization logo is automatically granted to designated members only as stipulated on the most recent Membership Privileges document (available at www.picmg.org) during the period of time for which their membership dues are paid. Nonmembers may not use the PICMG® organization logo.

The PICMG® organization logo must be printed in black or color as shown in the files available for download from the member's side of the Web site. Logos with or without the "Open Modular Computing Specifications" banner can be used. Nothing may be added or deleted from the PICMG® logo.

The use of the COM-HPC® logo is a privilege granted by the PICMG® organization to companies who have purchased the relevant specifications (or acquired them as a member benefit), and that believe their products comply with these specifications. Manufacturers' distributors and sales representatives may use the COM-HPC® logo in promoting products sold under the name of the manufacturer. Use of the logos by either members or non-members implies such compliance. Only PICMG Executive and Associate members may use the PICMG® logo. PICMG® may revoke permission to use logos if they are misused. The COM-HPC® logo can be found on the PICMG web site, www.picmg.org.

The COM-HPC® logo must be used exactly as shown in the files available for download from the PICMG® Web site. The aspect ratios of the logos must be maintained, but the sizes may be varied. Nothing may be added to or deleted from the COM-HPC® logo.

The PICMG® name and logo and the COM-HPC® name and logo are registered trademarks of PICMG®. Registered trademarks must be followed by the ® symbol, and the following statement must appear in all published literature and advertising material in which the logo appears:

PICMG, the COM-HPC® name and logo and the PICMG logo are registered trademarks of the PCI Industrial Computers Manufacturers Group.

1.7.   Intellectual Property

The PICMG Consortium draws attention to the fact that it is claimed that compliance with this specification may involve the use of a patent claim(s) ("IPR"). The PICMG Consortium takes no position concerning the evidence, validity or scope of this IPR.

The holder of this IPR has assured the Consortium that it is willing to license or sublicense all such IPR to those licensees (Members and non-Members alike) desiring to implement this specification. The statement of the holder of this IPR to such effect has been filed with the Consortium.

Attention is also drawn to the possibility that some of the elements of this specification may be the subject of IPR other than those identified below. The Consortium shall not be responsible for identifying any or all such IPR.

No representation is made as to the availability of any license rights for use of any IPR inherent in this specification for any purpose other than to implement this specification.

This specification conforms to the current PICMG® Intellectual Property Rights Policy and the Policies and Procedures for Specification Development and does not contain any known intellectual property that is not available for licensing under Reasonable and Non-discriminatory terms. In the course of Membership Review the following disclosures were made:

1.7.1.   Necessary IPR Claims (Referring to Mandatory or Recommended Features)

Samtec Inc. has the following patents, which may cover some aspects of the PICMG® COM-HPC® Module and Carrier Board Connectors. Contact Samtec Inc. at royalties@samtec.com for further information.

China   201480061913.2
China   202111274151.2
China   201921051845.8
China   202030159171.5
EPO   007814686-0001
EPO   007814686-0002
EPO   007814686-0003
EPO   007814686-0004
EPO   19830502.1
Taiwan 109138672
Taiwan M589915
Taiwan D209464
Taiwan 109304816
Taiwan 505072
US   29/70951
US   9374900

1.7.2.   Unnecessary Claims (Referring to Optional Features or Non-normative Elements)

US   9374900
CN   201480061913.2
TWM   505072

PCT   2021207390
TW   110112769

CN   11566924
US   17/817659

1.7.3.   Third Party Disclosures

(Note that third party IPR submissions do not contain any claim of willingness to license the IPR.)

- US 10,404,014 B2, FCI USA LLC, Sep. 3, 2019, "STACKING ELECTRICAL CONNECTOR WITH REDUCED CROSS TALK"

Refer to PICMG® IPR Policies and Procedures and the company owner of the patent for terms and conditions of usage.

PICMG® makes no judgment as to the validity of these claims or the licensing terms offered by the claimants.

THIS SPECIFICATION IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NON-INFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS SPECIFICATION SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE CONSORTIUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS SPECIFICATION.

Compliance with this specification does not absolve manufacturers of COM-HPC® equipment from the requirements of safety and regulatory agencies (UL, CSA, FCC, IEC, etc.).

PICMG®, CompactPCI®, AdvancedTCA®, ATCA®, AdvancedMC®, CompactPCI® Express, COM Express®, COM-HPC®, MicroTCA®, SHB Express®, and the PICMG, CompactPCI, AdvancedTCA, µTCA and ATCA logos are registered trademarks, and xTCA™, IRTM™ and the IRTM logo are trademarks of the PCI Industrial Computer Manufacturers Group. All other brand or product names may be trademarks or registered trademarks of their respective holders.

1.7.4.   Copyright Notice

PICMG has made every attempt to ensure that the information in this document is accurate yet the information contained within is supplied "as-is".

1.7.5.   Trademarks

Intel is a registered trademark of Intel Corporation. PCI Express is a registered trademark of Peripheral Component Interconnect Special Interest Group (PCI-SIG). COM-HPC® is a registered trademark of PCI Industrial Computer Manufacturers Group (PICMG). I2C is a registered trademark of NXP Semiconductors. Samtec is a registered trademark of Samtec Inc. All product names and logos referenced in this document are property of their owners. Thunderbolt is a registered trademark of the Intel corporation. Accelerate is a trademark of Samtec Inc.

1.8. Acronyms, Abbreviations and Definitions Used

Table 1: Acronyms, Abbreviations and Definitions Used

Term	Definition
10GBASE-KR	10 Gbps internal copper interface. Operates over a single lane and uses the same physical layer coding (defined in IEEE 802.3 Clause 49) as 10GBASE-LR (Single Mode Fiber 1310 nm) / ER (Single Mode Fiber 1550 nm) /SR (Multi Mode Fiber 850 nm)
25GBASE-KR	25 Gb/s internal copper interface using 25GBASE-R encoding over one lane in each direction
AC Coupled	This term means that series capacitors are inserted in the differential pair lines. This allows the transmit and receive lines to have their own, possibly separate DC common mode voltages.
ACPI	Advanced Configuration Power Interface
AMOLED	Active Matrix Organic (semiconductor) Light Emitting Diode (a flat panel display technology)
ARM	Advanced RISC Machine – a low power alternative CPU architecture widely used in mobile and embedded systems
ATX	Advanced Technology Extended – Industry standard PC Motherboard form factor and power supply definitions
BIDIR	Bidirectional (in reference to electrical signals)
BIOS	Basic Input Output System
BMC	Baseboard Management Controller or Board Management Controller – located on Carrier for COM-HPC, if implemented
Carrier Board	An application specific circuit board that accepts a COM-HPC Module
CCC	Current Carrying Capability
DDI	Digital Display Interface – an interface that can serve DisplayPort and HDMI/DVI,
DIMM	Dual In-line Memory Module – larger format SDRAM memory module used in desk top systems and server PCs
DisplayPort DP	DisplayPort is a digital display interface standard put forth by the Video Electronics Standards Association (VESA). It defines a new license free, royalty free, digital audio/video interconnect, intended to be used primarily between a computer and its display monitor.
DPLL	Digital Phase Locked Loop
DRAM	Dynamic Random Access Memory
DVI	Digital Visual Interface – a Digital Display Working Group (DDWG) standard that defines a standard video interface supporting both digital and analog video signals. The digital signals use TMDS.
EAPI	Embedded Application Programming Interface
EC	Embedded Controller
ECN	Engineering Change Notice
EEPROM	Electrically Erasable Programmable Read-Only Memory
Embedded DisplayPort eDP	Embedded DisplayPort (eDP) is a digital display interface standard defined by the Video Electronics Standards Association (VESA) for digital interconnect of Audio and Video within a closed system such as a laptop computer or a piece of laboratory instrumentation.
ESD	Electro Static Discharge
eSPI	Enhanced Serial Peripheral Interface
FAE	Field Application Engineer
FET	Field Effect Transistor – available as N-Channel (nFET) or P-Channel (pFET) devices
Flash	EEPROM memory used for code storage. It can be updated in place ("flashed").
FPGA	Field Programmable Gate Array
FR4	A type of fiber-glass laminate commonly used for printed circuit boards.
Gb	Gigabit
GbE	Gigabit Ethernet
Gbps	Gigabits per second
GPI	General Purpose Input
GPS	Global Positioning System
GPIO	General Purpose Input Output
GPO	General Purpose Output

Term	Definition
GPU	Graphics Processing Unit
Gtps	Giga Transfers per Second
HDA	Intel High Definition Audio (HD Audio) refers to the specification released by Intel in 2004 for delivering high definition audio.
HDMI	High Definition Multimedia Interface – digital display interface widely used in consumer electronics such as digital TVs
I2C	Inter Integrated Circuit – 2 wire (clock and data) signaling scheme allowing communication between integrated circuits, primarily used to read and load register values.
I2S	Inter IC Sound – a 5 wire serial data interface, used primarily for transmitting and receiving digital audio data
I3C	Improved Inter Integrated Circuit – builds on I2C and offers higher speeds and in-band interrupts
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IPR	Intellectual Property Rights
LAN	Local Area Network
Legacy Device	Relics from the PC-AT computer that are not in use in contemporary PC systems: primarily the ISA bus, UART-based serial port, parallel printer port, PS-2 keyboard, and mouse. Definitions vary as to what constitutes a legacy device. Some definitions include IDE as a legacy device.
LPC	Low Pin-Count Interface: a low speed interface used for peripheral circuits such as Super I/O controllers, which typically combine legacy-device support into a single IC.
LS	Least Significant
M.2	A small form factor add in card, for storage, WiFi, Cell Modems, etc. Interface options include PCIe x1, x2 or x4, SATA, USB and asynchronous serial. The standard is maintained by the PCI-SIG.
MAC	Media Access Control – in this document, MAC refers to the physical hardware bridge device between a CPU interface such as PCIe, and a network interface such as MDI or 10GBASE-KR or many others. A PHY is needed between the MAC and the Ethernet physical layer
MAFS	Term for Master Attached Flash Sharing where the Flash component is attached to the processor interface.
MDI	Media Dependent Interface between an Ethernet PHY and the system magnetics and copper twisted pairs
MDIO	Management Data Input/Output, or MDIO, is a 2-wire serial bus that is used to manage PHYs or physical layer devices in media access controllers (MACs).
ME	Management Engine – Intel term for a management microcontroller embedded into the chipset silicon. It is active before the main x86 CPU boots.
MIPI	Industry trade group that sets standards for mobile devices
MMC	Module Management Controller – a small microcontroller on the Module that works in conjunction with a Carrier BMC to implement IPMI functions. Implementation is optional.
MS	Most Significant
NA or N/A	Not Available, Not Applicable
NBASE-KR	Ethernet back plane signaling on PCB differential pairs. 'N' signifies the speed – 25Gbps or 10Gbps
NBASE-T	Ethernet physical layer signaling on twisted pairs. 'N' signifies the speed – 10Gbps, 5Gbps, 2.5Gbps, 1Gbps, 100Mbps or 10Mbps
NC	No Connect
NDA	Non-Disclosure Agreement
Nyquist Frequency	The critical frequency, sometimes called the "folding frequency", for a digital sampling system. It is (usually) ½ of the maximum data rate for the system.
NVME	Non Volatile Memory Express – non volatile memory with a PCIe interface – x1, x2 or x4 – often in an M.2 card form factor
OCXO	Oven Controlled Xtal (crystal) Oscillator
OEM	Original Equipment Manufacturer
OTP	One Time Programmable – an option offered by some silicon vendors to change IC parameters by programming or blowing device fuses once, before shipment.
PC-AT	"Personal Computer – Advanced Technology" – an IBM trademark term used to refer to Intel x86 based personal computers in the 1990s
PCB	Printed Circuit Board

Term	Definition
PCI	Peripheral Component Interconnect
PCI Express PCIe	Peripheral Component Interconnect Express – serialized point-to-point version of PCI
PEG	PCI Express Graphics
PHY	Physical layer device, usually used in the context of
PMD	Physical Medium Dependent – the physical layer of computer network protocols
POR	Plan of Record
PPS	Pulse Per Second (for Ethernet)
PTP	Precision Time Protocol (for Ethernet)
PU PD	Pull Up – A connection between a signal and a specified power rail, through a resistor Pull Down
R_a	Roughness Average – a measure of surface roughness, expressed in units of length.
ROM	Read Only Memory – a legacy term – often the device referred to as a ROM can actually be written to, in a special mode. Such writable ROMs are sometimes called Flash ROMs. BIOS is stored in ROM or Flash ROM.
RSVD	Reserved. If a pin is marked RSVD, nothing should be connected to it
RTC	Real Time Clock – battery backed circuit in PC-AT systems that keeps system time and date
S0, S1, S2, S3, S4, S5	System states describing the power and activity level S0 Full power, all devices powered S1 CPU powered, CPU and bus clocks off, not in common use S2 Not used S3 Suspend to RAM System context stored in RAM; only standby power rail present S4 Suspend to Disk System context stored on disk; only standby power rail present S5 Soft Off Main power rail off, only standby power rail present
SAFS	Term for Slave Attached Flash Sharing where the Flash component is attached behind a BMC component.
SATA	Serial Advanced Technology Attachment – serial interface standard for hard disks
SDP	Software Definable Pin
SKU	Stock Control Unit (a part number for a specific stockable item)
SGMII	Serial Gigabit Media Independent Interface
SMA	Sub Miniature type A – a small form factor circular connector used for miniature coax cables, for WiFi, GPS and Cell Modem antennas
SMBus	System Management Bus – a 3 wire bus – clock, data and alert – based in I2C – for system management
SOC	System On Chip
SO-DIMM	Small Outline Dual In-line Memory Module – small form factor SDRAM module
SPD	Serial Presence Detect – refers to serial EEPROM on DRAMs that has DRAM Module configuration information
SPI	Serial Peripheral Interface
Super I/O	An integrated circuit, typically interfaced via the LPC or eSPI bus that provides legacy PC I/O functions including PS2 keyboard and mouse ports, serial and parallel port(s) and a floppy interface.
TFT	Thin Film Transistor – refers to technology used in active matrix flat-panel displays, in which there is one thin film transistor per display pixel.
TMDS	Transition Minimized Differential Signaling – a digital signaling protocol between the graphics subsystem and display. TMDS is used for the DVI digital signals.
TPM	Trusted Platform Module, chip to enhance the security features of a computer system.
USB	Universal Serial Bus
Vgs	Gate to Source voltage to turn a FET on
WDT	Watch Dog Timer.
XAUI	10 Gbps Attachment Unit Interface.
XGMII	10 Gbps Media Independent Interface
XO	Xtal (crystal) Oscillator

1.9.   Applicable Documents and Standards

The list below is a partial list of documents and standards applicable to COM-HPC®. Many of the standards groups listed below (MIPI, PCI-SIG, USB, VESA etc.) have much more additional information available – ECNs, supplemental documents, test specifications, SI masks etc. These are too numerous to list here. Please explore the links below for additional documents that may be relevant.

- Advanced Configuration and Power Interface (ACPI) Specification Version 6.3, January 2019, Copyright © 2018, Unified Extensible Firmware Interface (UEFI) Forum, Inc. All rights reserved. https://uefi.org/specifications

- ATX Specification Version 2.2 © Intel Corp. 2004

- ATX12V Power Supply Design Guide, Version 2.2, March 2005 © Intel Corp.

- eSPI Enhanced Serial Peripheral Interface, Interface Base Specification Revision 1.0, Copyright © 2016, Intel Corporation. January 2016 https://downloadcenter.intel.com/download/27055/

- HDMI (High Definition Multimedia Interface) specifications. http://www.hdmi.org
  - High-Definition Multimedia Interface specification versions 1.3, 1.4b, 2.1
  - HDMI Alt Mode USB Type-C
- I2C Specifification
  - NXP UM10204 "I^{2}C-bus specification and user manual"
  - Rev 7 October 1, 2021
  - http://www.nxp.com use NXP site search tool to locate UM10204

- IEEE standards http://www.ieee.org
  - IEEE Std 802.3™-2018 (Revision of IEEE Std 802.3-2015), IEEE 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
  - IEEE1588 – 2008. IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, July 24, 2008, Copyright 2016

- Intelligent Platform Management Interface Specification Second Generation, v2.0, Document Revision 1.1, October 1, 2013 (c) Intel, Hewlett-Packard, NEC, Dell

  - An E7 red-line markup version of this document, dated April 21 2015, is available – see https://www.intel.com/content/www/us/en/servers/ipmi/ipmi-technical-resources.html

- Intel Low Pin Count (LPC) Interface Specification Revision 1.1, August 2002 Copyright © 2002 Intel Corporation. All rights reserved. https://www.intel.com/content/www/us/en/design/technologies-and-topics/low-pin-count-interface-specification.html

- MIPI Alliance specifications https://www.mipi.org
  - MIPI-CSI-2   Camera Serial Interface
  - MIPI-CSI-3   Camera Serial Interface
  - MIPI DSI   Display Serial Interface
  - MIPI DSI-2   Display Serial Interface
  - MIPI C-PHY   Physical layer spec for CSI-2 and DSI-2 (alternative)
  - MIPI D-PHY   Physical layer spec for CSI-2 and DSI-2
  - MIPI M-PHY   Physical layer spec for CSI-3
  - MIPI SoundWire   Serialized audio interface
  - MIPI I3C   Two wire serial data interface, successor to I2C

- MXM Graphics Module Mobile PCI Express Module Electromechanical Specification Version 3.0 Revision 1.1 (c) 2009 Nvidia Corporation Note: this document is not publicly available at the time of this writing but it does exist

- NEBS (Network Equipment – Building Systems)
  This is a collection of documents describing reliability criteria for telecom equipment. The NEBs documents are maintained by Telcordia / Ericsson
  https://telecom-info.telcordia.com

- PCI-SIG (Peripheral Component Interconnect Special Interest Group) specifications https://www.pcisig.com
  - PCI Express Base Specification Revision 5.0
  - PCI Express Card Electromechanical Specification Revision 4.0
  - PCI Express Mini Card Electromechanical Specification Revision 2.1
  - Add USB 3.0 to the Mini Card
  - PCI Express M.2 Specification Revision 4.0 V1.0
  - PCI Local Bus Specification Revisions 3,4 and 5.

- PICMG (PCI Industrial Computer Manufacturing Group) documents http://www.picmg.org/
  - PICMG COM.0 COM Express Module Base Specification Revision 3.0
  - PICMG EAPI Embedded Application Software Interface Specification Revision 1.0
  - PICMG EEEP Embedded EEPROM Specification Revision 1.0 (for COM-Express)
  - PICMG COM-HPC EEEP Embedded EEPROM Specification Revision 1.0
  - PICMG COM-HPC Carrier Design Guide Revision 2.2 (this document0
  - PICMG COM-HPC Base Specification Revision 1.2
  - PICMG COM Express Carrier Design Guide Revision 2.0
  - PICMG COM-HPC Platform Management Interface Specification Revision 1.0
  - PICMG Policies and Procedures for Specification Development, Revision 2.0

- Serial ATA Revision 3.5a Specification (March 2021) http://www.sata-io.org/

- SFP+, SFF-8083 Rev 3.1, SFF-8083 Specification for SFP+ 1X 10 Gb/s Pluggable Transceiver Solution (SFP10)
  Rev 3.1, Sep. 13, 2014 ftp://ftp.seagate.com/sff/SFF-8083.PDF

- SGET (Standardization Group for Embedded Technologies) standards and documents (www.sget.org)
  - SMARC Hardware Specification Revision 2.1.1 (Smart Mobility ARChitecture)
  - SMARC Design Guide Revision 2.1.1

- SPI, Serial Peripheral Interface Bus
  See http://elm-chan.org/docs/spi_e.html for some general information on SPI

- System Management Bus (SMBus) Specification Version 2.0, August 3, 2000 Copyright © 1994, 1995, 1998, 2000 Duracell, Inc., Energizer Power Systems, Inc., Fujitsu, Ltd., Intel Corporation, Linear Technology Inc., Maxim Integrated Products, Mitsubishi Electric Semiconductor Company, PowerSmart, Inc., Toshiba Battery Co. Ltd., Unitrode Corporation, USAR Systems, Inc. see http://www.smbus.org

- Trusted Computing Group Specifications https://www.trustedcomputinggroup.org
  Trusted Platform Module (TPM), Trusted Computing Group Specification 1.2 Revision 103, July 9, 2007
  TPM 2.0 Library Specification

- Underwriters Laboratories UL 1642 Standard for Safety for Lithium Batteries

- USB Specifications https://www.usb.org/
  - USB 2.0
  - USB 3.0, 3.1, 3.2
  - USB4 also known as "Thunderbolt 4"
  - USB Type-C Connector and Power Delivery specifications

- VESA (Video Electronics Standards Association) https://www.vesa.org
  - DisplayPort Interoperability Guideline Version 1.1a, February 5, 2009
    http://www.vesa.org/vesa-standards/free-standards/
  - DisplayPort Standard Version 1.4
  - DisplayPort Standard Version 2.0
  - Embedded DisplayPort (eDP) Specification Rev. 1.4b, Oct 10, 2015

2. COM-HPC Interfaces

2.1. COM-HPC Client and Server Pinout Differences

The complete listings of signal descriptions and connector pin assignments for the COM-HPC Client and Server pinout types are found in the ***PICMG COM-HPC Module Base Specification*** and are not repeated here. Table 2 below details the Module connector pin assignments that **differ** between the Client and Server types.

Table 2: Client and Server Type Pinout Difference Table

Pin	Row	Client	Server
20	A	DDI1_SDA_AUX-	ETH4_RX-
21	A	DDI1_SCL_AUX+	ETH4_RX+
23	A	DDI1_PAIR0-	ETH5_RX-
24	A	DDI1_PAIR0+	ETH5_RX+
26	A	DDI1_PAIR1-	ETH6_RX-
27	A	DDI1_PAIR1+	ETH6_RX+
29	A	DDI1_PAIR2-	ETH7_RX-
30	A	DDI1_PAIR2+	ETH7_RX+
32	A	DDI1_PAIR3-	RSVD
33	A	DDI1_PAIR3+	RSVD
35	A	eDP_AUX-	ETH4_TX-
36	A	eDP_AUX+	ETH4_TX+
38	A	eDP_TX0-	ETH5_TX-
39	A	eDP_TX0+	ETH5_TX+
41	A	eDP_TX1-	ETH6_TX-
42	A	eDP_TX1+	ETH6_TX+
44	A	eDP_TX2-	ETH7_TX-
45	A	eDP_TX2+	ETH7_TX+
47	A	eDP_TX3-	USB1_AUX-
48	A	eDP_TX3+	USB1_AUX+
19	B	I2S_LRCLK/SNDW_CLK3	RSVD
20	B	I2S_DOUT/SNDW_DAT3	RSVD
21	B	I2S_MCLK	RSVD
22	B	I2S_DIN/SNDW_DAT2	RSVD
23	B	I2S_CLK/SNDW_CLK2	RSVD
45	B	LID#	RSVD
46	B	SLEEP#	RSVD
20	C	SNDW_DMIC_CLK1	ETH0_TX-
21	C	SNDW_DMIC_DAT1	ETH0_TX+
23	C	SNDW_DMIC_CLK0	ETH1_TX-
24	C	SNDW_DMIC_DAT0	ETH1_TX+
26	C	DDI0_DDC_AUX_SEL	ETH2_TX-
27	C	DDI1_DDC_AUX_SEL	ETH2_TX+
28	C	DDI0_HPD	GND
29	C	DDI1_HPD	ETH3_TX-
30	C	eDP_HPD	ETH3_TX+
31	C	eDP_VDD_EN	GND
32	C	eDP_BKLT_EN	USB3_SSRX-
33	C	eDP_BKLTCTL	USB3_SSRX+
35	C	USB1_AUX-	USB2_SSRX-
36	C	USB1_AUX+	USB2_SSRX+
19	D	DDI0_SDA_AUX-	ETH0_RX-
20	D	DDI0_SCL_AUX+	ETH0_RX+
22	D	DDI0_PAIR0-	ETH1_RX-
23	D	DDI0_PAIR0+	ETH1_RX+
25	D	DDI0_PAIR1-	ETH2_RX-
26	D	DDI0_PAIR1+	ETH2_RX+

Pin	Row	Client	Server
28	D	DDI0_PAIR2-	ETH3_RX-
29	D	DDI0_PAIR2+	ETH3_RX+
31	D	DDI0_PAIR3-	USB3_SSTX-
32	D	DDI0_PAIR3+	USB3_SSTX+
34	D	AC_PRESENT	USB2_SSTX-
35	D	RSVD	USB2_SSTX+
3	E	DDI2_SDA_AUX-	RSVD
4	E	DDI2_SCL_AUX+	RSVD
6	E	DDI2_PAIR0-	RSVD
7	E	DDI2_PAIR0+	RSVD
9	E	DDI2_PAIR1-	RSVD
10	E	DDI2_PAIR1+	RSVD
12	E	DDI2_PAIR2-	RSVD
13	E	DDI2_PAIR2+	RSVD
15	E	DDI2_PAIR3-	RSVD
16	E	DDI2_PAIR3+	RSVD
18	E	DDI2_DDC_AUX_SEL	RSVD
19	E	DDI2_HPD	RSVD
69	E	RSVD	PCIe48_TX-
70	E	RSVD	PCIe48_TX+
71	E	RSVD	GND
72	E	RSVD	PCIe49_TX-
73	E	RSVD	PCIe49_TX+
74	E	RSVD	GND
75	E	RSVD	PCIe50_TX-
76	E	RSVD	PCIe50_TX+
77	E	RSVD	GND
78	E	NBASET1_CTREF	PCIe51_TX-
79	E	NBASET1_SDP	PCIe51_TX+
80	E	NBASET1_LINK_MID#	GND
81	E	NBASET1_LINK_ACT#	PCIe52_TX-
82	E	NBASET1_LINK_MAX#	PCIe52_TX+
84	E	RSVD	PCIe53_TX-
85	E	RSVD	PCIe53_TX+
87	E	ETH0_RX-	PCIe54_TX-
88	E	ETH0_RX+	PCIe54_TX+
90	E	ETH1_RX-	PCIe55_TX-
91	E	ETH1_RX+	PCIe55_TX+
1	F	RSVD	ETH2_SDP
2	F	RSVD	ETH3_SDP
3	F	RSVD	ETH4_SDP
4	F	RSVD	ETH5_SDP
5	F	RSVD	ETH6_SDP
6	F	RSVD	ETH7_SDP
7	F	RSVD	ETH4-7_I2C_CLK
8	F	RSVD	ETH4-7_I2C_DAT
9	F	RSVD	ETH4-7_INT#
10	F	RSVD	ETH4-7_MDIO_CLK
11	F	RSVD	ETH4-7_MDIO_DAT
12	F	RSVD	ETH4-7_PHY_INT#
13	F	RSVD	ETH4-7_PHY_RST#
14	F	RSVD	ETH4-7_PRSNT#
68	F	RSVD	PCIe48_RX-
69	F	RSVD	PCIe48_RX+
71	F	NBASET1_MDI0-	PCIe49_RX-
72	F	NBASET1_MDI0+	PCIe49_RX+
74	F	NBASET1_MDI1-	PCIe50_RX-
75	F	NBASET1_MDI1+	PCIe50_RX+
77	F	NBASET1_MDI2-	PCIe51_RX-
78	F	NBASET1_MDI2+	PCIe51_RX+

Pin	Row	Client	Server
80	F	NBASET1_MDI3-	PCIe52_RX-
81	F	NBASET1_MDI3+	PCIe52_RX+
83	F	RSVD	PCIe53_RX-
84	F	RSVD	PCIe53_RX+
86	F	ETH0_TX-	PCIe54_RX-
87	F	ETH0_TX+	PCIe54_RX+
88	F	ETH1_TX-	PCIe55_RX-
90	F	ETH1_TX+	PCIe55_RX+
95	F	RSVD	PCIe_CLKREQ3#
96	F	ETH0-1_PRSNT#	ETH0-3_PRSNT#
97	F	ETH0-1_PHY_RST#	ETH0-3_PHY_RST#
2	G	GND	RSVD
3	G	USB2_SSRX0-	RSVD
4	G	USB2_SSRX0+	RSVD
5	G	GND	RSVD
6	G	USB2_SSRX1-	RSVD
7	G	USB2_SSRX1+	RSVD
8	G	GND	RSVD
9	G	USB3_SSRX0-	RSVD
10	G	USB3_SSRX0+	RSVD
11	G	GND	RSVD
12	G	USB3_SSRX1-	RSVD
13	G	USB3_SSRX1+	RSVD
15	G	USB3_LSRX	RSVD
16	G	USB3_LSTX	RSVD
17	G	USB2_LSRX	RSVD
18	G	USB2_LSTX	RSVD
19	G	PEG_LANE_REV#	RSVD
69	G	RSVD	PCIe56_RX-
70	G	RSVD	PCIe56_RX+
72	G	CSI0_RX0-	PCIe57_RX-
73	G	CSI0_RX0+	PCIe57_RX+
75	G	CSI0_RX1-	PCIe58_RX-
76	G	CSI0_RX1+	PCIe58_RX+
78	G	CSI0_RX2-	PCIe59_RX-
79	G	CSI0_RX2+	PCIe59_RX+
81	G	CSI0_RX3-	PCIe60_RX-
82	G	CSI0_RX3+	PCIe60_RX+
84	G	CSI0_CLK-	PCIe61_RX-
85	G	CSI0_CLK+	PCIe61_RX+
87	G	CSI0_I2C_CLK	PCIe62_RX-
88	G	CSI0_I2C_DAT	PCIe62_RX+
89	G	CSI0_MCLK	GND
90	G	CSI0_RST#	PCIe63_RX-
91	G	CSI0_ENA	PCIe63_RX+
93	G	RSVD	PCIe_REFCLK3-
94	G	RSVD	PCIe_REFCLK3+
96	G	ETH0-1_I2C_CLK	ETH0-3_I2C_CLK
97	G	ETH0-1_I2C_DAT	ETH0-3_I2C_DAT
98	G	ETH0-1_PHY_INT#	ETH0-3_PHY_INT#
99	G	ETH0-1_INT#	ETH0-3_INT#
1	H	GND	RSVD
2	H	USB2_SSTX0-	RSVD
3	H	USB2_SSTX0+	RSVD
4	H	GND	RSVD
5	H	USB2_SSTX1-	RSVD
6	H	USB2_SSTX1+	RSVD
7	H	GND	RSVD
8	H	USB3_SSTX0-	RSVD
9	H	USB3_SSTX0+	RSVD

Pin	Row	Client	Server
10	H	GND	RSVD
11	H	USB3_SSTX1-	RSVD
12	H	USB3_SSTX1+	RSVD
13	H	GND	RSVD
14	H	USB2_AUX-	RSVD
15	H	USB2_AUX+	RSVD
16	H	GND	RSVD
17	H	USB3_AUX-	RSVD
18	H	USB3_AUX+	RSVD
68	H	RSVD	PCIe56_TX-
69	H	RSVD	PCIe56_TX+
71	H	CSI1_RX0-	PCIe57_TX-
72	H	CSI1_RX0+	PCIe57_TX+
74	H	CSI1_RX1-	PCIe58_TX-
75	H	CSI1_RX1+	PCIe58_TX+
77	H	CSI1_RX2-	PCIe59_TX-
78	H	CSI1_RX2+	PCIe59_TX+
80	H	CSI1_RX3-	PCIe60_TX-
81	H	CSI1_RX3+	PCIe60_TX+
83	H	CSI1_CLK-	PCIe61_TX-
84	H	CSI1_CLK+	PCIe61_TX+
86	H	CSI1_I2C_CLK	PCIe62_TX-
87	H	CSI1_I2C_DAT	PCIe62_TX+
88	H	CSI1_MCLK	GND
89	H	CSI1_RST#	PCIe63_TX-
90	H	CSI1_ENA	PCIe63_TX+
98	H	ETH0-1_MDIO_CLK	ETH0-3_MDIO_CLK
99	H	ETH0-1_MDIO_DAT	ETH0-3_MDIO_DAT

2.2. COM-HPC Client and COM-HPC Mini Pinout Differences

The pinout differences between the COM-HPC Client and the COM-HPC Mini are well covered in the COM-HPC Base Specification Revision 1.20 Section 10. Please refer to that document.

3. Reference Schematics and Block Diagrams

3.1. Sources for Technical Materials

The schematic diagrams, block diagrams and mechanical diagrams in this document were contributed by several companies and organizations, including Adlink, Advantech, Avnet Integrated, Bielefeld University, congatec, Intel, Kontron, Samtec and SECO. Hence the graphic styles vary a bit. An effort has been made to provide part numbers in the drawings that can be located in a web search (except for small generic parts).

3.2. Schematic Conventions

Schematic examples are drawn with signal directions shown per the Figure below. Signals that connect directly to the COM-HPC connector are flagged with the text "COM" in the off-page connect symbol, as shown in Figure 1 below. Nets that connect to the COM-HPC Module are named per the PICMG COM-HPC specification in almost all cases.

Figure 1: Schematic Conventions

Power nets shown in the sample schematics and drawings are labeled, for the most part, per the Table below. The power rail behavior under the various system power states is shown in the Table.

Table 3: Power Net Naming

Power Net	S0 On	S3 Suspend to RAM	S4 Suspend to Disk	S5 Soft Off	G3 Mechanical Off
+12V_S	12V	off	off	off	off
+5V_S	5V	off	off	off	off
+3.3V_S	3.3V	off	off	off	off
+1.5V_S	1.5V	off	off	off	off
+2.5V_S	2.5V	off	off	off	off
+5V_A	5V	5V	5V	5V	off
+3.3V_A	3.3V	3.3V	3.3V	3.3V	off
VCC_RTC	3.0V	3.0V	3.0V	3.0V	3.0V

3.3. Ethernet NBASE-T

A typical NBASE-T implementation is shown in Figure 2 below. The "N" refers to the link speed, and may be 10 Mbps, 100 Mbps, 1 Gbps, 2.5 Gbps, 5 Gbps or 10 Gbps. However not all speeds may be available on all Module designs. All COM-HPC Modules are required to support at least the 1 Gbps rate.

This example shows a "Mag-Jack" (an 8 pin RJ45 jack with integrated isolation magnetics) from Wurth Electronics, p/n 7499611420. This part is claimed by Wurth to support 10 Mbps, 100 Mbps, 1 Gbps and 10 Gbps data rates. There are many similar parts from Wurth and from vendors such as Bel-Fuse, Pulse Electronics and others. It may be advisable to check with your Module vendor on the suitability of chosen parts. Ethernet RJ45 jacks (including Mag-Jacks and jacks that require external magnetics) come in "tab-up" and "tab-down" versions. The Wurth part shown here is "tab-up". If a "tab-down" part is used then the PCB layout is impacted as the pin orientation is effectively flipped 180 degrees.

Implementing magnetics that are external to the jack is of course possible but it is trickier. It may be necessary in certain situations that require a higher than normal isolation between the Ethernet magnetics primary and secondary sides. This can be the case in safety critical designs such as medical equipment.

The colors and meanings of the colors used for NBASE-T LEDs are not standardized in the industry. The scheme shown in the diagram below is suggested for COM-HPC but not required

The diagram below shows ESD protection diode arrays (Texas Instruments TPD4E02B04) protecting the NBASE-T differential pairs. Many similar parts are available from other vendors. Make sure the selected part has a suitably low pin capacitance. It is very important that the parts (D43 and D44 in the Figure 2) are placed close to the connector and are routed in a "no stub" fashion. For example the net NBASET0+ in the Figure 2 should hit D43 pin 6 and continue under the D43 package to catch pin 5 and then on to the RJ45 connector. Pins may be swapped for easier routing, as long as the pairs are kept together and the no-stub routing is followed.

Figure 2:   NBASE-T
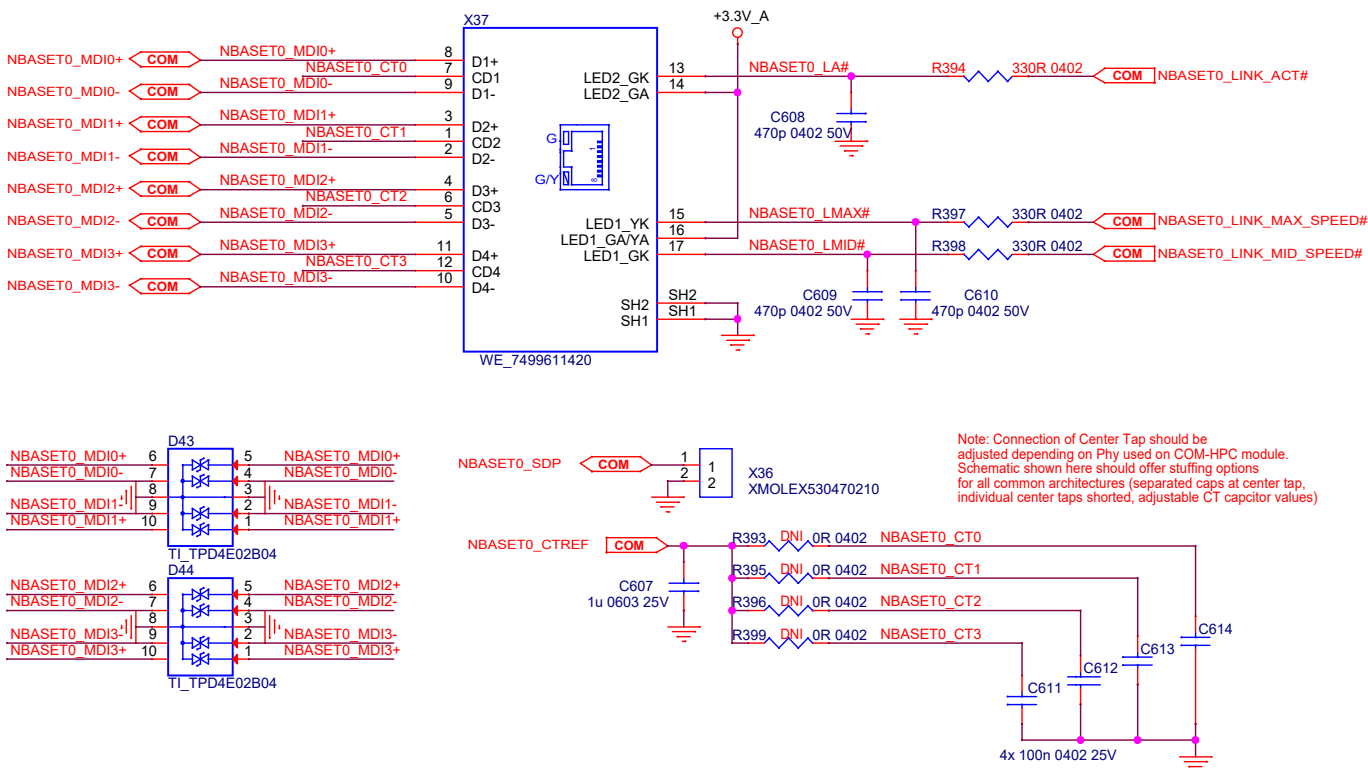
This example applies to the COM-HPC Mini as well. The LED voltage feeds are left at 3.3V to allow sufficient forward voltage to operate the LEDs.

3.4. Ethernet KR and KR4

This document section is not relevant for the COM-HPC Mini.

The Ethernet KR interfaces consist of a single TX pair and single RX pair. These pairs are capacitively coupled off of the COM-HPC Module – either on the Carrier board (for Module to PHY or Module to Module situations) or within the SFP assemblies.

Ethernet KR4 interfaces are comprised of four TX pairs and four RX pairs, capacitively coupled off of the COM-HPC Module, as per the KR interfaces.

In order to save pins, the side band signals for the 10G / 25G / 40G / 100G Ethernet KR interfaces are serialized on the Module silicon per an Intel convention known as CEI. This is an acronym for "Common Electrical Interface". The serialized CEI signals need to be deserialized on the Carrier Board. The block diagrams in this section describe which components are needed and what the functions of the deserialized nets are.

The Ethernet KR LED information is carried on one I2C bus per four Ethernet KR channels, known as a Quad. The I2C buses are named ETH0-3_I2C* (where the * indicates the final characters of the net name in that signal group) and ETH4-7_I2C* for the Server type. As the Client only supports 2 Ethernet channels, the group is named ETH0-1_I2C*.

There is one MDIO bus per Quad available to configure the PHYs on the Carrier Board.
The MDIOs are named ETH0-3_MDIO* and ETH4-7_MDIO* for the Server type. As the Client only supports 2 Ethernet channels, the group is named ETH0-1_MDIO*.

The Reset and Interrupt signals also follow the same naming convention.

The SDP signals are more critical in timing and are available directly.

3.4.1. Ethernet KR CEI Block Diagrams

Ethernet KR CEI concepts are illustrated in block diagram format in Figures 3 through 9 on the following pages. Many of the details of these implementations are vendor NDA protected. Some references to vendor document numbers for confidential material are listed after each Figure, if material is available. Designers interested in these materials need to contact the silicon vendors directly and work out the necessary NDAs.

Figure 3: *Intel SoC with CEI Boundary*
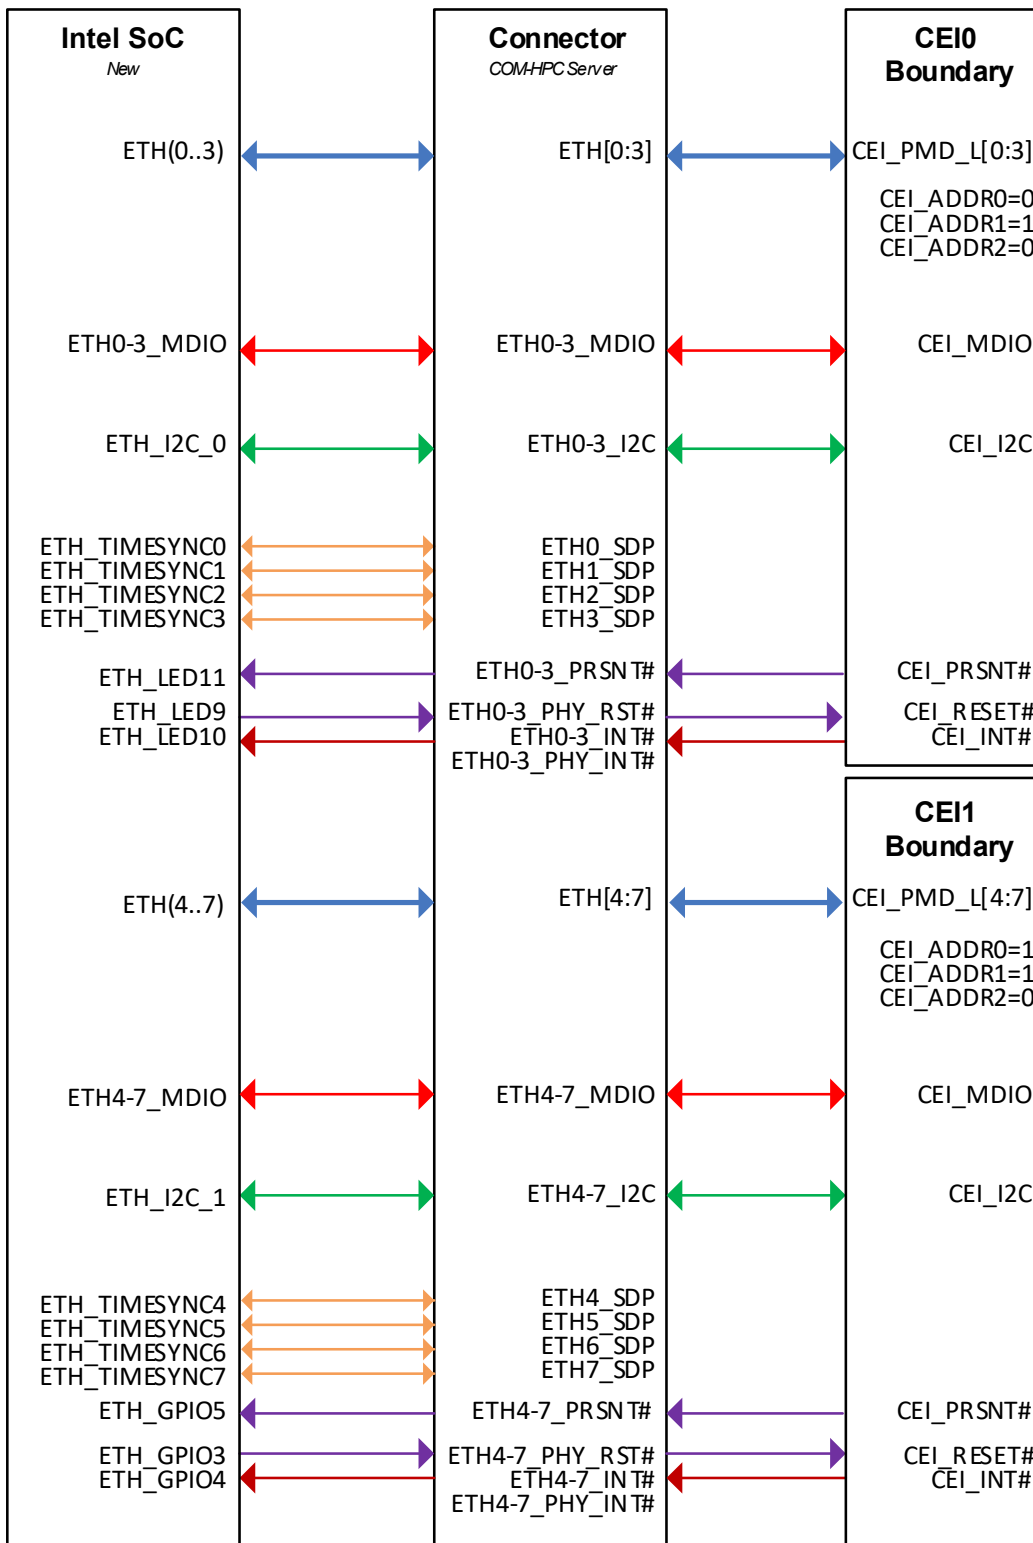
For further details on this configuration, refer to NDA protected Intel documents 620640 and 631178.

Figure 4:   **Intel CEI 4x SFP28**
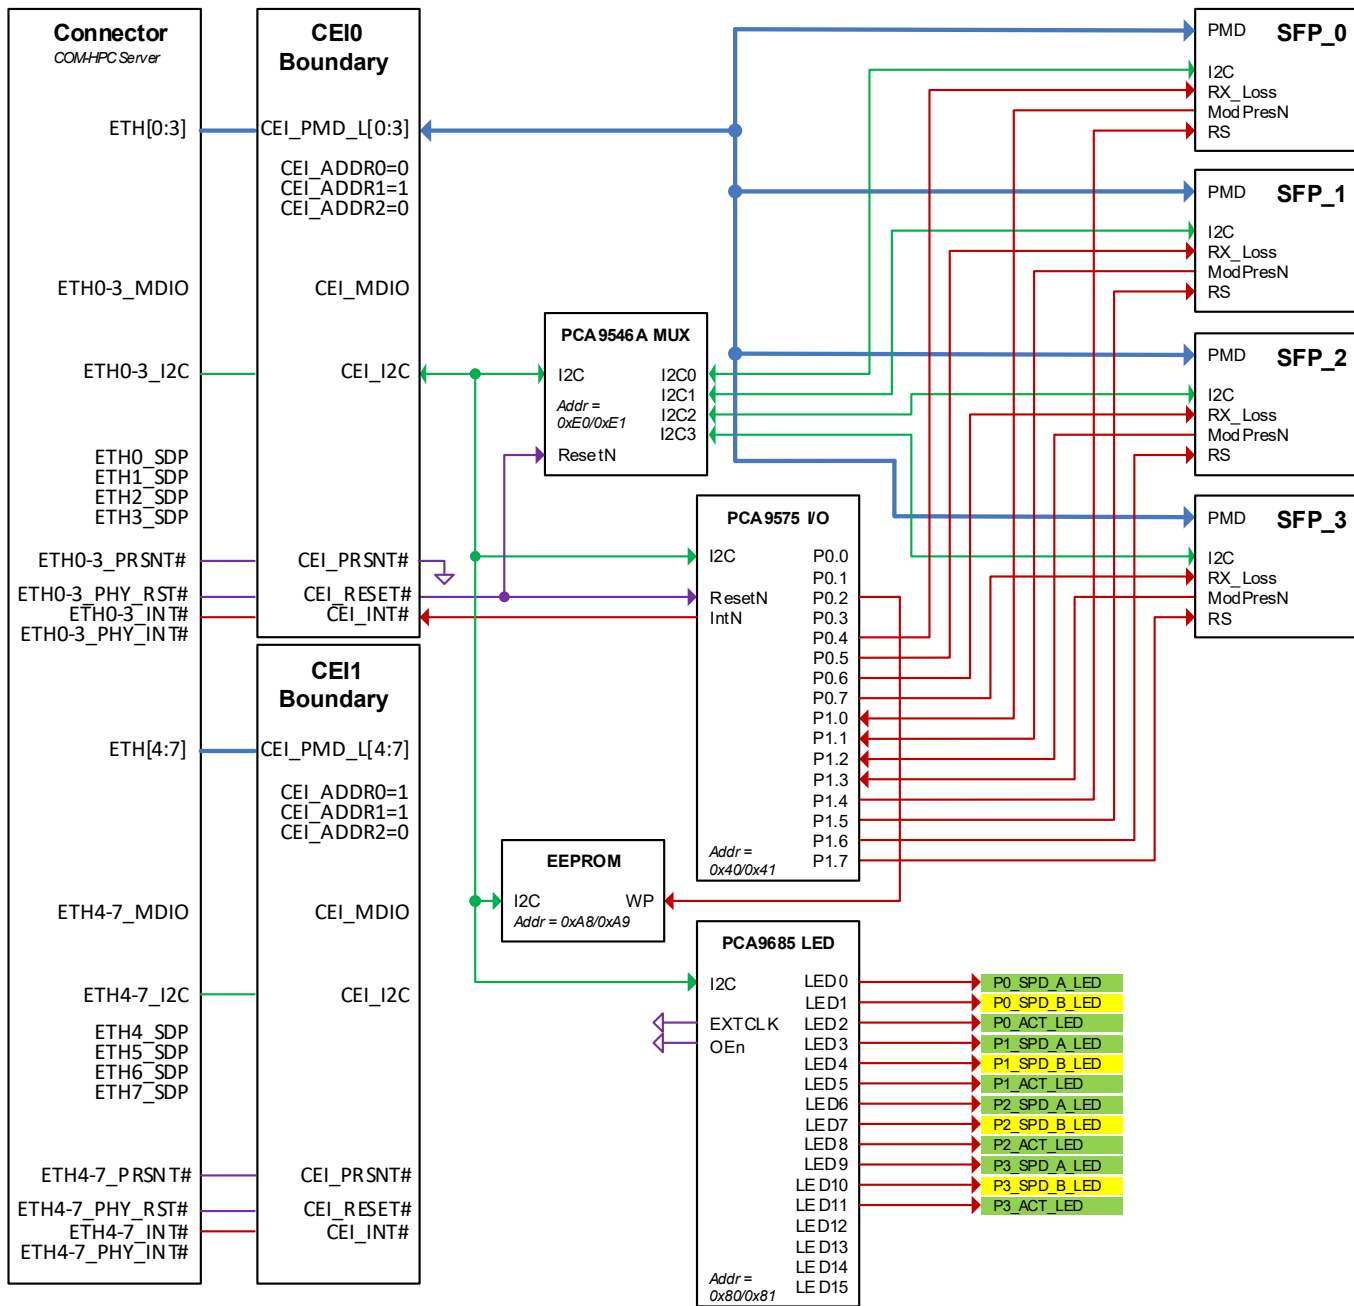
This configuration is not supported by additional Intel documentation at the time of this writing.

Figure 5: **Intel CEI 4x SFP28 with Retimer C827/XL827**
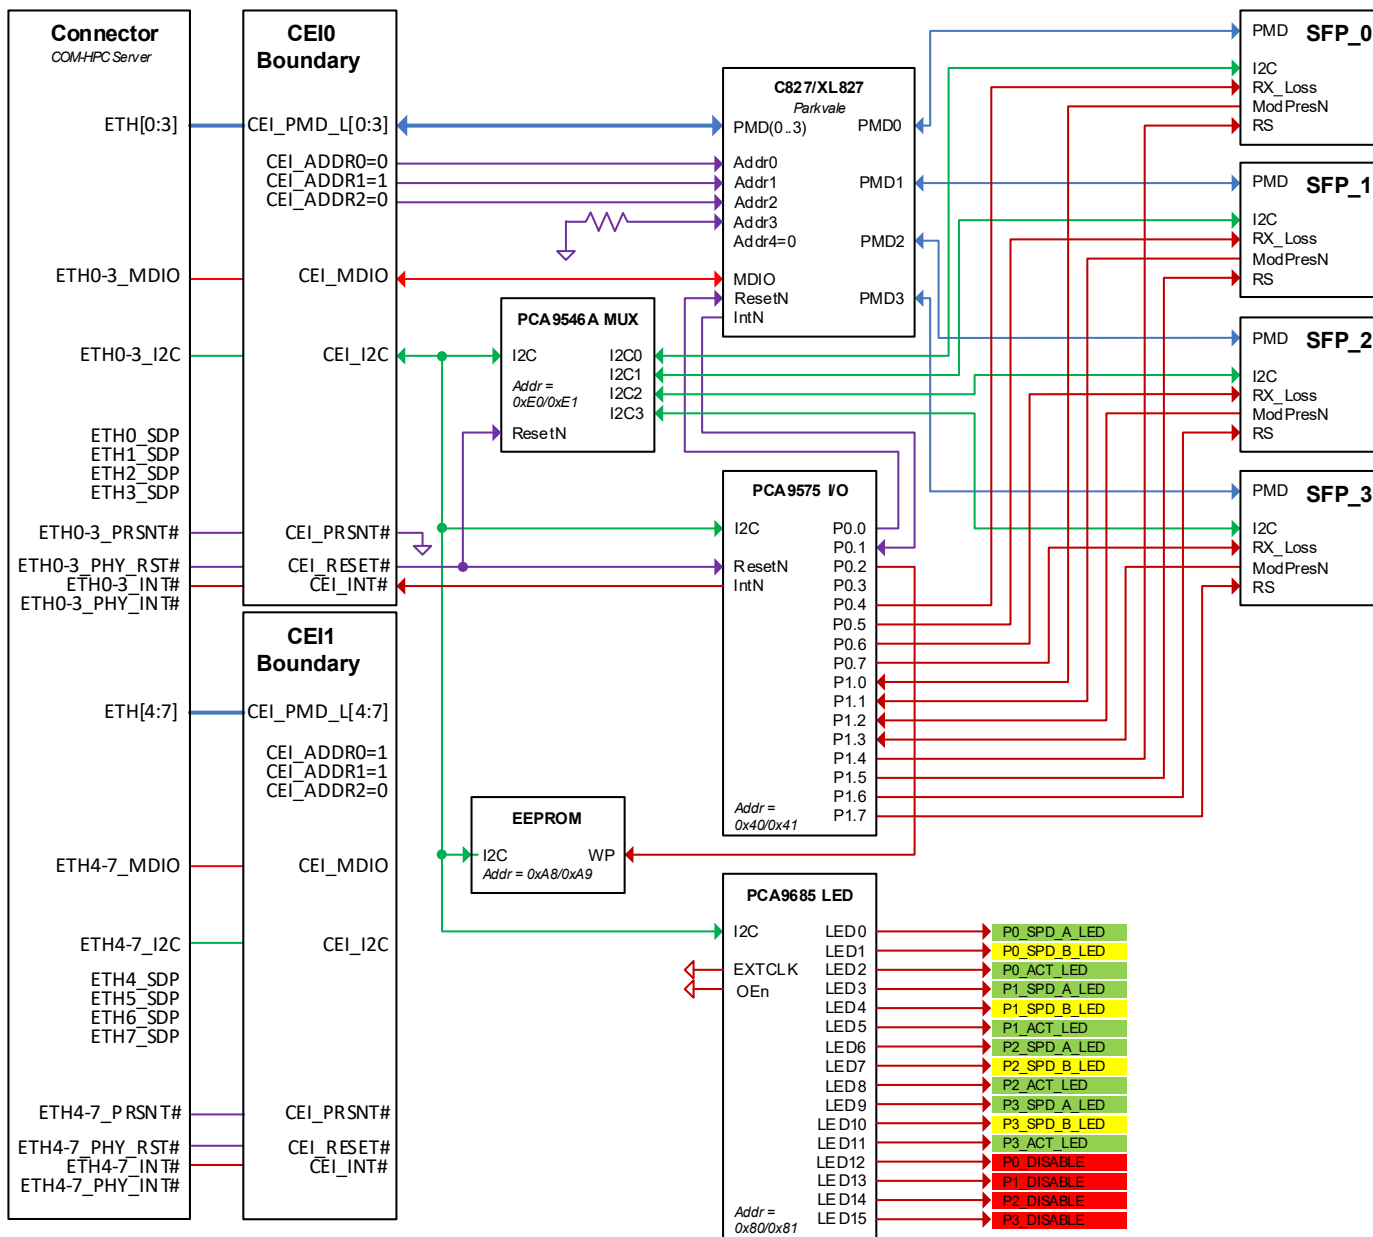
For further details on this configuration, refer to NDA Intel document 636564.

Figure 6: *Intel CEI QSFP28 with Retimer C827/XL827*
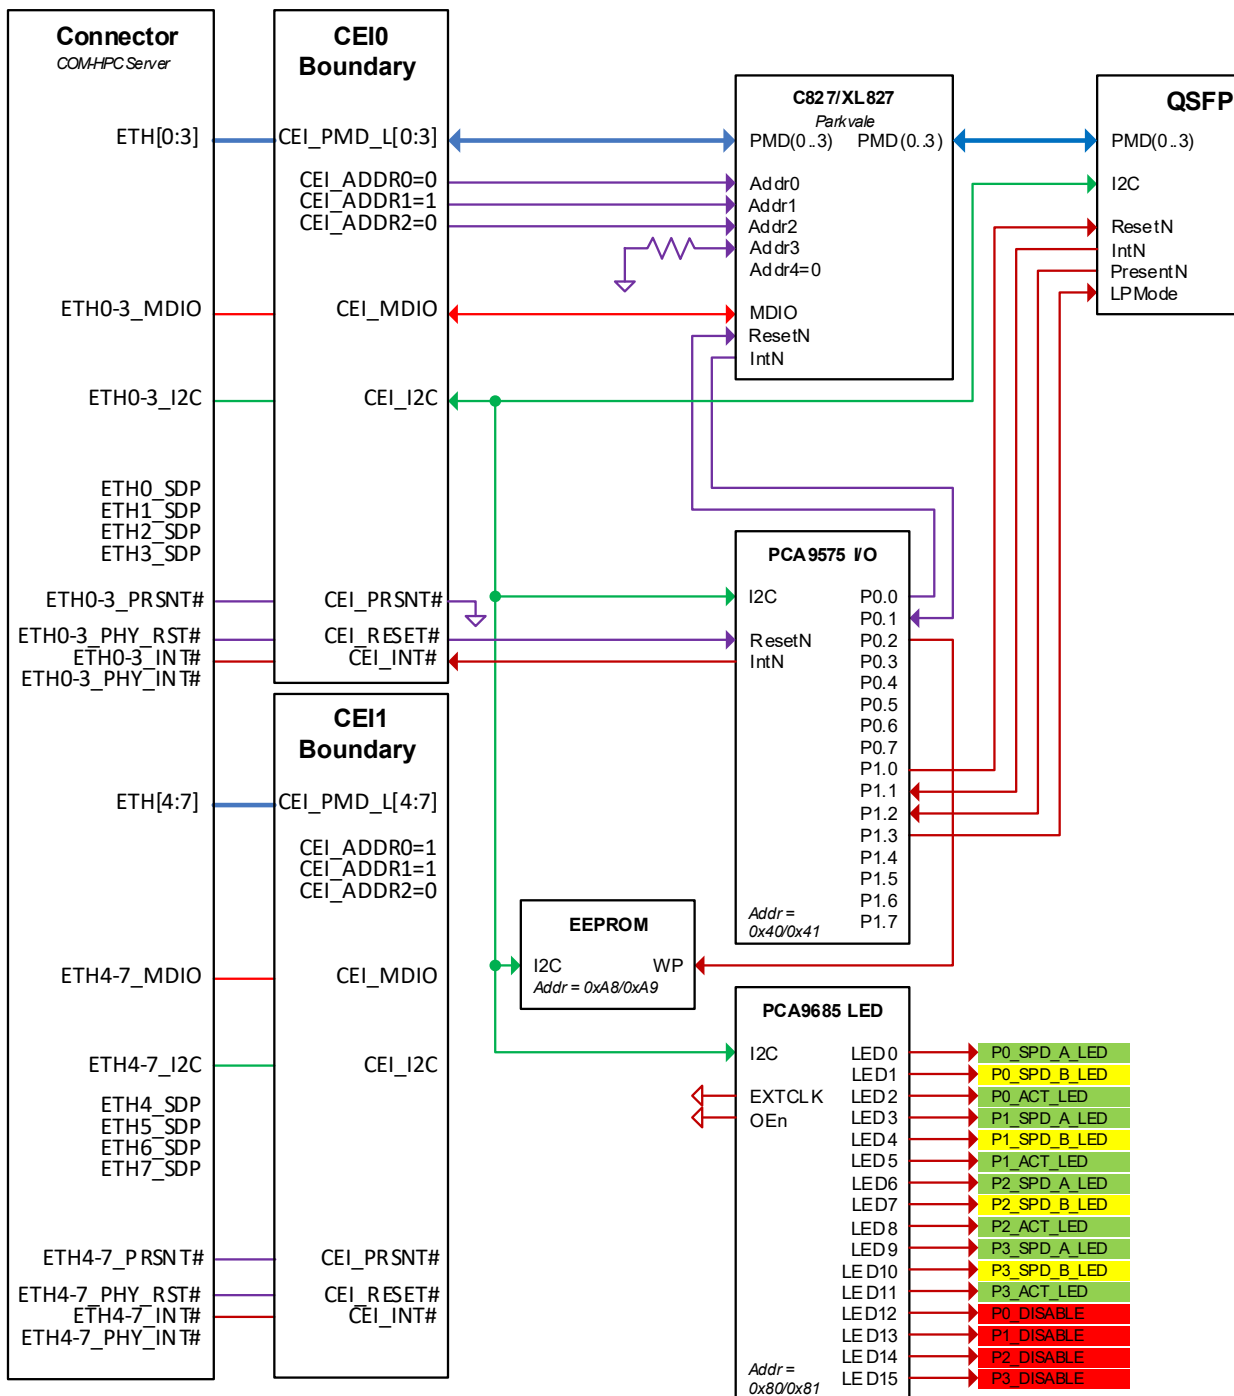
For further details on this configuration, refer to NDA Intel document 636566.

Figure 7: *Intel CEI 4x 10GBASE-T Copper PHY X557-AT4*
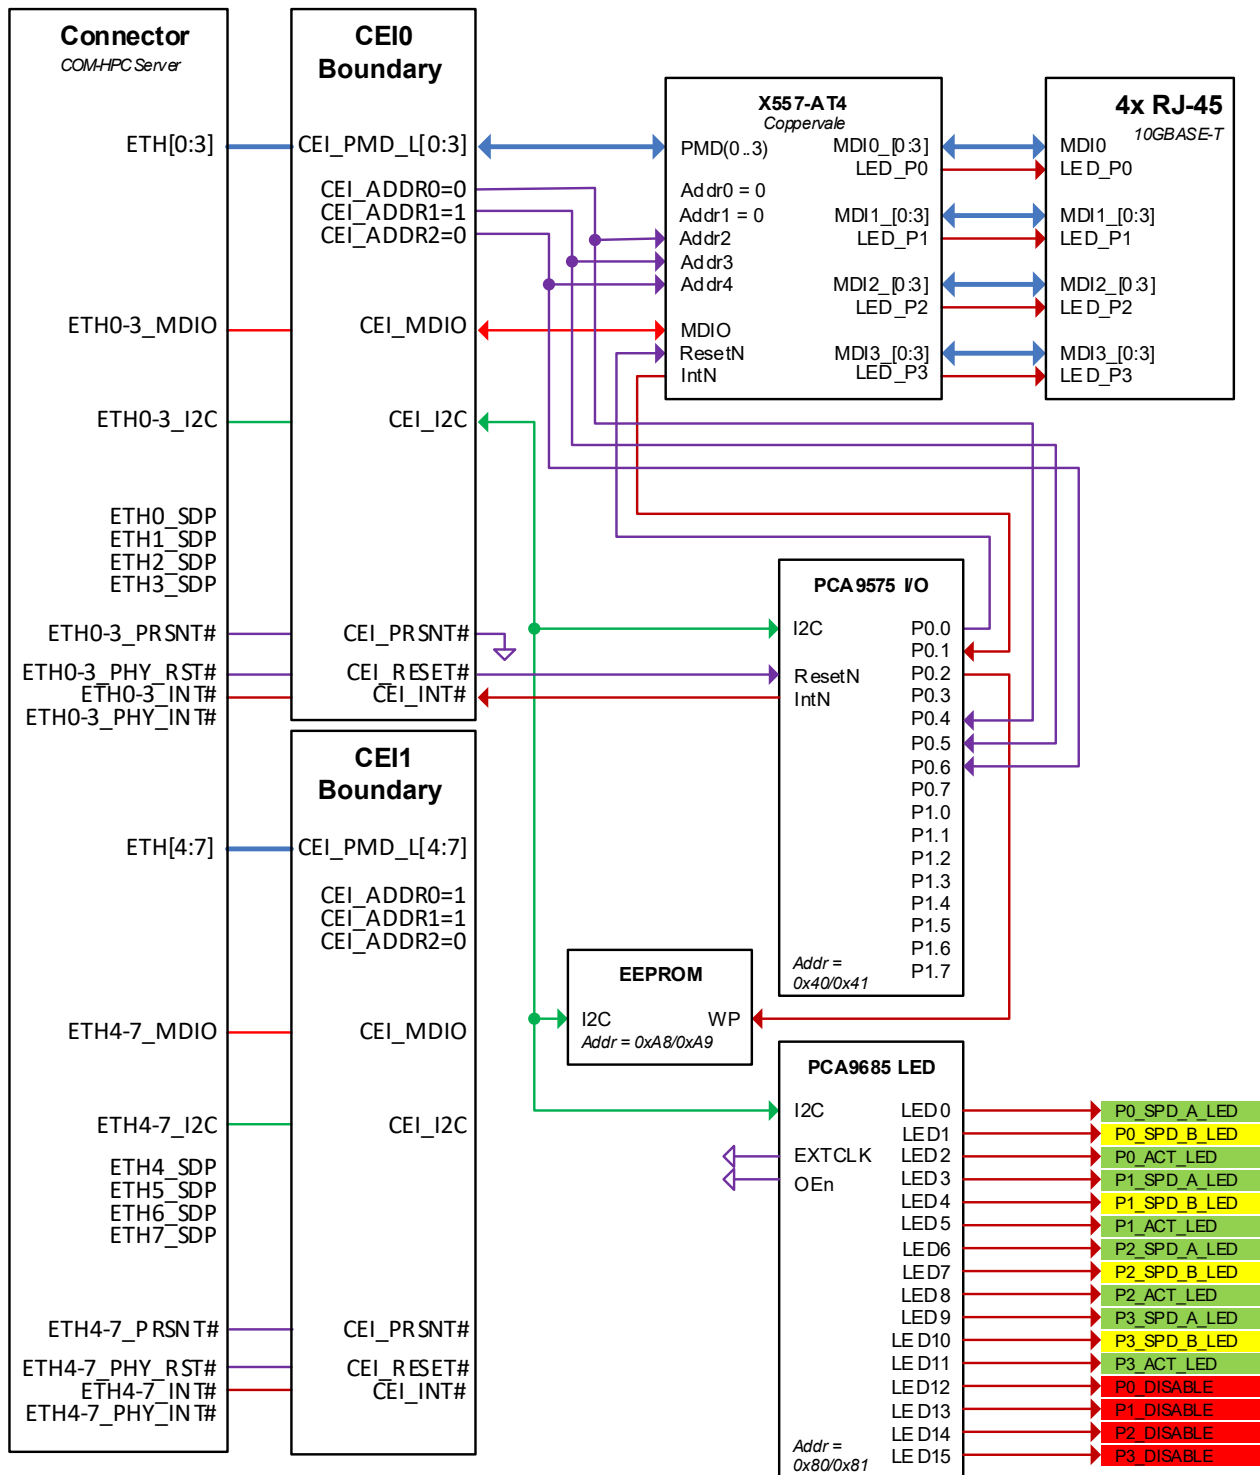
For further details on this configuration, refer to NDA Intel document 613899.

Figure 8: *Intel CEI 4x 1GBASE-T Copper PHY Marvell 88E1543*
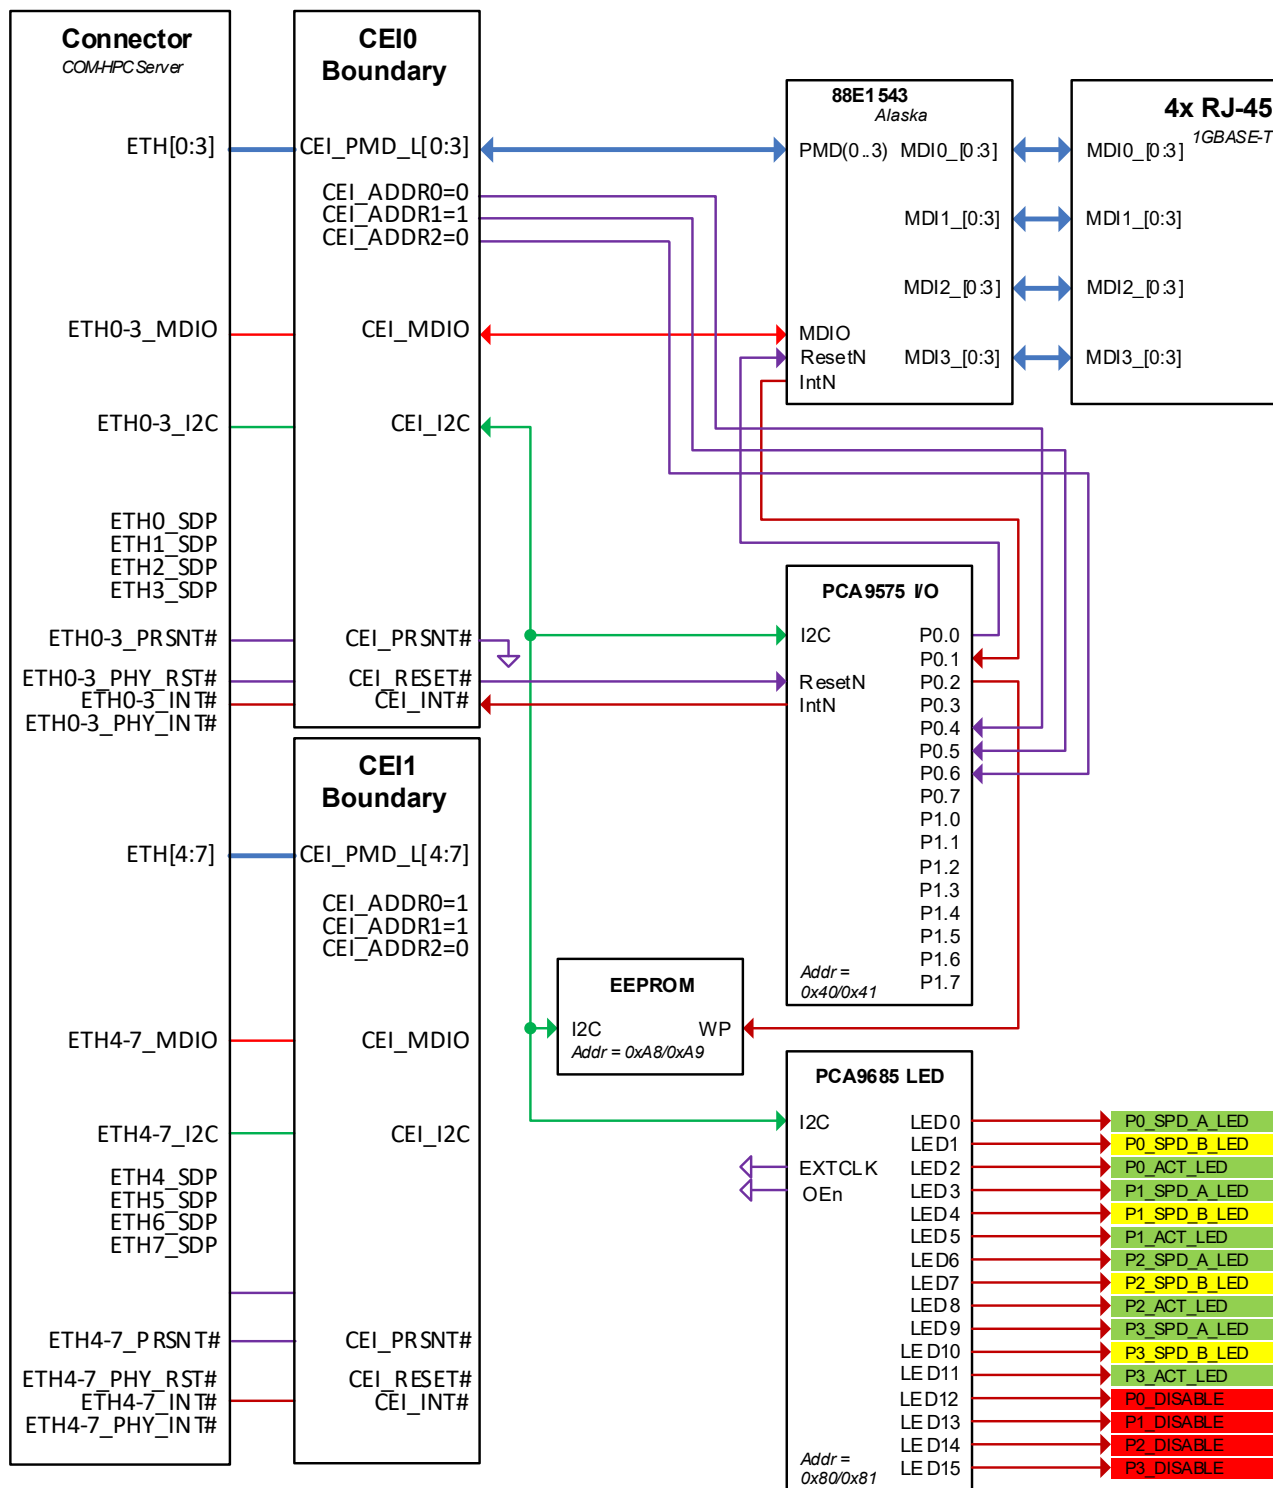
For further details on this configuration, refer to NDA Intel document 613900.

Figure 9:   AMD SoC 4xSFP+ with CS4223 Retimer
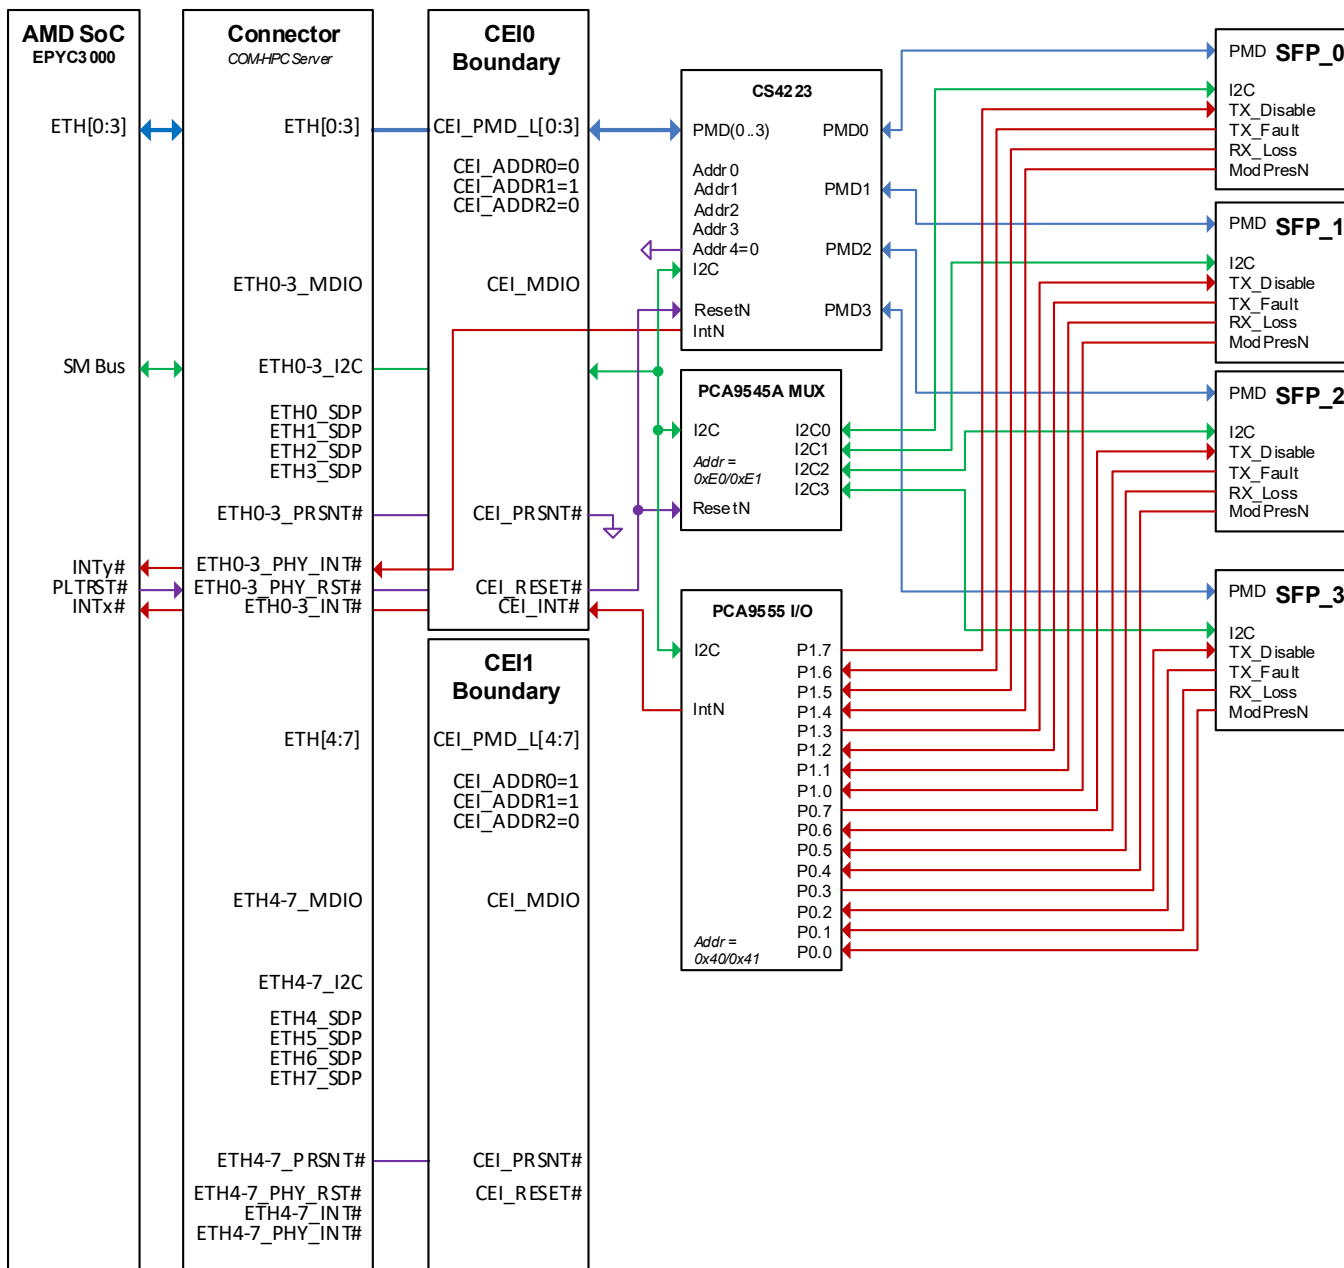
For further details on this configuration, contact AMD.

3.4.2. PHY Addresses

Table 4: *MDIO Addresses for Intel POR External PHYs*

PHY	MDIO Address (Decimal)	Ethernet Quad / Port
Intel "Parkvale"	2	Quad0
	3	Quad1
Intel "Coppervale" Marvell offers a similar PHY	8	Quad0 Port0
	9	Quad0 Port1
	10	Quad0 Port2
	11	Quad0 Port3
	12	Quad1 Port0
	13	Quad1 Port1
	14	Quad1 Port2
	15	Quad1 Port3

3.5. Serial ATA

3.5.1. Cabled SATA

The COM-HPC pinouts offer up to two SATA ports, designated SATA0 and SATA1. The implementation for a cabled interface is straightforward, as illustrated in Figure 10 below. No Carrier coupling capacitors are needed as they are specified to be present on the Module. The connections between the Module and the SATA connector are simple differential pairs. Some routing rules may be found in Section 4. below.

Two common connector styles used for cabled SATA implementations are shown in the Figure 10. The upper image shows a 7 pin data-only connector. Power to the SATA drive is handled separately in this case. The lower image in the Figure 10 shows a 22 pin connector that handles SATA data and power. There are three power rails defined on this connector, but all three are not necessarily used. Smaller format drives tend to use just one or two of these rails, Check the drive specifications.

Figure 10 below shows two typical COM-HPC cabled SATA. Note how the data pair polarity order flips along the connector pins: TX+ TX- GND RX- RX+ … this is not a mistake, but is part of the SATA specification.

Figure 10: Serial ATA – Cabled
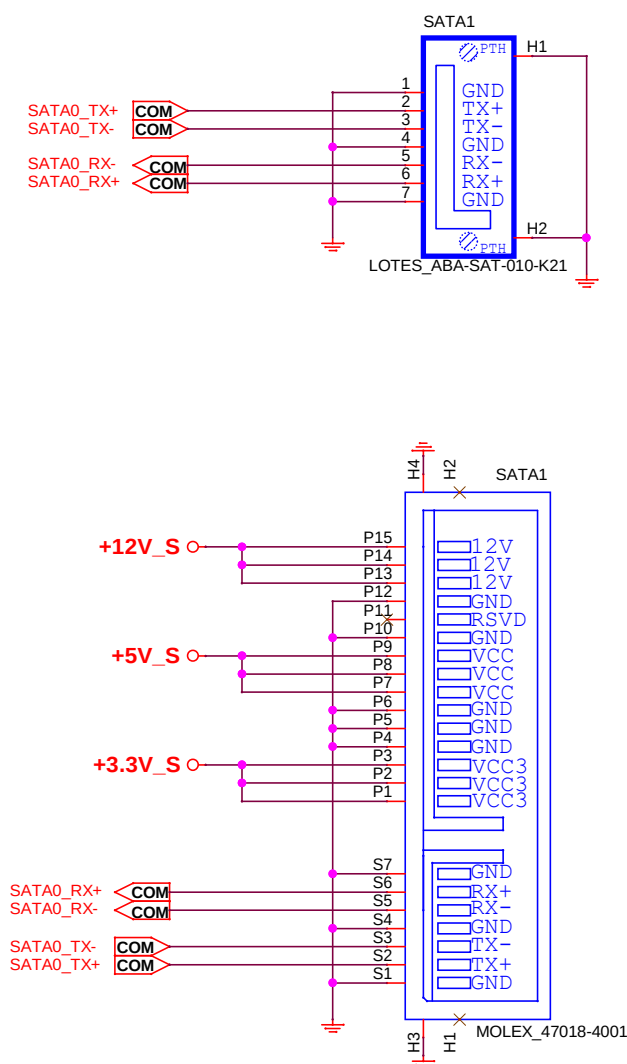
This Figure 10 applies to the COM-HPC Mini except that Carrier based coupling caps on the SATA TX and RX pairs (4 capacitors) are needed. See COM-HPC Base Specification Rev. 1.20 Section 10.

3.5.2. mSATA SSDs

The SATA specification defines a small form factor card for SATA SSDs that is almost identical in mechanical and electrical definition to the PCI-SIG miniPCIe format. The same card connector and mounting holes are used. Both half size and full size cards are in use. Sometimes dual mini-PCIe / mSATA implementations are executed. This involves multiplexing four signals per the following Table:

Table 5:   *mSATA Pin Mapping Relative to miniPCIe*

MiniPCIe Card Pin Name	MiniPCIe / mSATA Card Pin Number	PCIE Signal Relative to COM-HPC Module	SATA Signal Relative to COM-HPC Module
PETP0	33	PCIe TX+	SATA TX+
PETN0	31	PCIe TX-	SATA TX-
PERP0	25	PCIe RX+	SATA RX-
PERN0	23	PCIe RX-	SATA RX+

The SATA_RX- mapping to miniPCIe PERP0 and SATA_RX+ to PERN0 is intentional per the SATA specification.

The signals do not have to be multiplexed if a mSATA only or miniPCIe only implementation is desired.

SSD implementations are largely moving away from the miniPCIe format and to M.2 formats. In the M.2 formats, there are PCIe interfaces defined (x1, x2 and x4) and a SATA interface defined, similar to the miniPCIe / mSATA pin sharing format shown in the Table above. The M.2 PCIe x 4 format, sometimes referred to as NVMe, offer a much higher interface bandwidth than mSATA.

This mSATA information applies to the COM-HPC Mini as well. However the coupling capacitors on the SATA TX and RX pairs are moved off the Module for the Mini and onto the Carrier, close to the Mini Module connector pins. See the COM-HPC Base Specification Rev. 1.20 Section 10.

3.5.3. M.2 SATA SSDs

The PCI Express M.2 Specification defines several M.2 format SATA SSD options that may be used in COM-HPC systems. These include (but are not limited to):

- Socket 2 B-M Key (Table 3-23 in the PCI-SIG Version 4.0 M.2 document)

- Socket 3 M Key (Table 3-28)

These are not diagrammed here as SATA SSD implementations seem to be losing ground to PCIe based SSD implementations.

This M.2 SATA information applies to the COM-HPC Mini as well. However the coupling capacitors on the TX pair(s) are located on the Carrier, not on the Module, for the Mini. The coupling capacitors for the Module RX pair (the M.2 SATA card TX pair) are on the M.2 card. See COM-HPC Base Specification Rev. 1.2 Section 10.

3.6. PCI Express

3.6.1. General Notes

The COM-HPC PCIe resources are divided into 5 Groups:

- Group 0 Low (8 lanes)
  - Generally used for smaller links (x1, x2, x4) and slower PCIe link speeds (PCIe Gen 1,2,3).

- Group 0 High (8 lanes)
  - Recommended for use with one or two PCIe x4 NVME SSD instances
  - The COM-HPC specification recommends that higher bandwidth PCIe links be steered to this Group

- Group 1 (16 lanes)
  - Recommended for PEG use
  - The COM-HPC specification recommends that higher bandwidth PCIe links be steered to this Group

- Group 2 (16 lanes)
  - General purpose links – x16 or combinations of x8 and / or x4

- Group 3 (16 lanes)
  - Available on the Server pinout only

Each PCIe Group listed above has it's own 100 MHz PCIe Reference Clock pair from the COM-HPC Module. Additionally, there is a CLKREQ# (Clock Request) input to the Module for each PCIe Group.

There is one additional PCIe link available on both the Client and Server pinouts. This is a x1 link for use with a Carrier BMC (Board Management Controller). The BMC PCIe link makes use of the Group 0 PCIe Reference Clock pair.

If only a single PCIe link (of any link width – x1, x2, x4, x8 or x16) is used from a PCIe Group, then the COM-HPC PCIe Reference Clock pair may be used directly with the link target. If a Group uses more than one link (i.e. 2 or more links) then a Carrier Board PCIe Reference Clock buffer is needed for that Group. Many PCIe clock buffer products are available on the market. Buffers with 2,4,6 or 8 and more output pairs are available. The buffer must be appropriate for the fastest PCIe link in the Group (PCIe Gen 3, 4 or 5). Several examples are shown in the schematics below.

If the link's PCIe target is located on a slot card or a mezzanine board such as an M.2 site, the connector involved must be rated for the fastest PCIe link in use for that target. At the time of this writing, most such connectors are PCIe Gen 3 capable. Gen 4 and Gen 5 capable connectors are becoming available, most visibly from Amphenol FCI.

The COM-HPC Mini allows the implementation of PCIe Groups 0 Low and 0 High, but the coupling caps for the Mini Module TX PCIe Group 0 Low pairs are located on the Carrier. Coupling caps for Mini RX Module lines are located on the add-in card EXCEPT for MXM cards and Carrier "device down" situations in which case the PCIe RX coupling caps are on the Carrier.

The Mini further divides PCIe Group 0 low into two 4 lane groups, designated Group 0 Low A and Low B. See the COM-HPC Base Specification Revision 1.2 Section 10. COM-HPC Mini PCIe fill order guidelines are also covered there.

3.6.2. PCI Express Coupling Capacitor Locations

The proper positions for PCIe data pair coupling capacitors on COM-HPC Client and Server implementations is shown in Figure 11 below.

- COM-HPC Module TX pair coupling caps are on the COM-HPC Module.

- COM-HPC Module RX pair coupling caps are NOT on the COM-HPC Module.

  - For **most** Device Up mezzanine card implementations (Slot card, Mini-PCIe, and M.2 card) the coupling caps are up on the mezzanine card, close to the mezzanine target device TX pins.

    - The exception to this rule is with MXM-3 graphics cards: there are no PCIe coupling caps at all on a MXM-3 graphics card. The COM-HPC Module TX lines are AC coupled on the COM-HPC Module. The COM-HPC Module RX lines are AC coupled on the Carrier, near the MXM-3 Module TX pins.

  - For Device Down implementations, the coupling caps are down on the Carrier board, close to the target device TX pins.

Figure 11: PCIe Data Line Coupling Capacitor Positions (MXM-3 Cards Excluded)
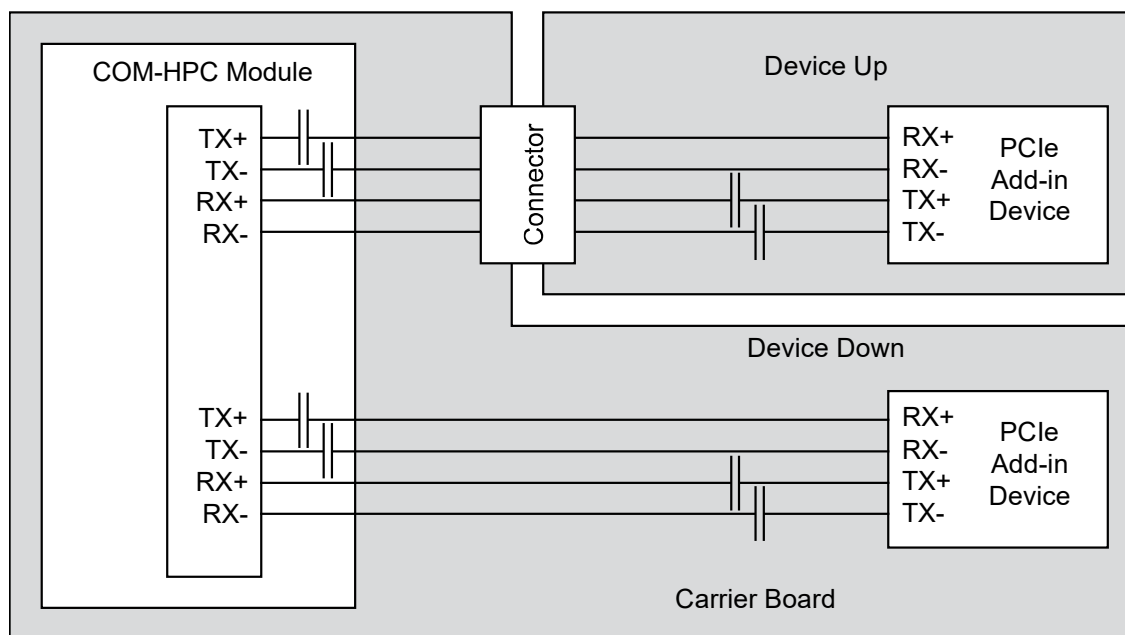
For COM-HPC Mini implementations, the TX pair caps for PCIe Group 0 Low move off the Module and onto the Carrier. See the COM-HPC Base Specification Revision 1.2 Section 10.

3.6.3. PCIe Group 0 Low Examples: Device Down, mini-PCIe, M.2 E-Key, M.2 B-Key

Figure 12: PCIe Device Down on Carrier – PCIe Group 0 Low – PCIe Lane 0
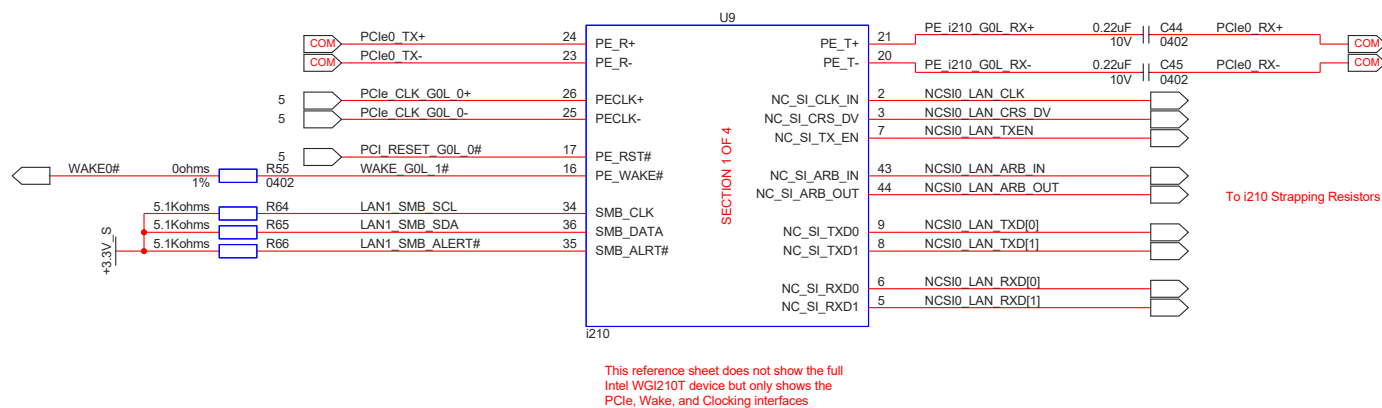
Figure 12 shows a portion of an Intel i210 Gigabit Ethernet implementation, the portion that is relevant to the COM-HPC Module interface. This Figure 12 shows the interface to the COM-HPC and to some Carrier board circuit elements such as an appropriate PCIe Clock buffer, shown later in this section.

The key point of this Figure is that coupling caps (C44 and C45 in the Figure) are needed on the Carrier, for the Carrier target device TX pair pins. These are the PCIe RX lines for the Module. These Carrier coupling caps are to be placed close to the i210 device, in a symmetric manner consistent with high speed PCB design practices.

Coupling caps for COM-HPC PCIe TX lines are on the Module.

If the PCIe target device is a slot or mezzanine card, the coupling caps for the target device are on the slot or mezzanine card and not on the Carrier. - except for MXM graphics card implementations. This case is discussed in Section 3.6.7. below.

The Figure 12 above applies to the COM-HPC Mini as well, except that series coupling capacitors are needed for the PCIe TX pair at the upper left portion of the Figure 12.

There may need to be level translation on the WAKE0# signal if it is destined for a 1.8V level I/O.

Figure 13:   Mini-PCIe Site – PCIe Group 0 Low – PCIe Lane 1
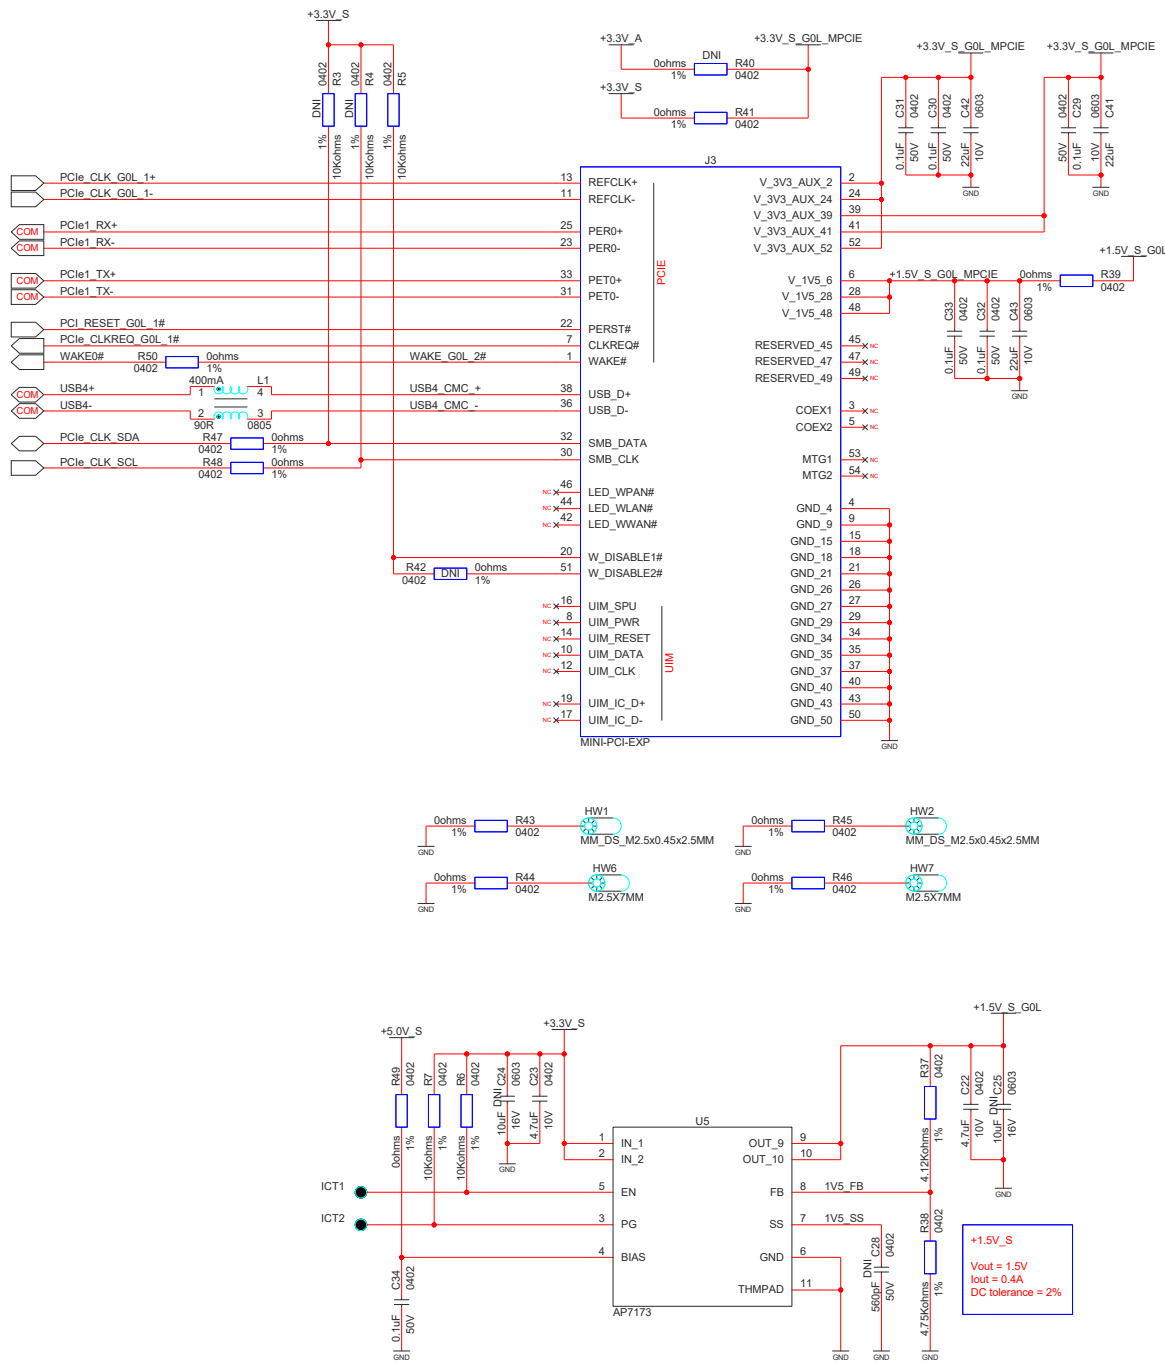
Figure 12 above applies to the COM-HPC Mini as well, except that series coupling capacitors are needed for the PCIe TX pair at the left portion of the Figure. Coupling caps for the PCIe RX pair are on the Mini-PCIe card.

There may need to be level translation on the some of the single ended sideband signals (e.g. WAKE0# etc.) if they destined for a 1.8V level I/O.

Figure 14: M.2 E-Key Site – WiFi Cards – PCIe Group 0 Low – PCIe Lane 2
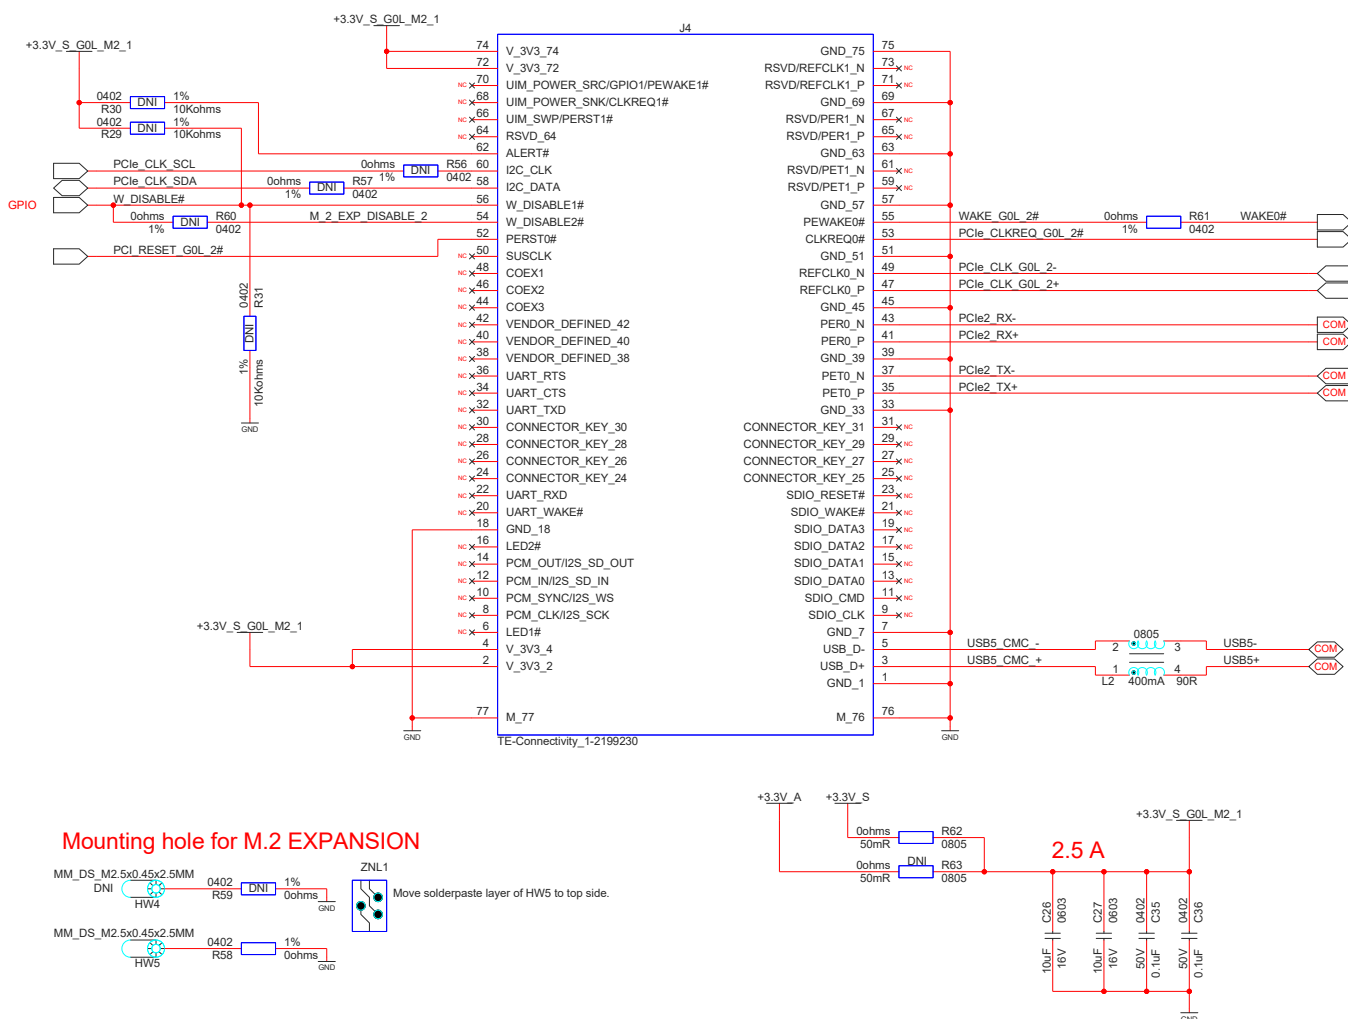
Figure 14 above applies to the COM-HPC Mini as well, except that series coupling capacitors are needed for the PCIe TX pair at the right side of the Figure 14. Coupling caps for the PCIe RX pair are on the M.2 card.

The single ended sideband signals will need level translation if they are sourced by or destined for the COM-HPC Mini Module. The PCIe Reset signal into the M.2 card needs to be at a 3.3V logic level.

Figure 15:   M.2 B-Key Site – Cell Modem Cards – PCIe Group 0 Low – PCIe Lane 3
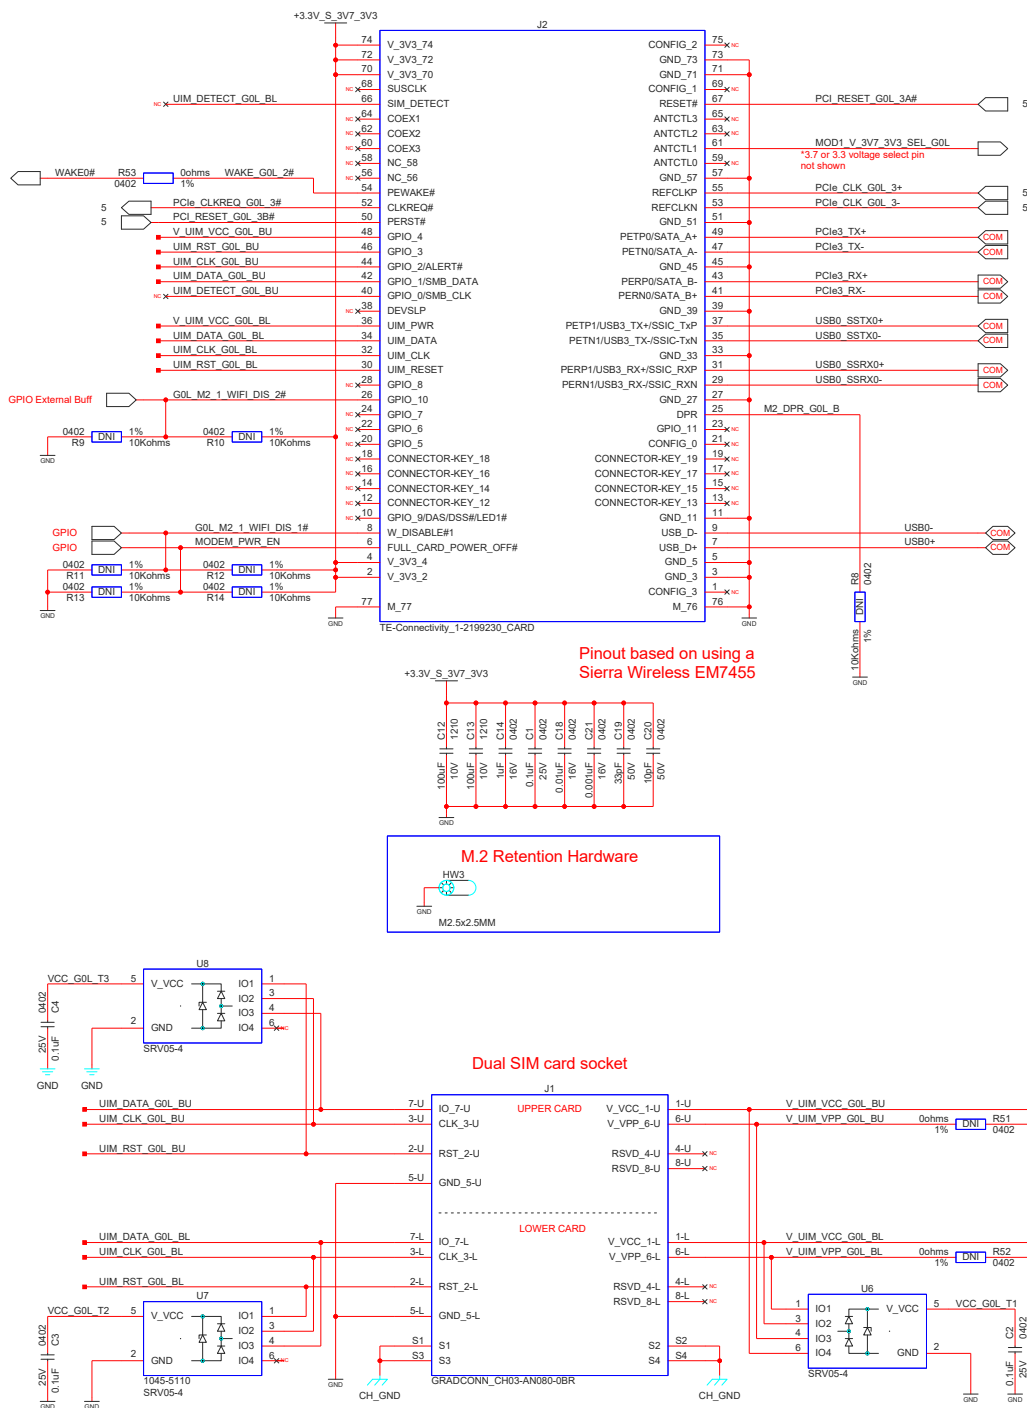
This example may be used on a COM-HPC Mini design, but Carrier board coupling capacitors have to be added for the PCIe TX pair shown at the right side of the Figure 15.

The single ended sideband signals will need level translation if they are sourced by or destined for the COM-HPC Mini Module. The PCIe Reset signal into the M.2 card needs to be at a 3.3V logic level. Other sideband signals may need level translation.

Figure 16:   Clock and Reset Buffers for PCIe Group 0 Low Example Circuits
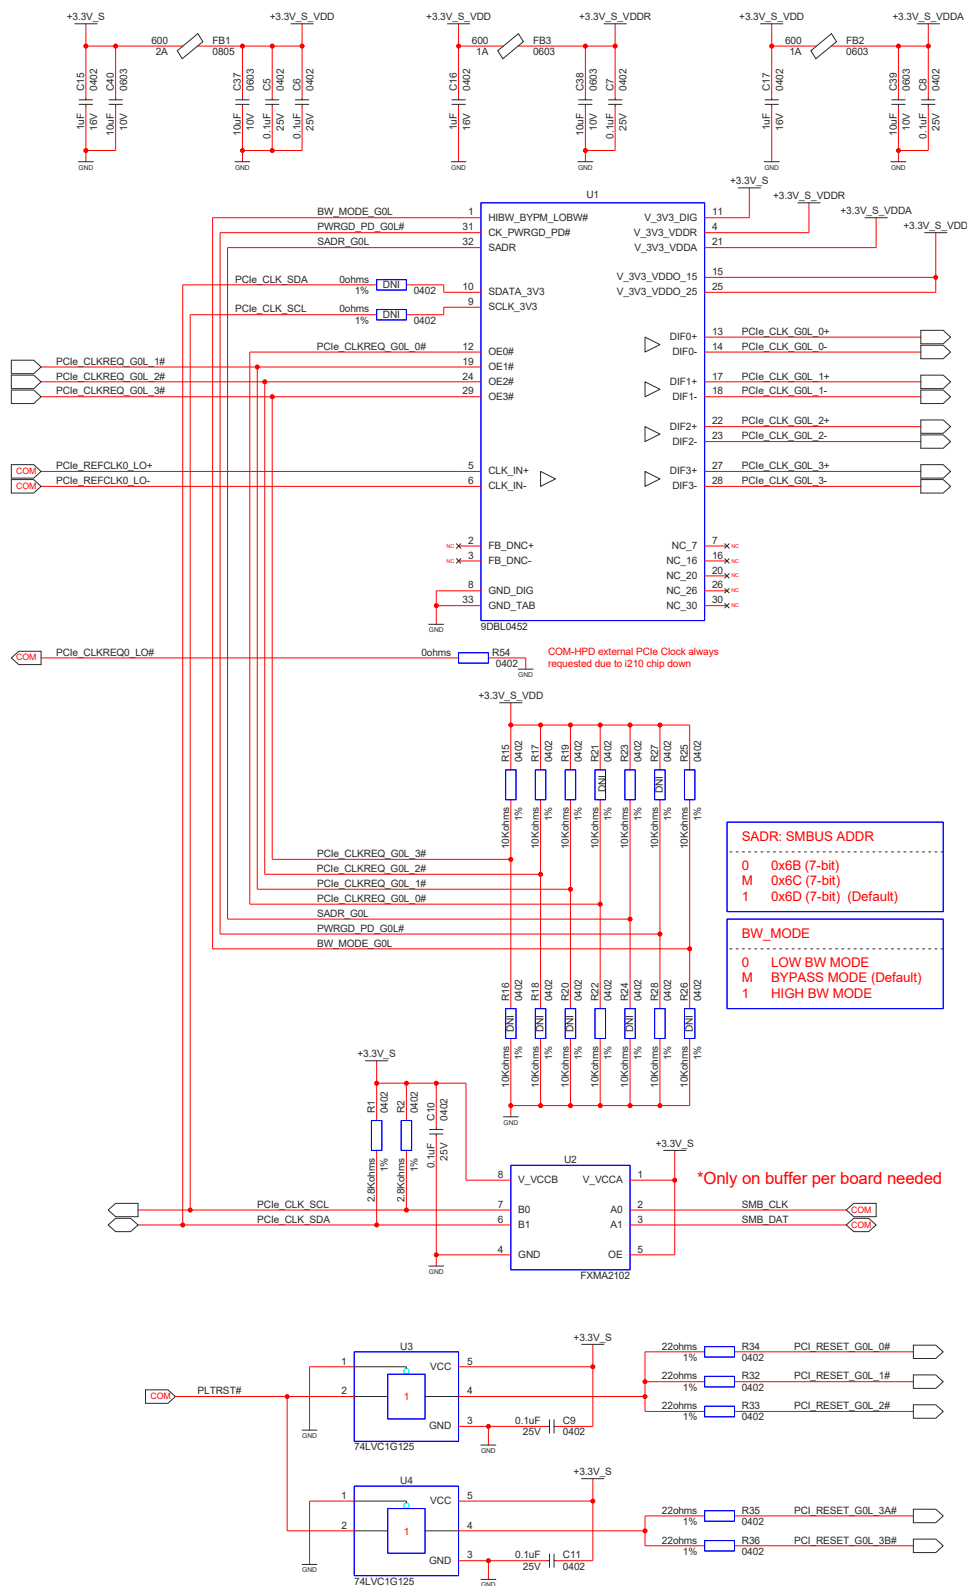
This example might be used in a COM-HPC Mini system, with some modifications – the circuits at the bottom of the page (U2, U3, U4) need to become 1.8V compliant to interface to the COM-HPC Mini Module.

It may be preferable but not strictly necessary to use a 1.8V I/O compliant clock buffer for COM-HPC Mini

designs. For example, the Diodes Inc PI6CB18200. Diodes Inc, Parade Technology, Renesas, Texas Instruments and others may have suitable parts.

The Mini PCIe Reference Clocks do not need level translation. They come out of the Mini at the requisite PCI SIG defined levels.

The PLTRST# signal comes out of the Mini at a 1.8V logic level and needs level translation if it is to drive 3.3V I/O level loads.

3.6.4. Dual PCIe x4 M.2 M Key NVME SSDs Examples on PCIe Group 0 High

The following three figures 17, 18 and 19 illustrate a dual PCIe x4 M.2 M Key deployment for NVME SSDs. The COM-HPC specification recommends that PCIe Group 0 High be used for this purpose.

Figure 17:   M.2 M-Key Site for NVME SSD Card #1 in Group 0 High PCIe Lanes 8:11
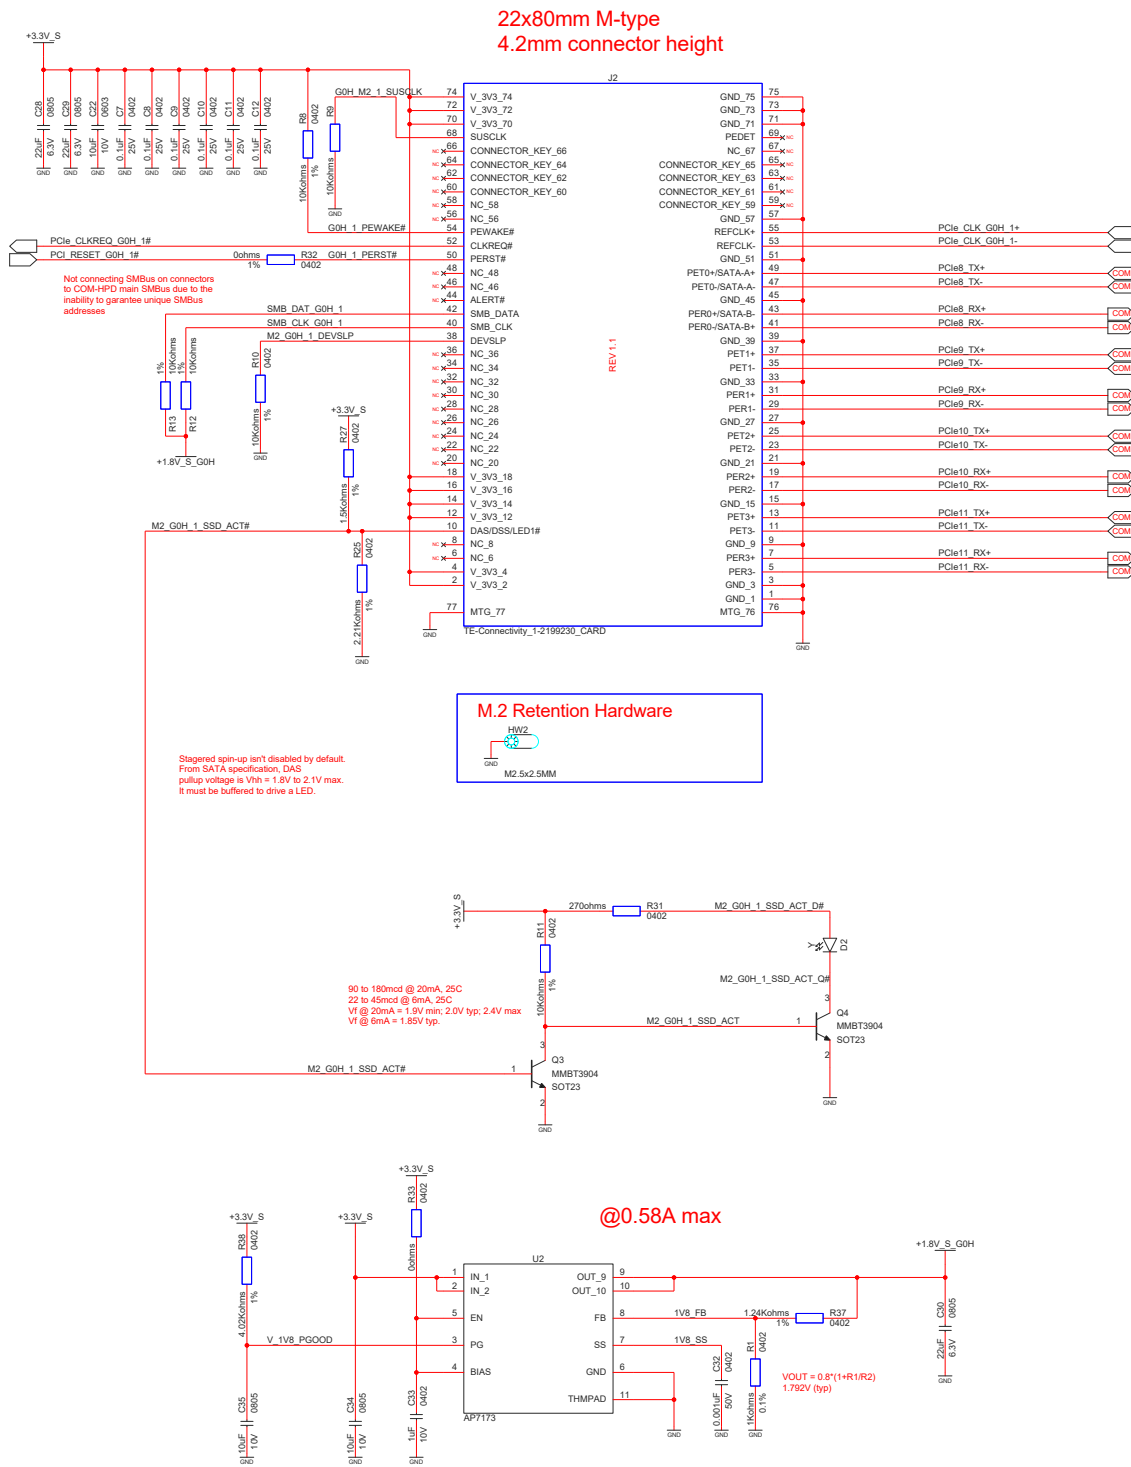
This example may be used on a COM-HPC Mini designs, using PCIe Group 0 High. The PCIe_CLKREQ_G0H_1# signal and the PCI_RESET_G0H_ 1# signal need level translation if they are sourced from or destined to the Mini Module.

Figure 18: M.2 M -Key Site for NVME SSD #2 in Group 0 High PCIe Lanes 12:15

The M.2 M-Key connector used must have a PCIe "speed rating" at least as high as the link speed being used (PCIe Gen 3, 4 or 5).

PCIe Gen 3 and Gen 4 capable M.2 connectors are common. At the time of this writing, PCIe Gen 5 capable M.2 connectors are not common. PCIe Gen 3, 4, and 5 capable M.2 connectors are available from Amphenol FCI.

This example may be used on a COM-HPC Mini design. The PCIe_CLKREQ_G0H_2# signal and the PCI_RESET_G0H_ 2# signal need level translation if they are directly sourced from or destined to the Mini Module.

Figure 19: *Clock Buffer and Reset for PCIe Dual M.2 NVME SSD PCIe Group 0 High*
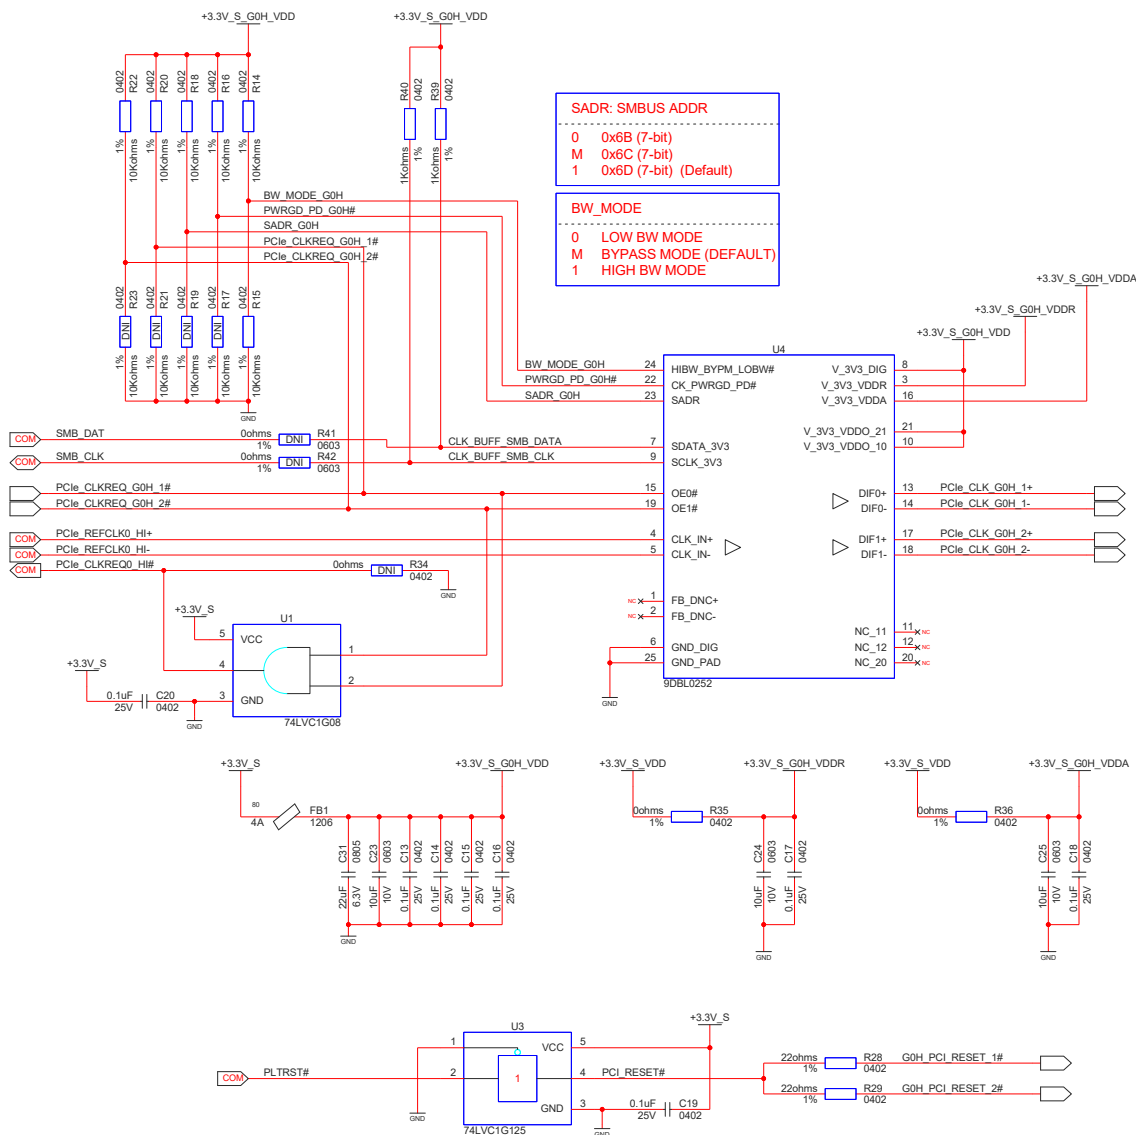
A dual channel clock buffer is used as there are two PCIe x4 links implemented in this PCIe Group 0 High example. The COM-HPC PCIe_CLKREQ0_HI# signal is driven by logic gate U1 in the Figure 19 above, resulting in a clock request if either one or both of the NVMe cards are present. Alternatively, U1 could be removed and the COM-HPC Group 0 High clock request line held low by R34, in which case the Group 0 High PCIe clock pair would always run. The clock buffer shown is the 9DBL0252 from Renesas / IDT. It is PCIe Gen 1,2,3,4 and 5 capable.

The PLTRST# buffer shown, U3, is a 74LVC1G125 device that tolerates a signal input between 0 and 5.5V even in the absence of the VCC to the device. The PLTRST# signal is in the S5 power domain; the VCC applied to U3 is in the S0 domain.

This example is not directly applicable to the COM-HPC Mini, due to most of the Mini I/O voltage levels being defined to be at 1.8V on the Mini rather than 3.3V. The concepts above can be converted to a 1.8V implementation. It may be preferable but not strictly necessary to use a 1.8V I/O compliant clock buffer for COM-HPC Mini designs. For example, the Diodes Inc PI6CB18200. Diodes Inc, Parade Technology, Renesas, Texas Instruments and others may have suitable parts.

3.6.5. PCIe x16 Slot Card Site on PCIe Group 1

Figure 20: *PCIe x16 Slot Card Site on PCIe Group 1 PCIe Lanes 16:31*
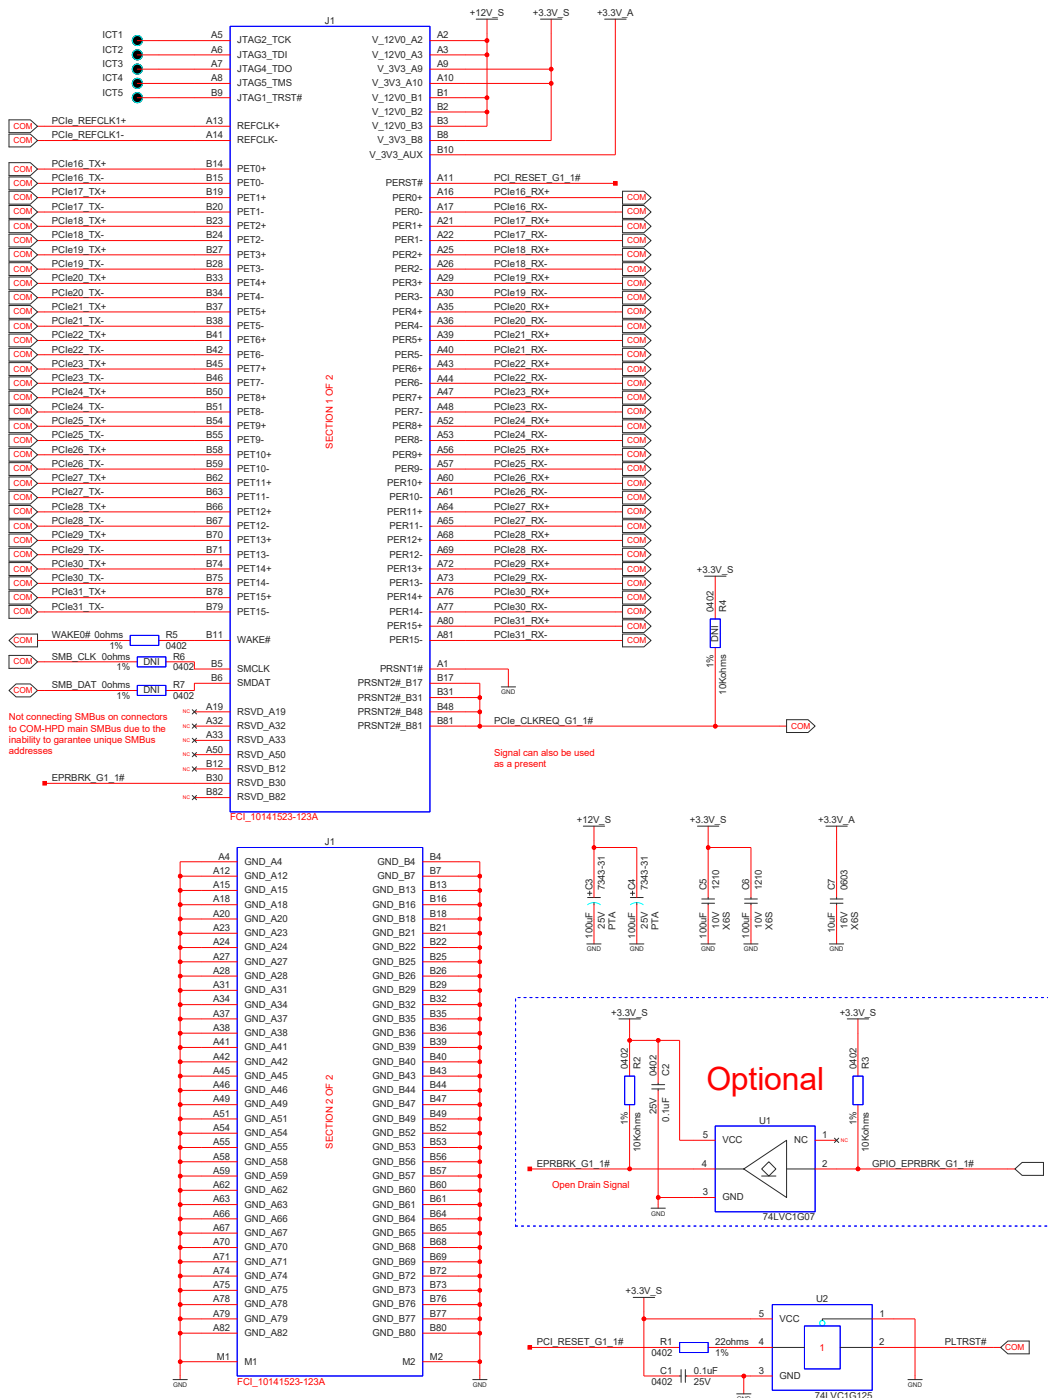
No PCIe clock buffer is needed as there is only one PCIe link in this example. The COM-HPC Group 1 PCIe clock pair is used directly. If the group is split into two or more links, then a PCIe clock buffer would be required.

The slot connector used must be rated and qualified for the PCIe link speed expected. Slot connectors rated for PCIe Gen 3 and 4 are common. Connectors rated for PCIe Gen 5 are at the time of this writing are still new. Such parts are available from Amphenol FCI and others.

This example does not directly apply to the COM-HPC Mini as there is no PCIe Group 1 on the Mini.

3.6.6. PCIe Group 2

Figures 21 through 24 below illustrate the implementation of 3 PCIe slots on COM-HPC PCIe Group 2, along with a PCIe clock buffer appropriate for use with PCIe Gen 4 and below. The slot RESET# signals come from a buffer in Figure 24. Additional notes are provided after the last Figure in this series.

These examples do not apply to the COM-HPC Mini.

Figure 21: PCIe x8 Slot Card Site on PCIe Group 2 PCIe Lanes 32:39
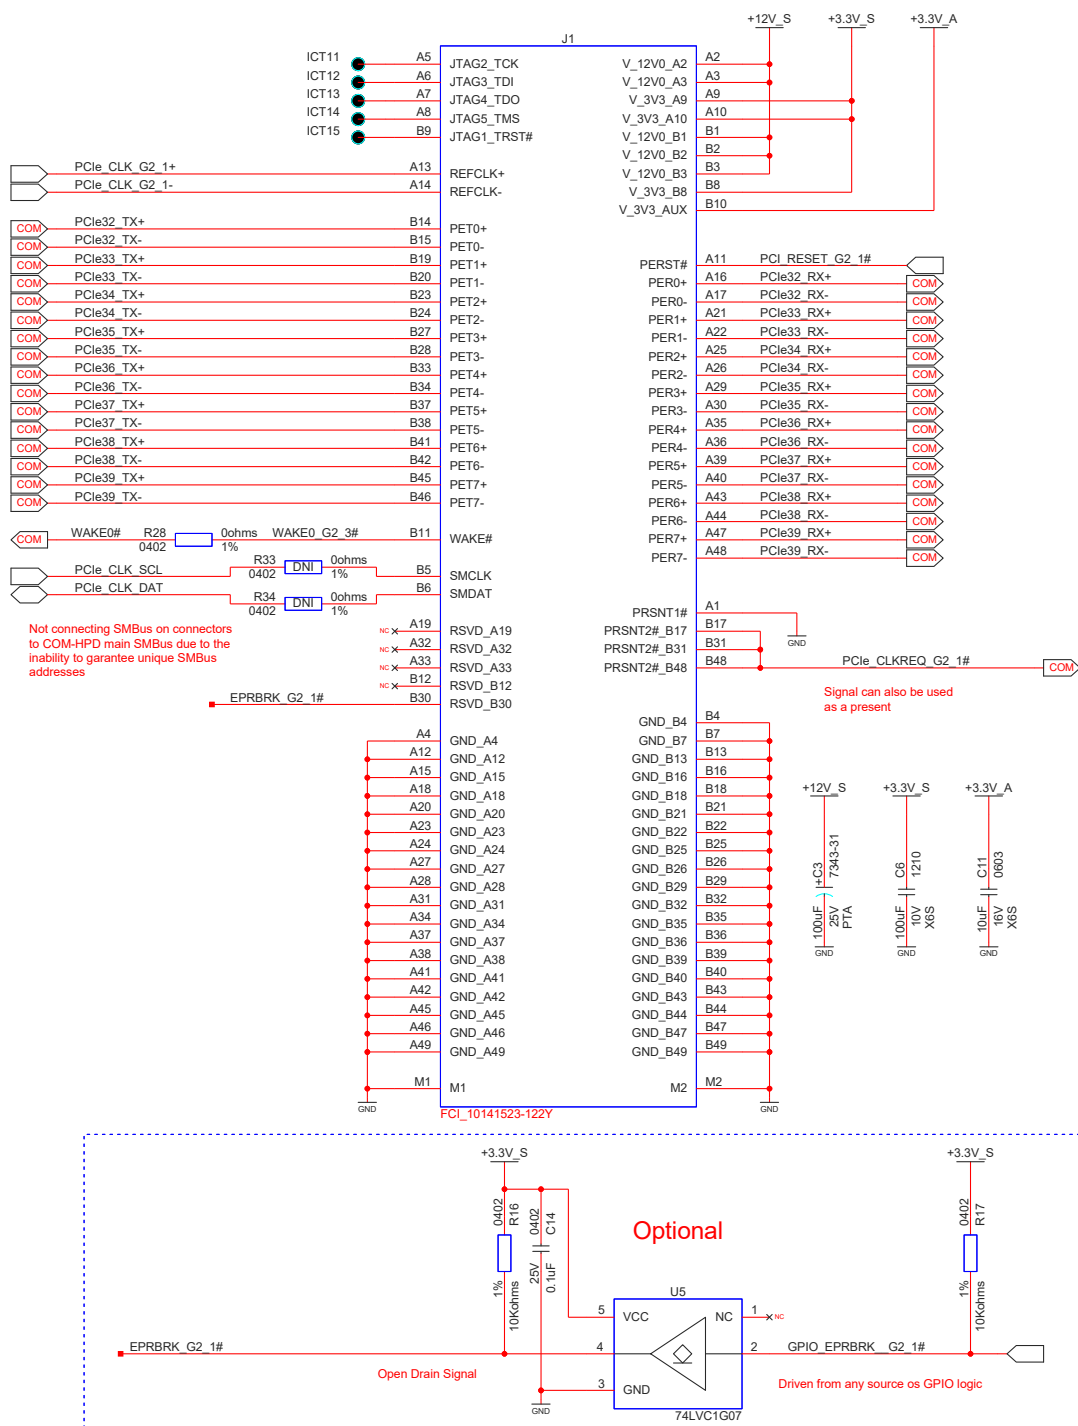
Figure 22:   PCIe x4 Slot Card Site on PCIe Group 2 PCIe Lanes 40:43
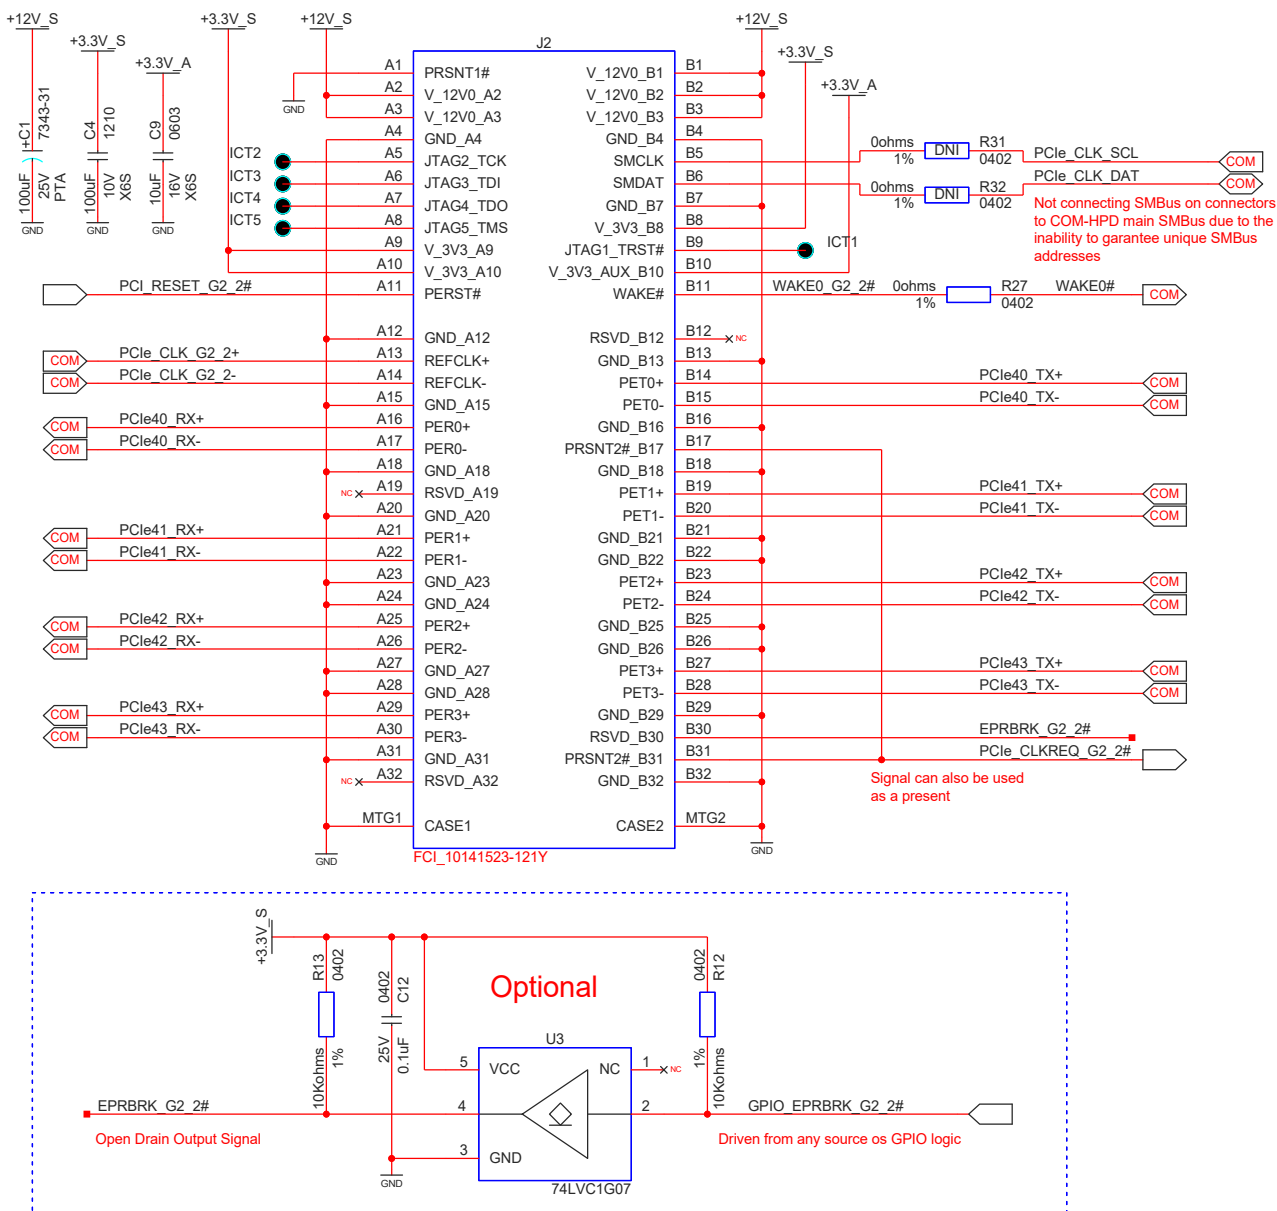
This example does not apply to the COM-HPC Mini.

Figure 23: PCIe x4 Slot Card Site on PCIe Group 2 PCIe Lanes 44:47
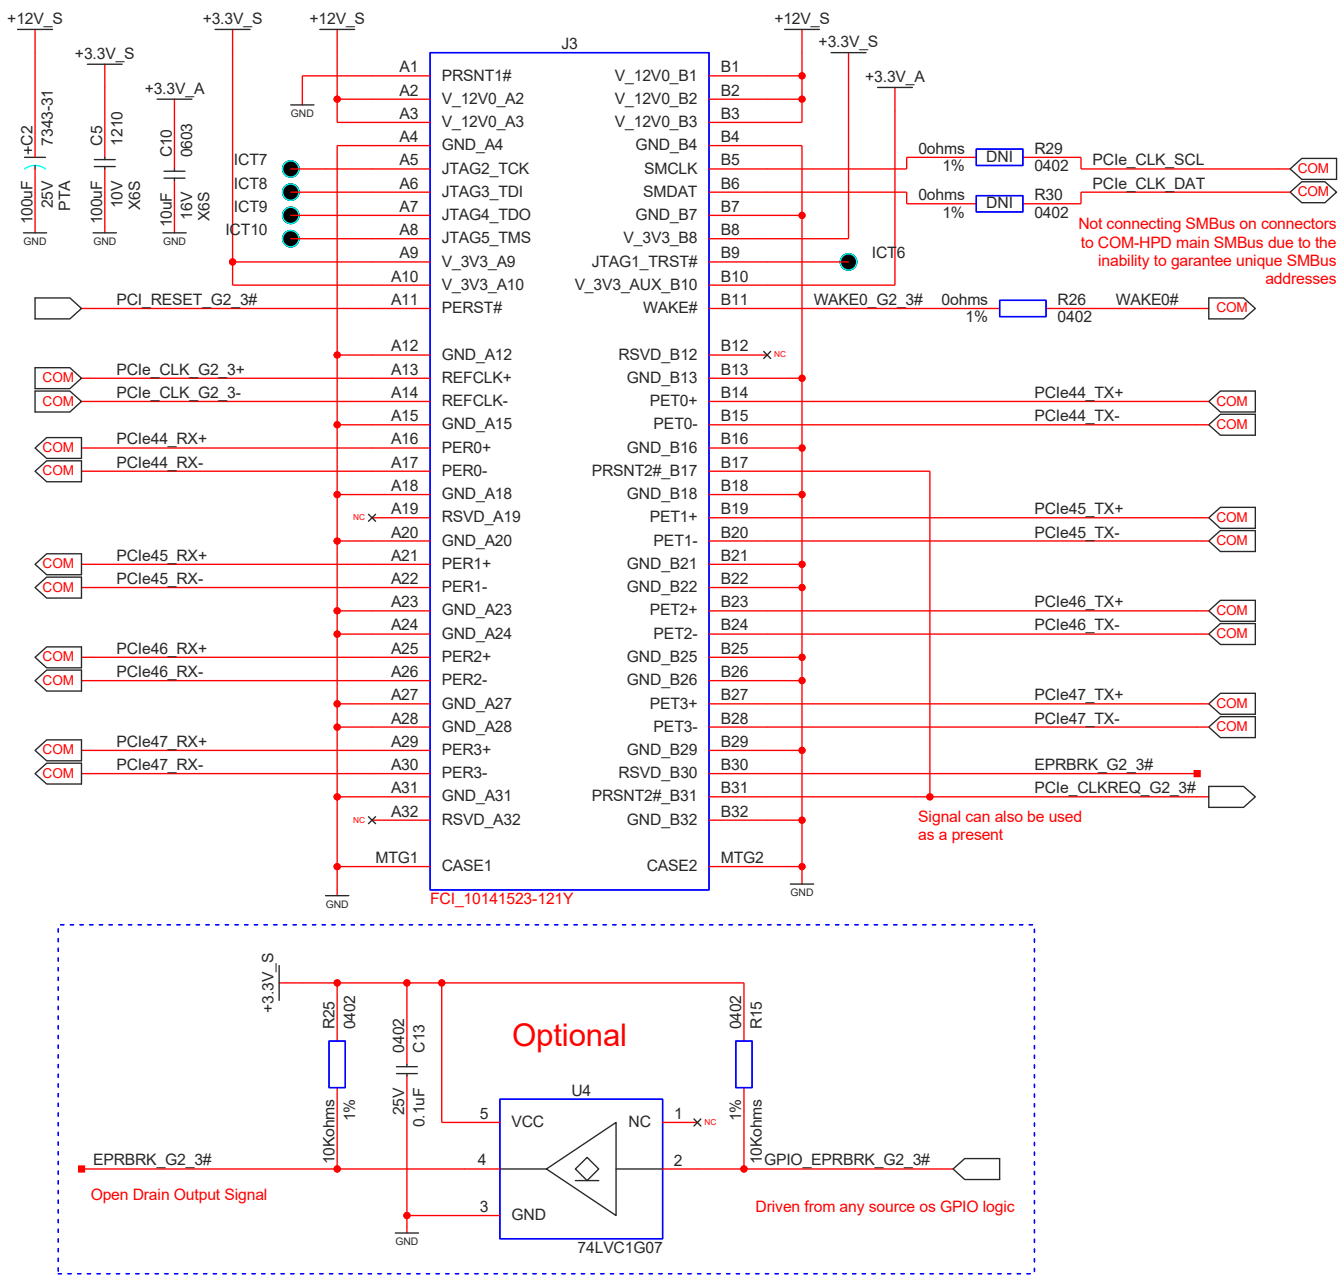
This example does not apply to the COM-HPC Mini.

Figure 24: PCIe Clock Buffer and Reset Buffer for PCIe Group 2 Example
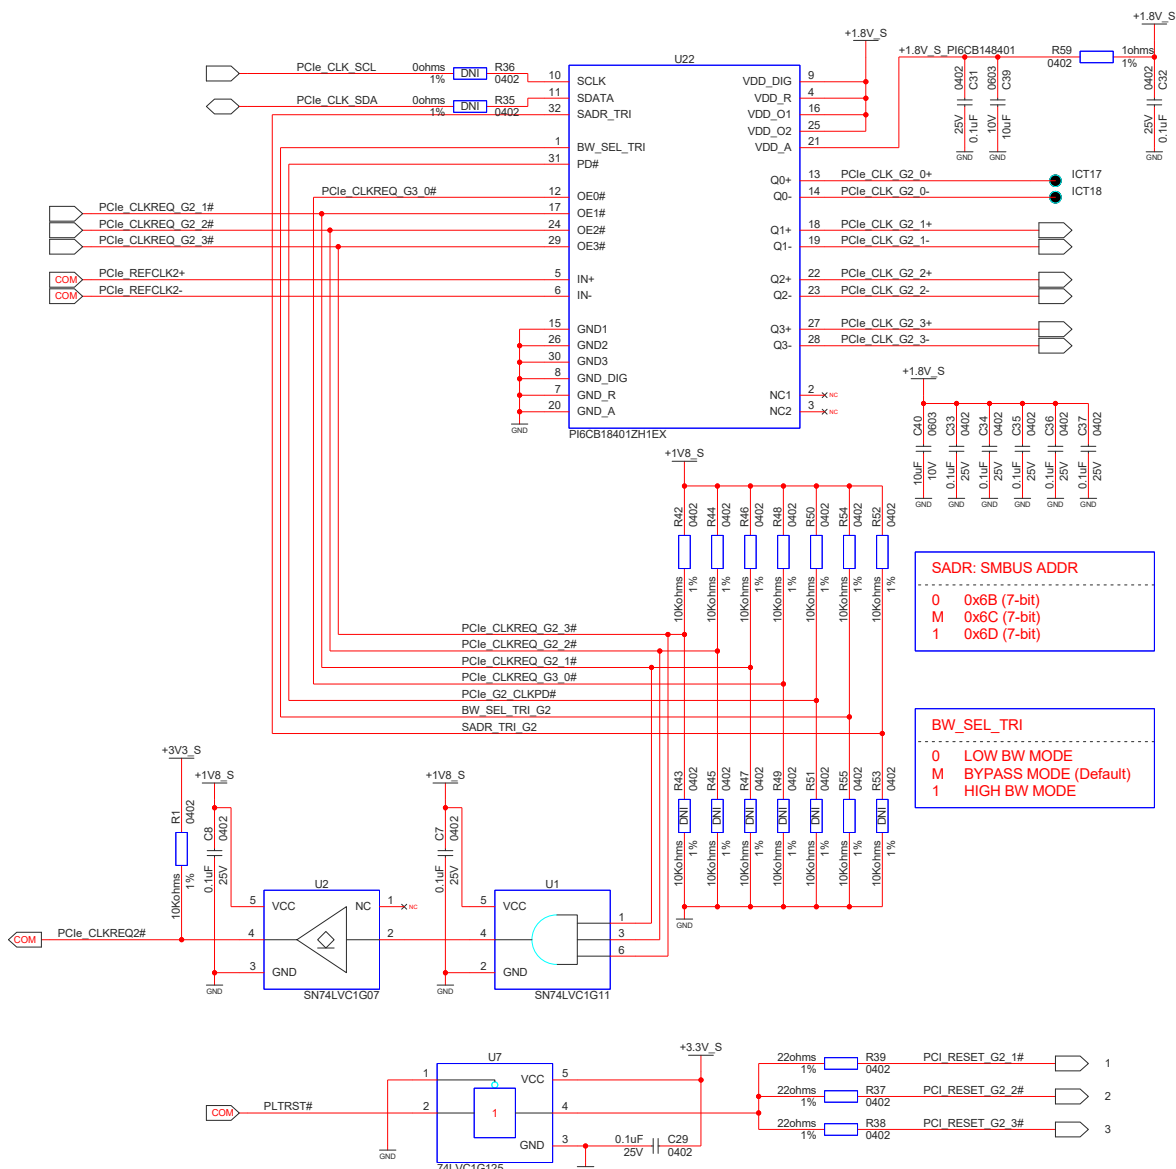
*Max 4 outputs per buffer
*Push/Pull design use LVC07 wih PU resistor for OC designs

This example does not apply to the COM-HPC Mini.

3.6.7. MXM-3 Graphics Card Module on Carrier

The **COM Express Carrier Design Guide** Rev 2.0 Section 2.6 has a good schematic example of a MXM-3 graphics card implementation on COM Express Carriers. The net names used are not the same as the COM-HPC net names, but the correlation is straightforward.

MXM-3 cards use a 16 lane PCI Express interface. Most MXM cards also allow x8 implementations (and all PCIe devices must work in x1 configurations per the PCI-SIG specification). COM-HPC recommends but does not require that COM-HPC PCIe Group 1 be used for PEG (PCI Express Graphics). No Carrier PCIe clock buffer is needed for the MXM card (assuming that the MXM card is only PCIe device used in the COM-HPC PCIe Group). The COM-HPC PCIe reference clock that goes with the COM-HPC PCIe group can be used directly with the MXM card.

Coupling capacitors for the COM-HPC Module PCIe RX pairs (MXM card PCIe TX pairs) must be present on the Carrier, preferably close to the MXM connector. Use discrete 0402 or 0201 package size parts.

The MXM-3 specification document is currently hard to find. The document was created and is owned by Nvidia but is not publicly available. Information found in the **COM Express Carrier Design Guide**, from the MXM-3 connector vendors (Aces, Amphenol / FCI, Foxconn, JAE, Yamaichi), from the MXM GPU card vendors and from the COM-HPC Module vendor should be sufficient to carry out a design.

The Amphenol / FCI MXM-3 connector part number 10151114-001TLF supports all PCIe signal rates up to and including Gen 5.

Note that MXM cards operate in the S0 power domain only. Any signals that are active in the S5 (suspend) power state must be isolated, in the S5 power state, from the MXM card.

Note also that MXM-3 connectors have 314 individual pins, but the MXM-3 specification gangs multiple connector pins together for power delivery. The 314 individual connector pins are grouped together in the NVIDIA MXM-3 specification into PWR and GND blocks labeled E1, E2, E3 and E4, and then the pins left over are numbered 1 through 281. The MXM-3 connector drawings from the connector vendors usually illustrate this.

MXM graphics cards may be used in a system with the COM-HPC Mini, but series coupling capacitors for any PCIe Group 0 Low TX pairs and all of the PCIe RX pairs (Group0 Low and High) must be present on the Carrier. They are not present on the MXM card. Any Carrier TX caps should be placed close to the COM-HPC Module pins. The RX caps (RX relative to the COM-HPC Module; TX from MXM) should be placed close to the relevant MXM pins.

I/O signals between the MXM card and the COM-HPC Mini may need to be level translated for COM-HPC Mini signals that are at a 1.8V level.

3.6.8. PCIe Reference Clocks

COM-HPC, like COM Express and like most commercial and embedded PCI Express implementations, uses the "Common Clock" architecture defined in the PCI-SIG **PCI Express Base Specification**, Revisions 1 through 5**.** In this arrangement, there is a common 100 MHz reference clock source that feeds the PCIe Root complex and copies are fed to each PCIe Target device serviced by that Root. The maximum skew allowed between any two PCIe reference clocks, at their destinations, is 10 ns for PCIe Gen 1 and 12 ns for Gen 2 through 5.

In most cases, the 100 MHz reference source is integrated into the SOC or chipset and the reference clock routing to the Root is internal to the SOC or chipset. In some cases, a clock generator IC that is external to the SOC or chipset is used. In any case, the 100 MHz reference source for a COM-HPC Root device is on the COM-HPC Module, either internal or external to the SOC or chipset.

The SOC or chipset may provide one or multiple copies of the PCIe reference clock. If the SOC, chipset or Module does not provide enough copies of the reference clock for the Carrier PCIe targets then one or more PCIe clock buffers are used, on and / or off Module, depending on the situation.

The PCIe targets use the 100 MHz reference clock copy, along with the clocking information embedded into the PCIe data stream to quickly form a local copy of the appropriate high frequency clock (2.5 GHz for Gen1, 5 GHz for Gen 2, and so on) needed to correctly interpret the incoming PCIe data stream, and to correctly time and encode the target's outgoing data stream.

Table 6 below lists the maximum clock jitter allowed for each PCIe generation, per the PCI-SIG source specifications, for the Common Clock architecture. Note the ever shrinking jitter allowance as the generations advance. For example, the Gen 5 jitter allowance is only 15% of the Gen 3 allowance. However, PCIe Gen 5 uses a different filtering transfer function than Gen 3 and Gen 4, so the comparison is more nuanced than indicated here.

Table 6: PCIe Maximum Allowable Clock Jitter

PCIe Generation	Signaling Rate	Reference Clock Max Jitter Allowed	Notes
1	2.5 Gbps	86 ps PTP	PTP is Peak to Peak
2	5.0 Gbps	3.1 ps RMS	RMS is Root Mean Square
3	8.0 Gbps	1.0 ps RMS	
4	16.0 Gbps	0.5 ps RMS	
5	32.0 Gbps	0.15 ps RMS	
6	64.0 Gbps	0.10 ps RMS	

Table 7 on the following page defines some PCIe Clock Buffer mode terminology.

Table 7: PCIe Clock Buffer Modes

Mode	Description	Pros and Cons
Zero Delay	A Clock Buffer PLL keeps the output clock copies in phase with the input clock. Also known as ZDB (Zero Delay Buffer) mode or as PLL mode. Some clock buffers have High Bandwidth and Low Bandwidth PLL modes.	**Pro:** Zero delay makes it easier to meet the maximum skew of 12 ns between any two clocks, especially if the PCIe traces are very long, or if buffers are cascaded The PLL tends average out the jitter contribution from the source clock (see jitter discussion later in this document section, following Table 8) **Con:** PLLs buffers may have trouble with a Spread Spectrum source (see discussion following Table 8)
Fan Out or PLL Bypass	No PLL used. The output clock copies are an exact frequency copy of the input but are not in phase with the input.	**Pro:** lower jitter from the buffer itself, in most cases (but the source clock jitter must be added to the that of the fan out buffer, per discussion following Table 8) Fan out buffers track a SpreadSpectrum clock source easily **Con:** may be harder to meet 12 ns max clock skew Although the fan out buffer jitter itself is low, the source clock jitter adds to the fan out buffer jitter

Timing Delay Discussion For Various PCIe Clock Buffer Scenarios

Regarding the max PCIe Reference Clock skew of 12 ns (or 10 ns for PCIe Gen 1) and the use of Fan Out (non – PLL) based clock buffers: a modern Fan Out buffer will have a worst case skew of well under 5 ns (several vendors claim 3 ns max, and at least one claims 1.5 ns max). Signals propagate at about 6 inches per ns, so a system with a 5 ns buffer delay and about 12 inches of PCB trace (2 ns delay) would have a worst case skew of 7 ns which is comfortably within the 12 ns Gen 2 through Gen 5 skew limit. Very long PCIe trace situations might need the Zero Delay Buffer – be aware of the possible issues with spread spectrum sources.

However .. if there is a clock buffer on the COM-HPC Module in-between the Root complex PCIe reference clock and the clock(s) going out to the COM-HPC pins, then there will likely be an additional delay time that factors into the analysis in the previous paragraph. Check with your Module vendor on that. As: "what is the skew between the PCIe Reference Clock to the CPU or SOC Root Complex, and the COM-HPC PCIe Clock Reference pins" … it could be anywhere from 0 ns to 5 ns, depending on Module design details. Also: "what is the jitter contribution of a Module PCIe clock buffer" - if there is one.

PCIe Clock Buffer Options – Keep Them Open

A PCIe Clock Buffer IC usually has pin-strap(s) and / or SMBus options allowing the Clock Buffer operational modes to be set. It is best to keep access to these options open as sometimes issues can be resolved late in the design cycle (i.e. during regulatory and compliance testing) by changing the operational mode of the PCIe Clock Buffer. For example, Spread Spectrum PCIe reference clock operation may work with some but not all of the Clock Buffer modes. Additionally, some PCIe Clock Buffer devices have mechanisms (such as SMBus registers or OTP ordering options) to change parameters such as output clock slew rate, signal amplitude and / or the output termination values.

Sample PCIe Clock Buffer List

A sample collection of PCIe Clock Buffers is given in Table 8 below. Of course this is just a snapshot of what is available and appropriate at the time of this writing. Fan Out Mode jitter is additive (meaning the Module source jitter needs to be added together with the Fan Out buffer jitter). PLL Mode jitter is not additive, hence is marked as Total in the Table. The Module source jitter tends to get averaged out in the PLL. This is described in more detail on the page following Table 8.

Table 8: PCIe Clock Buffer Vendors and Part Numbers

Vendor	Part Numbers	Notes	RMS Jitter (picoseconds)
Diodes Inc. (Pericom)	PI6CB18200 (dual, no internal term) PI6CB18401 (quad, internal term) PI6CB18601 (hex, internal term) PI6CB18801 (octal, internal term) An internal termination value for 100 ohm differential traces is implied in the data sheet but not explicitly stated.	PCIe Gen 4 capable 1.8V supplies OE# on each output SMBus configuration option Pin strap configuration option Zero Delay Buffer modes High BW PLL Low BW PLL PLL Bypass (aka Fan Out) Mode	**PLL Mode (Total)** Gen 1 5.0 Gen 2 0.3 Gen 3 0.1 Gen 4 0.05 **Fan Out Mode** Values not shown in public data sheet
Diodes Inc. (Pericom)	PI6CB33202 (dual, 85 ohm internal term) PI6CB33402 (quad, 85 ohm internal term) PI6CB33602 (hex, 85 ohm internal term) PI6CB33802 (octal, 85 ohm internal term) PI6CB33201 (dual, 100 ohm internal term) PI6CB33401 (quad, 100 ohm internal term) PI6CB33601 (hex, 100 ohm internal term) PI6CB33801 (octal, 100 ohm internal term)	PCIe Gen 5 capable 3.3V power supply OE# on each output SMBus configuration option Pin strap configuration option Zero Delay Buffer modes High BW PLL Low BW PLL PLL Bypass (aka Fan Out) Mode	**PLL Mode (Total)** Gen 1 0.05 Gen 2 0.05 Gen 3 0.05 Gen 4 0.05 Gen 5 0.05
Diodes Inc. (Pericom)	PI6CB332001A (20 outputs, 85 ohm internal term)	PCIe Gen 5 capable 3.3V power supply OE# for 8 outputs SMBus, Side-Band interface support 20 HCSL outputs with On-chip Termination	**Fan Out (Additive)** Gen 1 0.03 Gen 2 0.03 Gen 3 0.03 Gen 4 0.03 Gen 5 0.12
Renesas (IDT)	9DBL0252 (dual, 85 ohm internal term) 9DBL0452 (quad, 85 ohm internal term) 9DBL0651 (hex, 85 ohm internal term) 9DBL0851 (octal, 85 ohm internal term) 9DBL0242 (dual, 100 ohm internal term) 9DBL0442 (quad, 100 ohm internal term) 9DBL0641 (hex, 100 ohm internal term) 9DBL0841 (octal, 100 ohm internal term)	PCIe Gen 5 capable 3.3V power supplies OE# on each output SMBus configuration option Pin strap configuration option Zero Delay Buffer modes High BW PLL Low BW PLL PLL Bypass (aka Fan Out) Mode	**Fan Out (Additive)** Gen 1 5.0 Gen 2 0.428 Gen 3 0.149 Gen 4 0.156 Gen 5 0.05 **PLL Mode (Total)** Gen 1 33 Gen 2 1.9 Gen 3 0.53 Gen 4 0.48 Gen 5 0.149

Vendor	Part Numbers	Notes	RMS Jitter (picoseconds)
Renesas (IDT)	9DBL0255 (dual, 85 ohm internal term) 9DBL0455 (quad, 85 ohm internal term)	Ultra low jitter PCIe Gen 5 capable 100 ohm option with ext resistors 3.3V power supplies OE# for each output	**Fan Out (Additive)** Gen 3 0.033 Gen 4 0.033 Gen 5 0.012
Renesas (IDT)	9ZXL0451E (quad, 85 ohm internal term) 9ZXL0651E (hex, 85 ohm internal term) 9ZXL0851E (octal, 85 ohm internal term) 9ZXL1251E (12 out, 85 ohm internal term)	PCIe Gen 5 capable 3.3V power supplies OE# on each output SMBus configuration option Zero Delay Buffer modes High BW PLL Low BW PLL PLL Bypass (aka Fan Out) Mode	**Fan Out (Additive)** Gen 1 1.9 Gen 2 0.126 Gen 3 0.062 Gen 4 0.062 Gen 5 0.024 **Low BW PLL Mode (Total Jitter)** Gen 1 6.8 Gen 2 0.12 Gen 3 0.07 Gen 4 0.07 Gen 5 0.018
Skyworks (Silicon Labs)	Si53204-A02 (quad, 85 ohm internal term) Si53208-A02 (octal, 85 ohm internal term) Si53212-A02 (12 out, 85 ohm internal term) Si53204-A01 (quad, 100 ohm internal term) Si53208-A01 (octal, 100 ohm internal term) Si53212-A01 (12 out, 100 ohm internal term) Silicon Labs has many other PCIe Clock Buffers, too numerous to list here.	PCIe Gen 5 capable 1.8V power supplies OE# on each output SMBus configuration option Fan Out Mode only (no PLL)	**Fan Out (Additive)** Gen 1 17 (PTP) Gen 2 0.2 Gen 3 0.06 Gen 4 0.06 Gen 5 0.021
Texas Instruments	LMK00334 (quad output, external term)	PCIe Gen 4 capable 3.3V and 2.5V supplies Single OE# Fan Out Mode only	**Fan Out (Additive)** Gen 3 0.15 Gen 4 0.05
Texas Instruments	LMK00338 (octal output, external term)	PCIe Gen 3 capable 3.3V and 2.5V supplies Single OE# Fan Out Mode only	**Fan Out (Additive)** Gen 3 0.15
Texas Instruments	CDCB2000 (20 outputs, 85 ohm int term)	PCIe Gen 5 capable 3.3V supplies OE# for 8 outputs SMBus configuration option Side Band Interface config option	**PLL Mode (Total)** Gen 1 5.0 Gen 2 0.2 Gen 3 0.15 Gen 4 0.08 Gen 5 0.03
Texas Instruments	CDCDB800 (octal output, 85 or 100ohm software selectable term)	PCIe Gen 5 capable 3.3V supplies OE# on each output SMBus configuration options Fan Out Mode only Propagation delay 0.5 ns typically 3 ns max	**Fan Out (Additive)** Gen 3 0.1 Gen 4 0.1 Gen 5 0.025

There may be more subtleties in the jitter numbers than is immediately apparent here. For example, some buffers allow the clock output slew rate to be adjusted, but the slew rate may in turn affect the jitter values. Check the vendor data sheets and make use of the vendor application engineers.

Note on the Additive and Total Jitter values in Table 8 Above:

These values are taken from silicon vendor data sheets and are meant here as a rough guide. The jitter values for all operational modes (e.g. high PLL BW, low PLL BW etc) of the clock buffer devices may not be shown here. The values in the Table for Fan Out buffers are the "Maximum Additive Jitter" values listed in the vendor data sheets. For PLL buffers, the typical values are usually much lower, often less than half, of the maximum values. Jitter analysis can be tricky. Designers should consult the actual vendor data sheets and vendor application notes before making design decisions.

Note on Fan – Out Buffer Jitter vs. PLL or Zero Delay Buffers

If a Fan-Out mode PCIe clock buffer is used, then the clock jitter at the target is the square root of the sum of the squares of the COM-HPC Module clock source and of the Fan-Out buffer jitter, per the expression shown here:

$$\sqrt{J^2_{source} + J^2_{buffer}}$$

For PLL or Zero – Delay mode buffers, the clock jitter at the target is simply the jitter of the PLL buffer as listed in the vendor data sheet. The PLL buffer tends to average out the source clock jitter, unless it is extreme.

In other words, with regard to the jitter values shown in Table 8 above, the Fan-Out buffer jitter values are not the full story, as the clock generator source jitter values need to be factored in per the equation above. The source clock generator jitter values need to be obtained from the Module vendor or the SOC vendor data sheets.

Note on the Internal Termination Impedances in Table 8 Above:

PCIe Clock Buffer internal output termination values suitable for both 85 ohm and 100 ohm differential traces are shown as being available in Table 8 above. The **COM-HPC Base Specification** recommends an 85 ohm differential impedance for the PCIe Reference Clocks coming off the Module, and hence into the Carrier Clock Buffer. Designers are free to choose either 85 ohms or 100 ohm differential impedances for their Clock Buffer output distribution. The phrasing "85 ohm internal term" in the Table above means that the device internal termination is appropriate for 85 ohm differential pairs, and similarly for "100 ohm internal term".

Spread Spectrum Clock (SSC) Operation

SSC profiles for different PCIe generations are different, so a PLL based buffer with support for SSC needs to have an appropriate loop bandwidth for the PFD (Phase Frequency Detector) within the PLL. Therefore, it is advisable to check whether a PLL based buffer supports SSC for the PCIe generation it is to be used with.

Some clock buffer vendors recommend against using a Spread Spectrum Clock source with their PLL mode parts and recommend the use of a Fan Out buffer instead. Check with your clock buffer vendor and allow for a PLL bypass mode (Fan Out Mode) option if possible.

COM-HPC Reference Clocks vs COM Express / Use of Clock Buffers

COM Express Rev 3.0 defines a single PCIe Reference Clock in it's pinout.

COM-HPC Rev 1.0 allows up to five PCIe Reference Clock pairs – there is one COM-HPC Module PCIe Reference Clock pair for each of the five COM-HPC PCIe groups, as outlined in Section 3.6.1. above.

3.6.9.   PCIe Redrivers and Retimers

PCIe maximum trace length guidelines are given in Section 4.3.4. below and are presented along with loss budgets and more context in the **COM-HPC Base Specification**. However, PCIe Gen 3, 4, and 5 implementations may well need a redriver or retimer on the Carrier to make up for signal degradation.

A **redriver** is an analog circuit that reshapes the PCIe signal using sophisticated analog techniques. A nearly closed PCIe signal eye can become a compliant open eye with a redriver. Redrivers may have a digital section in the form of I2C accessible registers or strap pins to set redriver parameters.

A **retimer** is a digital and analog circuit that clocks in the PCIe signal using the PCIe 100 MHz reference clock and an on-chip PLL and reissues the reclocked signal in pristine form. Two popular vendors for PCIe redriver and retimer products are Diodes Inc. (formerly Pericom) (www.diodes.com) and Texas Instruments (www.ti.com). A retimer may possibly yield better results than a redriver, at a cost.

Table 9: PCIe Redrivers and Retimers

Vendor	P/N	Notes
Diodes Inc	PI3EQX16904GL	PCIe Gen 4 capable quad lane **redriver** (4 lanes in one direction)
	PI3EQX16908GL	PCIe Gen 4 capable octal lane **redrive**r (8 lanes in one direction)
Texas Instruments	DS160PR410	PCIe Gen 4 capable quad lane **redriver** (4 lanes in one direction)
	DS160PR810	PCIe Gen 4 capable octal lane **redrive**r (8 lanes in one direction)
	DS160PT801	PCIe Gen 4 capable 16 lane **retimer** (8 lanes TX and 8 lanes RX)
	DS320PR810	PCIe Gen 5 capable octal lane **redrive**r (8 lanes in one direction)
	DS320PR822	PCIe Gen 5 capable quad 2x2 crosspoint **redrive**r

The items shown in Table 9 above represent only a small sample of parts available on the market. Texas Instruments, for example, has quite a few additional redriver and retimer parts not listed here. Some of the unlisted parts incorporate redriver or retimer functions along with analog multiplexer and cross-point switch functions.

3.7. USB

3.7.1. USB Terms and General Information

Table 10: USB.org Branding Term Summary

USB.org Current Branding	Nominal Bit Rates	Notes	USB.org Former Branding
USB 2.0	480 Mbps (High Speed mode)	Single half duplex DC coupled pair Also supports slower USB 1.1 and 1.0 legacy modes	USB 2.0
USB 3.2 Gen 1	5 Gbps	Dual simplex AC coupled transmit pair and a receive pair. Also requires a specific USB 2.0 link, on a separate set of conductors.	USB 3.0 USB 3.1 Gen 1
USB 3.2 Gen 2	10 Gbps	Dual simplex AC coupled transmit pair and a receive pair. Also requires a specific USB 2.0 link, on a separate set of conductors.	USB 3.1 USB 3.1 Gen 2
USB 3.2 Gen 2x2	10 Gbps (per lane) 20 Gbps (two lanes)	Two AC coupled transmit pairs and two receive pairs allowing 20 Gbps operation in each direction. Also requires a specific USB 2.0 link, on a separate set of conductors.	
USB4 Gen 2x2	10 Gbps (per lane) 20 Gbps (two lanes)	Incorporates USB 3.2 Gen 2x2 and USB 2.0 features, along with additional features such as DisplayPort operation, USB Type-C (reversible) connector, and Thunderbolt 4 support.	
USB4 Gen 3x2	20 Gbps (per lane) 40 Gbps (two lanes)	Features 20 Gbps bit rate per lane and uses 2 lanes TX and 2 lanes RX. Also includes USB 2.0 features, along with additional features such as DisplayPort operation, USB Type-C (reversible) connector, and Thunderbolt 4 support.	
USB SuperSpeed USB SuperSpeed+	5 Gbps 10 or 20 Gbps	The high speed interface used in USB 3.2 Gen 1, Gen 2, Gen 2x2 and USB4 is referred to as the SuperSpeed or SuperSpeed+ interface. A USB 3.2 Gen 1, Gen 2, Gen 2x2 or USB4 implementation require both SuperSpeed / SuperSpeed+ support and USB 2.0 support. The SuperSpeed / SuperSpeed+ interface is implemented on a separate set of pins from the USB 2.0 interface. However, every SuperSpeed implementation needs a specific companion USB 2.0 interface.	

Actual payload data rates are lower than what is implied by the "Nominal Bit Rates" in the chart above, due to the encoding methods used in the serialized data stream.

The most common connector for USB host ports is the Type-A connector. Figure 25 below is a view looking into a USB 3 Type-A host receptacle (the Carrier connector is receptacle, the cable connector is the plug). Some points about this illustration:

- A USB 2.0 Type-A connector only has pins 1 through 4 present. Pinout details are in Table 11 below.
- A USB 3 Type-A connector has 9 pins:
  - Pins 1 through 4 from the USB 2.0 definition are used for power, GND and a USB 2.0 data pair.
  - Pins 5 through 9 are used for SuperSpeed or SuperSpeed+ TX and RX pairs and a GND.
- A USB 2.0 cable plug may be used with a USB 3 receptacle, but only the USB 2.0 link will function.
- A USB 3 cable plug may be used with a USB 2.0 receptacle, but only the USB 2.0 link will function.
  - The USB 3 pins 5 through 9 are cleverly positioned so that they are invisible to the USB 2.0 plug.
- USB 2.0 target devices are allowed to consume up to 500 mA at 5V on a Type-A connector.
- USB 3 target devices are allowed to consume up to 900 mA at 5V on a Type-A connector.

Figure 25: **USB 3 Type-A Connector Receptacle – Looking Into the Receptacle**

Table 11: USB Type-A Pin-Out

Pin	Signal	Signal Description	Notes
1	VBUS	5V current limited USB target power	500mA (USB 2) or 900mA (USB 3)
2	D-	USB 2.0 differential signal (-)	
3	D+	USB 2.0 differential signal (+)	
4	GND	GND for USB 2.0 pair and power	
5	SSRX-	USB SuperSpeed RX(-)	Not used / not present on USB 2.0
6	SSRX+	USB SuperSpeed RX(+)	Not used / not present on USB 2.0
7	GND	GND for SuperSpeed RX and TX cable drain wire	Not used / not present on USB 2.0
8	SSTX-	USB SuperSpeed TX(-)	Not used / not present on USB 2.0
9	SSTX+	USB SuperSpeed TX(+)	Not used / not present on USB 2.0

Type-A Connector Electrical Distinctions

There are three general categories of USB Type-A connectors:

- USB 2.0   480 Mbps USB 2.0 signaling (no USB 3) – 4 pin connector
- USB 3.2 Gen 1   5 Gbps SuperSpeed signaling (along with USB 2.0) - 9 pin connector
- USB 3.2 Gen 2   10 Gbps SuperSpeed+ signaling (along with USB 2.0) – 9 pin connector

For USB 2.0 and USB 3.2 Gen 1 Type-A connectors, there are many vendors and styles (R/A, vertical, single, dual, quad combinations, combinations with other standards such as GbE etc).

For USB 3.2 Gen 2 (10 Gbps pair signaling), there are not many Type-A connector parts available as of this writing. Amphenol FCI is a connector vendor that has several 10 Gbps capable Type-A connectors available. Amphenol FCI GSB4111312HR, for example is a single R/A version of such a part.

Most USB 3.2 Gen 2 implementations use a Type-C connector rather than Type-A. Type-C implementations are covered in Sections 3.7.5. through 3.7.11. below.

3.7.2. USB 2.0 Type-A Example
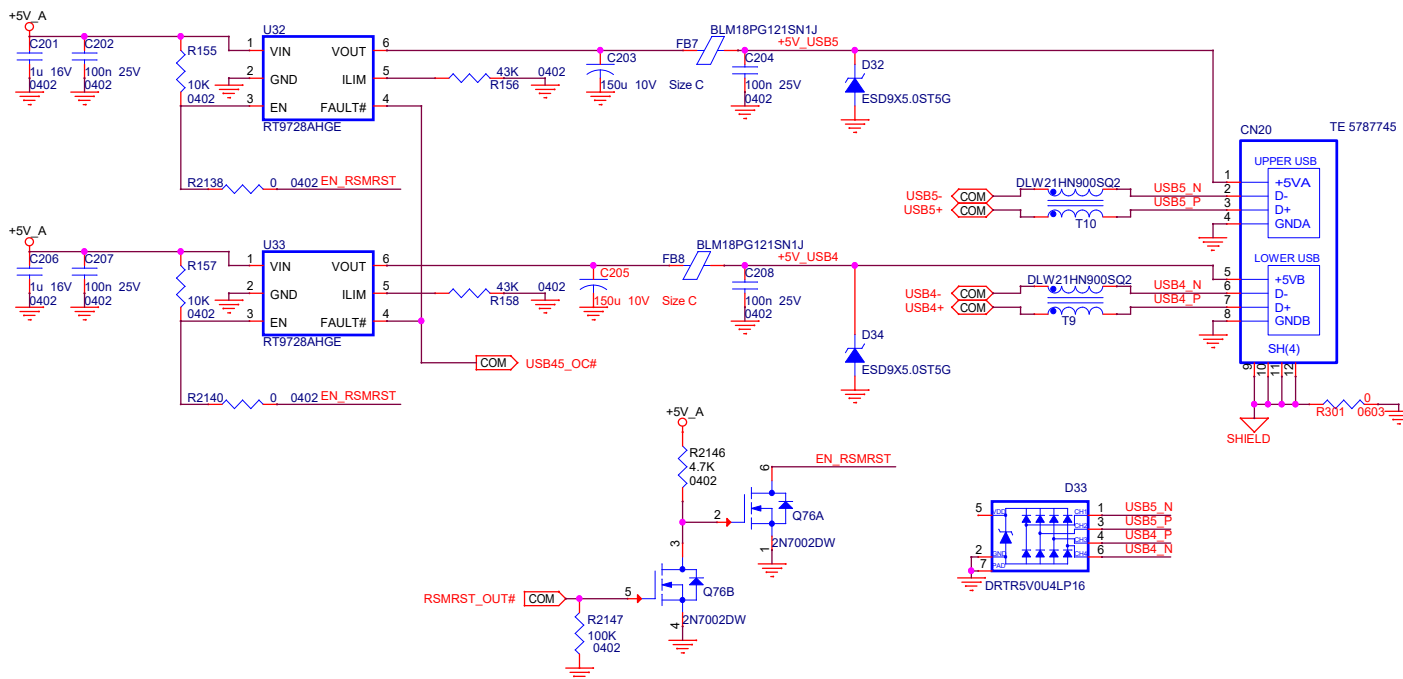
Figure 26: USB 2.0 Type-A Example

Figure 26 above illustrates a typical USB 2.0 implementation on a Carrier.

- The USB 2.0 data lines must be routed as differential pairs, in a no-stub fashion.

- Components T9 and T10 are common-mode chokes that are an EMI mitigation measure.

- Component D33 is a ESD protection diode array. Pins 1,3,4,6 may be exchanged if needed to provide the easiest no-stub routing.

- Components U32 and U33 are USB power switches and current limiters. For USB 2.0, the current delivered to a USB target device is to be limited to about 500 mA.

  - The power switch / current limiter shown is from Richtek. There are many similar parts available from Texas Instruments, Micrel, Microchip and others.

- The current limiter IC FAULT# pins are tied to the COM Module USB port 4 and 5 over-current input. There are 4 such inputs (for USB 0,1 and USB 2,3 and USB 4,5 and USB 6,7).

- The 5V power traces involved between the USB power switches, through FB7 and FB8 and on to connector CN20 must be sized to carry the 1A current (and this should be increased to 1.5 or 2 A to allow for a safety factor).

- The power switches are enabled by the COM-HPC RSMRST_OUT# signal. This signal going high indicates that the +5V_A power rail is stable.

This example applies to the Mini, except that the RSMRST_OUT# signal is a 1.8V signal and the gate threshold of the 2N7002 FET (Q76B) can be as high as 2.5V. A different FET with a lower threshold needs to be used, or the RSMRST_OUT# needs to be level shifted using an alternate method. The Rohm RA1C030LD is a suitable part with a Vgs gate threshold between 0.5V and 1.5V.

The USB over-current signal (net USB45_OC#) is not a problem for the Mini as the current limiting IC outputs driving this net are open drain pins on U32 and U33 and the signal is pulled up on the Mini Module.

USB 2.0 Allocation Notes – Client and Server Modules

The COM-HPC Client and Server pinouts allow up to eight USB 2.0 ports each. Note however that the first four USB 2.0 ports (COM-HPC USB0+/- through USB3+/-) are paired with the corresponding USB SuperSpeed ports (COM-HPC USB0_SSTX0+/- and USB0_SSRX0+/- through USB3_SSTX0+/- and USB3RX0+/-). A USB SuperSpeed port needs a specific companion USB 2.0 pair for certain setup functions.

Thus … if the Carrier needs one or more USB 2.0 only ports (no SuperSpeed) in addition to the four SuperSpeed capable ports, the above pairings need to be considered. COM-HPC USB4+/- through USB7+/- are USB 2.0 only ports.

USB 2.0 Allocation Notes – COM-HPC Mini Modules

USB 2.0 resource allocations are different on the COM-HPC Mini than on the Client and Server. The allocations on the Mini are well described in the COM-HPC Base Specification Section 10. USB 2.0 channels 6 and 7 are reserved for USB 2.0 use only on the Mini.

3.7.3.   USB 3.2 Gen 1 and Gen 2 Type-A

A USB 3.2 Gen 1 example on a Type-A connector is given in the ***COM Express Carrier Design Guide*** Rev 2.0 Section 2.9. At the time that the COMe Design Guide was written, USB 3.2 Gen 1 (single SuperSpeed TX pair and single RX pair, 5 Gbps signaling, plus a USB 2.0 pair) was referred to as USB 3.0.

A USB 3.2 Gen 2 Type-A connector implementation (single SuperSpeed+ TX pair and single RX pair, 10 Gbps signaling) is basically the same as a Gen 1 implementation, except that the components involved may need an upgrade for the 10 Gbps signaling: lower capacitance ESD diodes, different common mode choke choices and a connector receptacle appropriate for 10 Gbps signaling. An additional consideration is that there may be a greater need for a redriver. Most USB 3.2 Gen 2 implementations use a Type-C connector rather than Type-A. Type-C implementations are covered in Sections 3.7.5. through 3.7.11. below.

No USB 3 redriver is shown in the COM Express Carrier Design Guide example. If the traces from the COM-HPC Module connector to the Type-A host receptacle are more than a few inches, then a Carrier redriver may be advisable.

3.7.4. USB 3 Redrivers

Table 12: USB 3 Redrivers

Vendor	P/N	Notes
Diodes Inc	PI3EQX7841	USB 3.1 Gen 1 capable single port redriver (1 TX pair and 1 RX pair) 5 Gbps per pair
	PI3EQX1004E	USB 3.1 Gen 2 capable dual port redriver (2 TX pairs and 2 RX pairs) 10 Gbps per pair
Texas Instruments	TUSB522P	USB 3.2 Gen 1 capable single port redriver (1 TX pair and 1 RX pair) 5 Gbps per pair
	TUSB1002A	USB 3.2 Gen 2 capable single port redriver (1 TX pair and 1 RX pair) 10 Gbps per pair
	TUSB1004	USB 3.2 Gen 2 capable dual port redriver (2 TX pairs and 2 RX pairs) 10 Gbps per pair May be used to support two USB 3.2 Gen 2 ports

The items shown in Table 12 above represent only a small sample of such parts available on the market. USB Type-C Port Multiplexers, which may include redriver and retimer capabilities, are listed in Table 14 below. USB Type-C implementations are covered in Sections 3.7.5. through 3.7.10. below, and USB4 in Section 3.7.11. .

USB Hubs – May Serve as Retimers

USB 2 and USB 3 hubs are plentiful and may be considered as a form of a USB retimer: they clock the USB 2 and 3 signals in, process them and clock them out in fresh form. Of course the downstream bandwidth is shared, if more than one downstream hub port is used. Microchip Technologies (www.microchip.com) seems to be the dominant USB hub supplier and has dozens of offerings. Granted, there may be some software subtleties concerning the use of USB hubs versus true USB retimers (a true retimer should be invisible to software apart from possible setup; a hub has to be enumerated by the operating system, etc.).

3.7.5. USB Type-C Overview

USB Type-C refers to a small form factor reversible connector definition (reversible cable plug, no polarity, can be inserted with either orientation), and to the USB and other data and negotiated power delivery formats that it supports. Some highlights include:

- Polarity free operation
    - Cable plug can be used in either orientation
- USB 2.0   (480 Mbps signaling)
- USB 3.2 Gen 1 x1 and Gen 2 x 1   (5 Gbps, 10 Gbps signaling) (single lane)
- USB 3.2 Gen 1 x2 and Gen 2 x 2   (10 Gbps, 20 Gbps signaling) (2 lanes)
- "Alternate Modes" including
    - DisplayPort (2 lanes) + USB 3.2
    - DisplayPort (4 lanes)
    - HDMI
    - Intel Thunderbolt
    - Other vendor specific Alternate Modes
- USB4, described Section 3.7.11. below.
- USB Power Delivery (PD) protocol and implementation
    - Allows negotiated power delivery, from 5V up to 20V and up to 100W.
- Active cable support (electronics within the USB cable assembly)

For an excellent explanation of USB Type-C features, capabilities and details on how they work, see the Microchip Technologies Application Note **AN1953 Introduction to USB Type-C.** Much of the information in this section has been adapted from this note.

A typical Type-C receptacle is shown in Figure 27 below, at the left. A typical cable plug is shown at the right. The connector is fairly small, with an overall width less than 9mm and body height just under 3mm. These dimensions are similar but slightly larger than the Apple Computer "Lightening" connectors that are popular on consumer cell phones. The USB Type-C connector system has more capabilities than the "Lightening" system.

The receptacle and corresponding cable plugs are mechanically symmetrical and the cable plug can be used in either orientation. The connector pinout, presented on the following page, is almost completely symmetrical.

There are some locking versions of the USB Type-C connector available.

Figure 27:   USB Type-C Receptacle and Plug Images

Figure 28:   USB Type-C Receptacle Pin-Out – Looking Into Carrier Receptacle

A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12
GND	TX1+	TX1-	VBUS	CC1	D+	D-	SBU1	VBUS	RX2-	RX2+	GND
GND	RX1+	RX1-	VBUS	SBU2	D-	D+	CC2	VBUS	TX2-	TX2+	GND
B12	B11	B10	B9	B8	B7	B6	B5	B4	B3	B2	B1

Note that the Type-C receptacle connector pinout is mostly symmetrical with respect to flipping the plug connector. If the connector plug is inserted "right side up" (plug A1 to receptacle A1 etc.), all plug and receptacle signals match. If the connector plug is inserted "upside down" (plug A1 to receptacle B1 etc) then a few things must be sorted out by Carrier hardware, as explained on the following pages.

Table 13: USB Type-C Connector Pinout

Pin	Signal Name	Signal Description	Pin	Signal Name	Signal Description
A1	GND		B12	GND	
A2	TX1+	SuperSpeed TX1+	B11	RX1+	SuperSpeed RX1+
A3	TX1-	SuperSpeed TX1-	B10	RX1-	SuperSpeed RX1-
A4	VBUS	Bus Power to Peripheral USB device	B9	VBUS	Bus Power to Peripheral USB device
A5	CC1	Configuration Channel 1 or VCONN	B8	SBU2	Side Band Use 2
A6	D+	USB 2.0 D+	B7	D-	USB 2.0 D-
A7	D-	USB 2.0 D-	B6	D+	USB 2.0 D+
A8	SBU1	Side Band Use 1	B5	CC2	Configuration Channel 2 or VCONN
A9	VBUS	Bus Power to Peripheral USB device	B4	VBUS	Bus Power to Peripheral USB device
A10	RX2-	SuperSpeed RX2-	B3	RX1-	SuperSpeed TX2-
A11	RX2+	SuperSpeed RX2+	B2	RX1+	SuperSpeed TX2+
A12	GND		B1	GND	

For the yellow shaded cells in the Table above, the 'A' and 'B' signals have complete symmetry and nothing at all needs to be done to sort them out if the cable plug is "upside down". The non-shaded signals need some help in the "upside down" case – to get the plug signals to the correct Carrier destinations. Multiplexers are involved, and, to reduce stubs and preserve signal integrity, the "right side up" signals are usually routed through multiplexers along with the "upside down". signals.

USB 2.0 D+ and D-

The USB 2.0 D+ and D- differential pair data lines are arranged in a symmetrical block in the USB Type-C pinout definition. This arrangement has the result that no signal multiplexing is needed for "right side up" and "upside down" cable plug insertions. However, as a consequence of this arrangement, there are some short signal stubs. Since the USB 2.0 signaling rate is relatively low (480 Mbps), this is not a problem.

VBUS

VBUS is the power source provided by the host system (the COM-HPC Carrier in this case) to the attached downstream port. It can be the traditional fixed 5V current limited supply per USB 2.0 or USB 3.x, or it can be a higher voltage supply, up to 20V, and up to 100W, as negotiated by implementations following the *USB Type-C Power Delivery Specification.* The Power Delivery (PD) negotiation and implementation capability is optional, but necessary for higher powered peripherals. The PD negotiation happens over one of the two CC lines. Note that there are four VBUS pins and four GND pins. All eight pins should be used, to handle the possibly high power and current levels.

VCONN

VCONN is 5V nominal 1W max power source for active USB cables. Active cables have internal electronics that boost the signals carried, allowing longer cable assembles. The electronics in an Active cable may take their power from VCONN or VBUS. VCONN is routed to the receptacle CC2 pin if the plug is "right side up" or to the receptacle CC1 pin if the plug is "upside down".

CC1 and CC2 Configuration Channel Signals

The CC1 and CC2 signals serve several purposes in USB Type-C implementations:

- The CC1 and CC2 pins are used by the host system to identify whether the cable plug is inserted "right side up" or "upside down", through an analog detection process, relying on certain resistor values on the host side and on the downstream port side.

- The CC1 and CC2 signals are also used to identify the basic host power delivery requirements to the downstream peripheral. A resistor scheme and analog measurements are used to identify 5V 500 mA. 1.5A and 3A possibilities.

- The receptacle CC1 pin (if the plug connector is "rightside up") or the receptacle CC2 pin (plug connector is "upside down") may be used to negotiate the USB Type-C Power Delivery using a one – wire protocol defined in the *USB Power Delivery Specification*. This is optional but necessary if the peripheral needs a VBUS voltage over 5V.

- VCONN power is distributed to the "unused" CC pin (CC2 for plug "rightside up" and CC1 for plug "upside down").

- The Microchip application note AN1953 explains the CC1 and CC2 operational details very well.

SuperSpeed TX1+, TX1-, RX1+, RX1-

- If the cable plug is "rightside up" then these pins are used for the USB 3.2 Gen 1 and Gen 2 single lane SuperSpeed signals, or for the first lane of a two lane implementation.

- If the cable plug is "upside down" then the cable TX1+, TX1-, RX1+ and RX1- signals land on the receptacle TX2+, TX2-, RX2+ and RX2- pins. In this case, Carrier hardware has to route these signals to the proper TX1+, TX1-, RX1+ and RX1- destinations on the Carrier board.

- In practice, a Carrier Board multiplexer is used to route the receptacle TX1 and RX1 pairs to the proper Carrier destination, as the signals are high speed and stubs must be avoided.

- In some cases, the TX1 and RX1 high speed pairs are used for "Alternate Mode" purposes. Alternate Mode use is negotiated as part of the USB Power Delivery protocol.

SuperSpeed TX2+, TX2-, RX2+, RX2-

- If the cable plug is "rightside up" then these pins may be used for the second lane set of a USB 3.2 Gen 1x2 or Gen 2x2 implementation.

- If the cable plug is "upside down" then the cable TX2+. TX2-, RX2+ and RX2- signals land on the receptacle TX1+, TX1-, RX1+ and RX1- pins. In this case, Carrier hardware has to route these signals to the proper TX2+. TX2-, RX2+ and RX2- destinations on the Carrier board.

- In practice, a Carrier Board multiplexer is used to route the receptacle TX2 and RX2 pairs to the proper Carrier destination.

- In some cases, the TX2 and RX2 high speed pairs are used for "Alternate Mode" purposes. Alternate Mode use is negotiated as part of the USB Power Delivery protocol.

- A common Alternate Mode usage of these pairs is for a DisplayPort implementation.

SBU1 and SBU2

- SBU is an acronym for Side Band Use.

- These are optional signals, not needed for USB only implementations.

- For the DisplayPort Alternate Mode, these signals are used for the DisplayPort Aux Channel pair.

- For an HDMI Port Alternate Mode, these signals are used for the HDMI I2C channel.

3.7.6. USB Type-C Port Multiplexers

Selecting a USB Type-C Port Multiplexer can be tricky. It is important to understand what it does, and does not do, and what software support is available. Tables 14 lays out some of the possibilities. It is best to work with the silicon vendor and Module vendor FAEs on the details. It's pretty difficult, but not impossible, to cover all the possible USB Type-C modes in a single design. All the USB Type-C Port Multiplexers listed in this Table incorporate redriver or retimer functions, reducing part count.

Table 14: USB Type-C Port Multiplexers – Possible Modes

Ref	Mode	Notes	Possible Part(s)	Part Notes
1	USB 3.2 Gen 1x1	5 Gbps signaling single SuperSpeed TX pair and single RX pair used	TUSB542 TUSB1042 TUSB1104 TUSB1142 TUSB1146 Intel JHL9040R	TUSB542 is 5 Gbps Others are 10 Gbps capable
2	USB 3.2 Gen 2x1	10 Gbps signaling single SuperSpeed+ TX pair and single RX pair used	TUSB1042 TUSB1044 TUSB1046 TUSB1104 TUSB1142 TUSB1146 Intel JHL9040R	10 Gbps capable parts
3	USB 3.2 Gen 1x2	5 Gbps signaling per pair dual SuperSpeed TX pairs and dual RX pairs used 10 Gbps net TX speed, 10 Gbps net RX speed	TUSB1104	
4	USB 3.2 Gen 2x2	10 Gbps signaling per pair dual SuperSpeed+ TX pairs and dual RX pairs used 20 Gbps net TX speed, 20 Gbps net RX speed	TUSB1104	
5	DisplayPort Alternate Mode 2 DP lanes + USB 3 Separate DP Source	USB 3.2 Gen 1x1 or Gen 2x1 on TX1 / RX1 Two DP pairs on TX2 / RX2 (RX2 used as DP TX pair) DP sourced externally, from GPU pins	TUSB546A-DCI TUSB1044 TUSB1046 TUSB1046A-DCI TUSB1146	TUSB546A-DCI is 5 Gbps Others are 10 Gbps capable
6	DisplayPort Alternate Mode 4 DP lanes Separate DP Source	No USB 3 at all (USB 2 remains) Four DisplayPort pairs on TX1,RX1,TX2,RX2 DP sourced externally, from GPU pins	TUSB546A-DCI TUSB1046A-DCI TUSB1046 TUSB1146	TUSB546A-DCI is 5 Gbps Others are 10 Gbps capable
7	DisplayPort Alternate Mode 2 DP lanes + USB 3	USB 3.2 Gen 1x1 or Gen 2x1 on TX1 / RX1 Two DP pairs on TX2 / RX2 (RX2 used as DP TX pair) DP multiplexed with USB 3 within chip-set	TUSB544 TUSB1044 Intel JHL9040R	
8	DisplayPort Alternate Mode 4 DP lanes	No USB 3 at all (USB 2 remains) Four DisplayPort pairs on TX1,RX1,TX2,RX2 DP multiplexed with USB 3 within chip-set	TUSB544 TUSB1044 Intel JHL9040R	
9	HDMI Alternate Modes	Similar to DP Alternate Modes	TUSB546	
10	USB4: 20 Gbps only	All USB3 modes USB4: 20 Gbps signaling using 2 lanes DP Alternate Modes	Intel JHL9040R	
11	Thunderbolt Modes	All USB 3 modes USB4: 20 Gbps signaling, 40 Gbps using 2 lanes DP Alternate Modes PCIe Alternate Mode	Intel JHL8940R	

Notes on Table 14:

- All the "TUSB" prefixed parts listed above are from Texas Instruments
- The Intel JHL9040R is a USB4 retimer part, formerly known as the "Hayden Bridge".

3.7.7. USB Type-C Power Delivery Controllers

The USB Type-C specification is an ambitious specification with many features. For Power Delivery, the specification allows up to 100W of power, over a voltage range from 5V to 20V, to be delivered either **out** of the device in question or accepted **into** the device. For example, a laptop computer might want to provide power to an external display or printer in some situations. In a different situation, the same laptop may want to accept power from an external charger for battery recharging.

The Power Delivery options are negotiated over the USB Type-C CC lines. If there is no negotiation, than a simple old style USB 3.0 or USB 2.0 Type-A power delivery **out** of the COM-HPC host is assumed

Sections 3.7.10. and 3.7.11. below, and more specifically in Figures 34 and 35 below show a USB Type-C Power Delivery solution that allows 15W max power at 5V, **out** of the COM-HPC carrier to an external device. The Texas Instruments TPS65994 Power Delivery controller is shown. This is actually a dual part that could support two USB Type-C ports. Only one port is used in the Section 3.7.10. USB 3.2 Gen 2x2 example, and similarly for the Section 3.7.11. USB4 port example.

Higher power levels (up to 100W, voltages over 5V to 20V range), either **out** or **into** the COM-HPC Carrier are possible with other PD controllers. For example, the Texas Instruments TPS65987D device allows up to 100W power delivery, over a 5V to 20V range, **out** of or **into** the system, using integrated power FETs. The voltage level, the current level and the current direction are negotiated over the CC lines before power is applied to, or accepted from, the USB Type-C VBUS.

The TPS65994 device used in the USB 3,2 Gen 2x2 and USB4 design examples below has a provision, using external power FETs, for up to 100W to come **in** to the design, but this capability is not used in these examples.

There are many additional USB Type-C Power Delivery controllers available from Cypress Semiconductor (now part of Infineon), Microchip Technologies, NXP, On Semiconductor, Texas Instruments and others.

3.7.8. USB Type-C Port Protection Components

It is important to protect USB Type-C port pins against accidental exposure to 20V VBUS contact, and against ESD events. The USB 3.2 Gen 2x2 and USB4 schematic examples (Sections 3.7.10. and 3.7.11. below) illustrate this.

The example here uses a Texas Instruments TPD6S300 USB Type-C Port Protector to protect the Type-C CC lines (2 pins), SBU lines (2 pins) and the USB 2.0 lines (4 pins, in the Type-C implementation). This is shown in Figure 35 below.

Note: since these examples were created, Texas Instruments has upgraded their Type-C Port Protector to the TPD6S300A and that should be used for new designs.

The high speed data pairs (2 TX pairs and 2 RX pairs, for USB, DP, HDMI etc) are protected separately in these schematic examples, using discrete low capacitance ESD diodes. This is shown in Figure 35 below.

There are many other possible USB Type-C Port Protection components, from Texas Instruments, Microchip Technologies, On Semiconductor, NXP and others.

If the COM-HPC is implementing a battery powered option, then there are battery charging and dead battery concerns to consider. Refer to the Texas Instruments and Microchip Technologies data sheets and application notes for more technical information on this.

3.7.9.   USB 3.2 Gen 2x1 Type-C Basic Implementation

A basic USB Type-C implementation that supports USB 2.0, USB 3.2 Gen 1 x1 and USB 3.2 Gen 2x1 is straightforward. An example is presented in Figure 29 on the following page. Although this example is in block diagram format and does not include the many passive components needed for a complete design, it only requires two small ICs. This Type-C example is hardly any more complex than a traditional USB 3 Type A , design, especially if a redriver is included in the Type A design.

The example in Figure 29 includes a USB Type-C port multiplexer and USB 3 redriver in a single IC package that can be placed close to the Type-C connector receptacle to best launch the signal over the USB cable. Note that TX line coupling capacitors are needed on the redriver output pins.

The example also includes a USB Type-C Power Source controller that performs cable detection, provides cable orientation information, provides VBUS power and VCONN power and current limiting for both, along with fault detection. This part does not implement the full USB Power Delivery protocol – this is not necessary here as the VBUS power is limited to traditional USB 3 values of 5V nominal, 1A operational and 1.5A fault current.

An implementation that allows the full USB Type-C Power Delivery gamut (5V to 20V, up to 100W) requires a more complex Power Source or Delivery part, that implements the one-wire negotiation on the CC1 or CC2 lines (depending on cable plug insertion polarity).

There are many useful parts for USB Type-C support available from Texas Instruments, Microchip Technology, Diodes Inc. and other vendors.

Figure 29 below uses the COM-HPC USB0 port as an example (for USB 2.0 and USB 3.2 signals). Any of the first four COM-HPC USB ports (USB0 through USB3) may be used. Remember that COM-HPC USB 2 and USB 3 ports are paired together. See the notes on this in the *COM-HPC Base Specification V1.0* Table 15.

Figure 29:   USB Type-C Basic Implementation: USB 3.2 Gen 1 and Gen 2
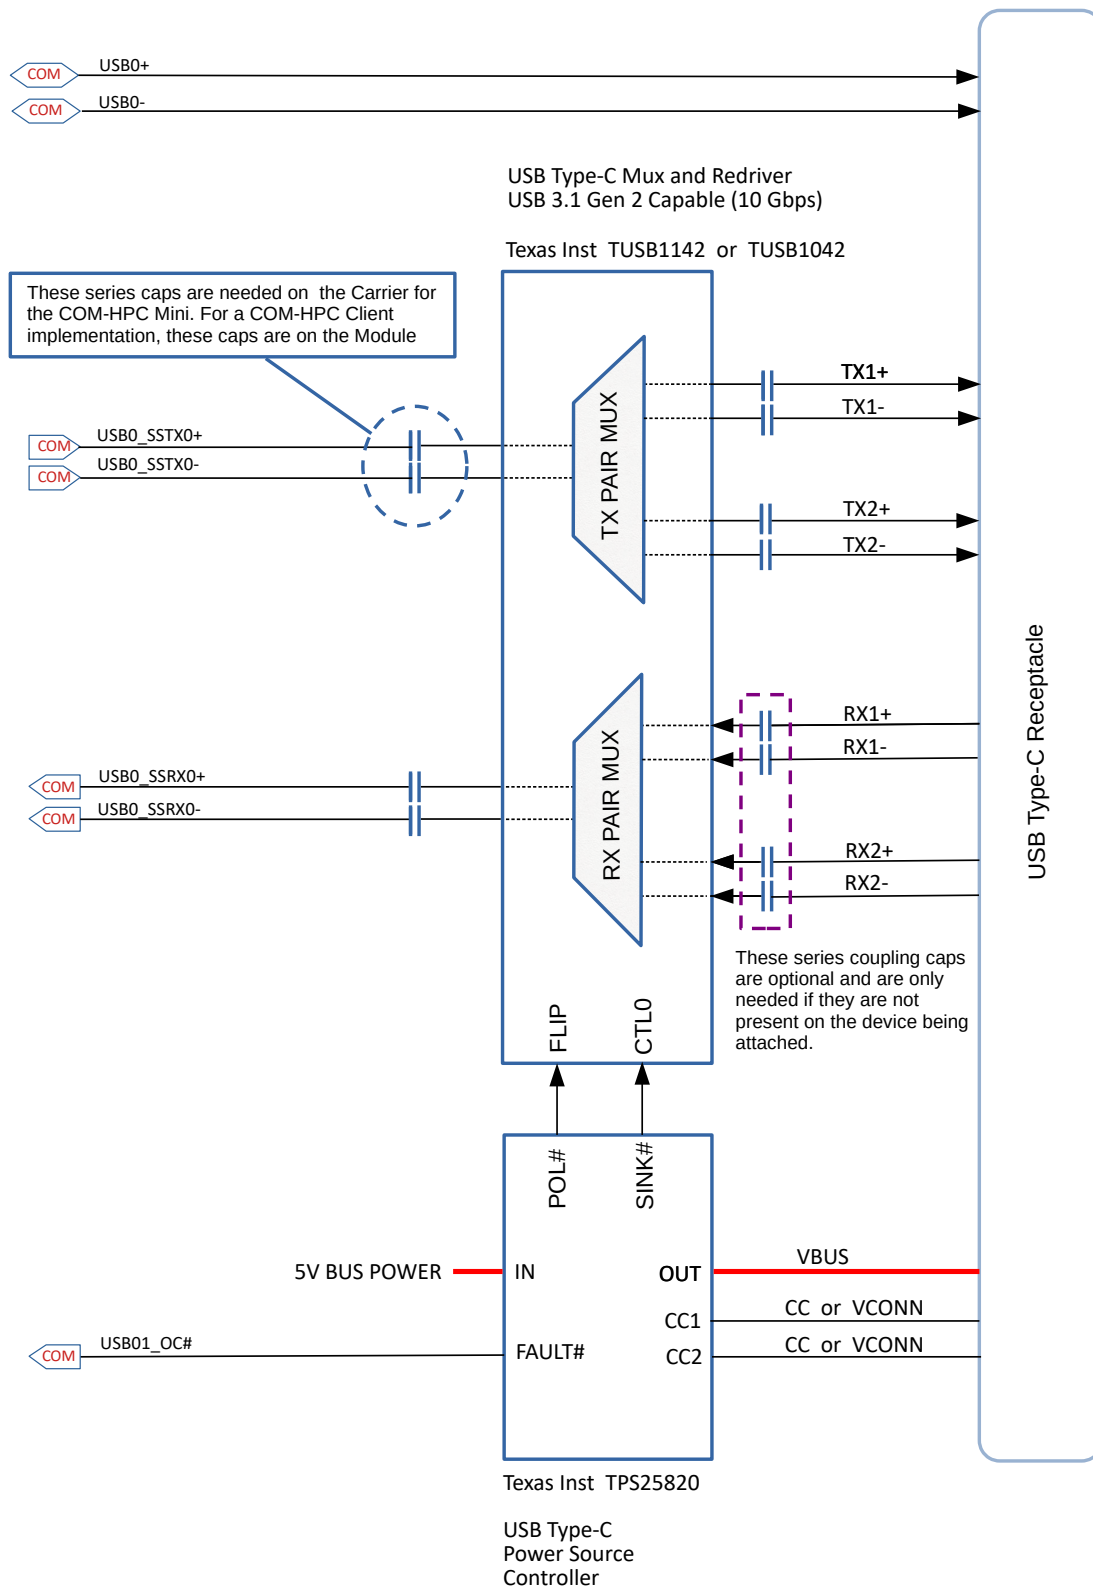
3.7.10.   USB 3.2 Gen 2x2 Type-C Example Implementation

A detailed schematic example of a USB 3.2 Gen 2x2 implementation (meaning two USB SuperSpeed TX pairs and two RX pairs, each pair capable of 10 Gbps signaling) is shown in Figures 30 through 35 below. The net TX signaling over two pairs is 20 Gbps, and the net RX signaling over two pairs is 20 Gbps This is delivered over a Type-C reversible connector.

A Type-C port multiplexer shown in Figure 31 below Texas Instruments (TI) TUSB1104 is an optimal part for this application.

A single TX and RX SuperSpeed pair Type-A connector option is implied by some resistor stuffing options in Figure 30 below (R5W6 through R5W9, not populated) but the Type-A connector details are not shown in this Figure set. The Type-A implementation is discussed in Section 3.7.3. above.

This example includes a USB Power Delivery controller, Texas Instruments TPS65994, in Figure 34 below. In this example, the power delivery is **out** of the COM-HPC Carrier, at 5V and at up to 3A. See Section 3.7.7. above more some discussion on Power Delivery controllers.

Figure 30: USB 3.2 Gen 2x2 Type-C (1 of 6): Option Resistors for Type-C or Type-A
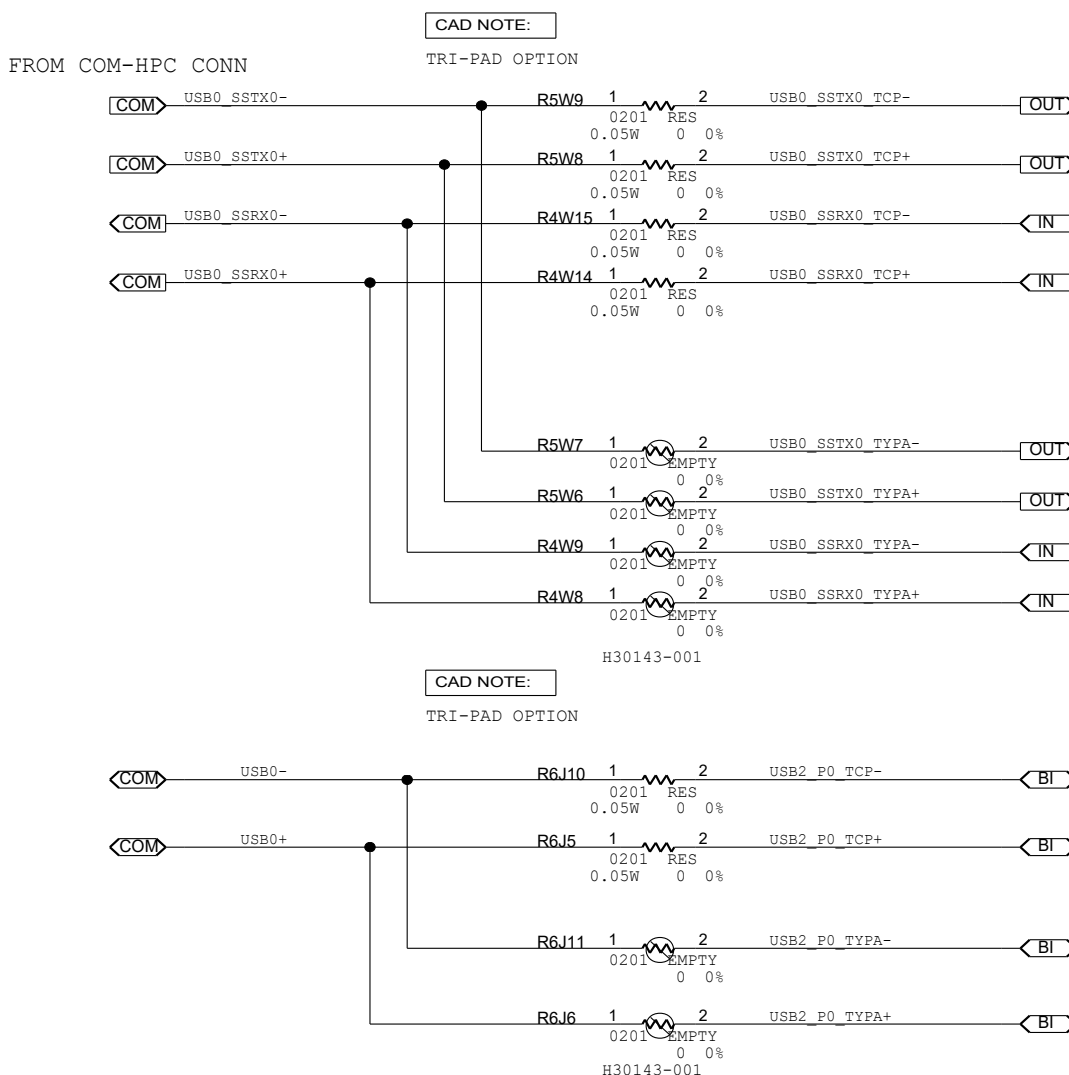
Figure 31: USB 3.2 Gen 2x2 Type-C (2 of 6): Port Multiplexer and Redriver
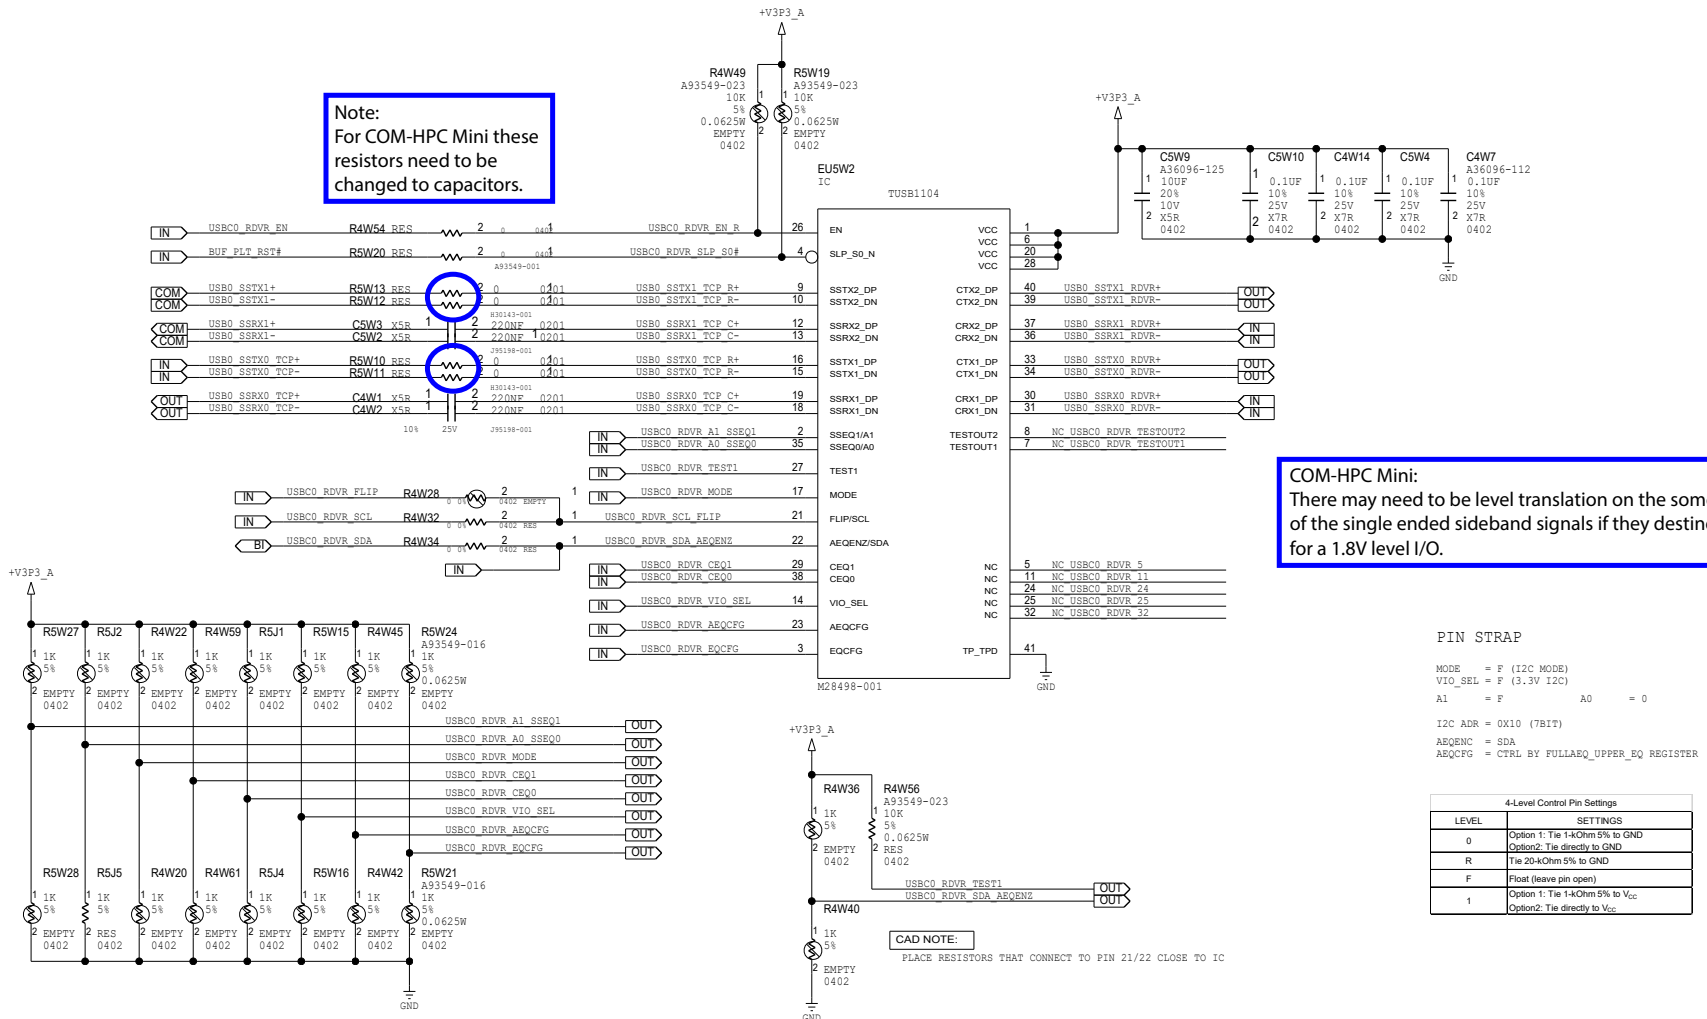
Figure 32:   USB 3.2 Gen 2x2 Type-C (3 of 6): EMI Mitigation and ESD Protection
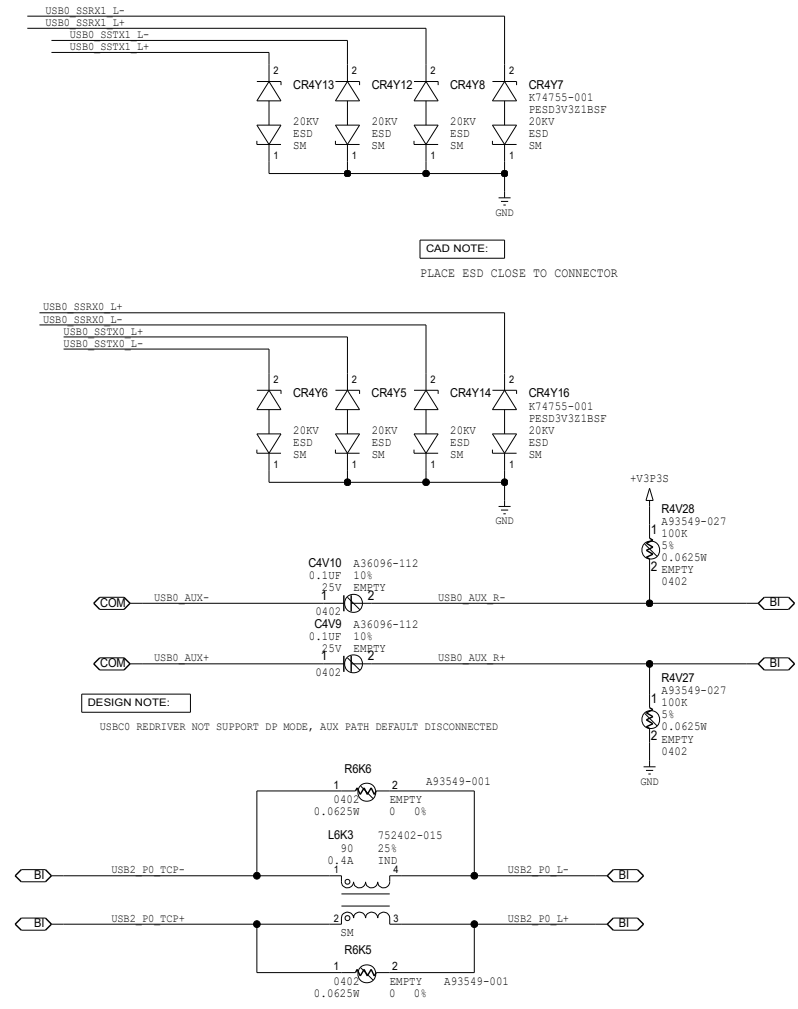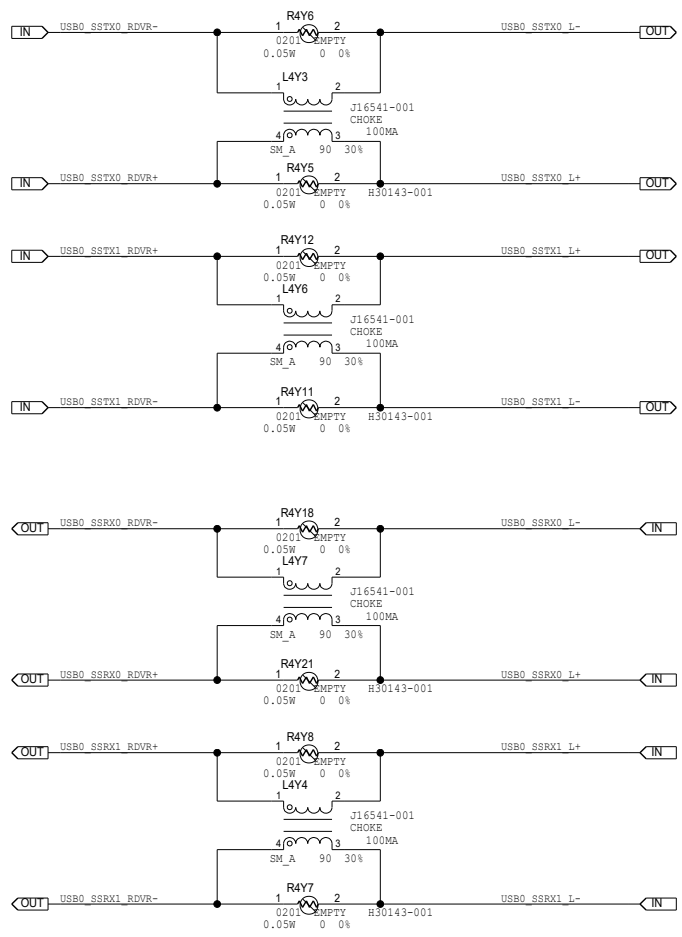
Figure 33:   USB 3.2 Gen 2x2 Type-C (4 of 6): Port Port Mux / Redriver Coupling Capacitors
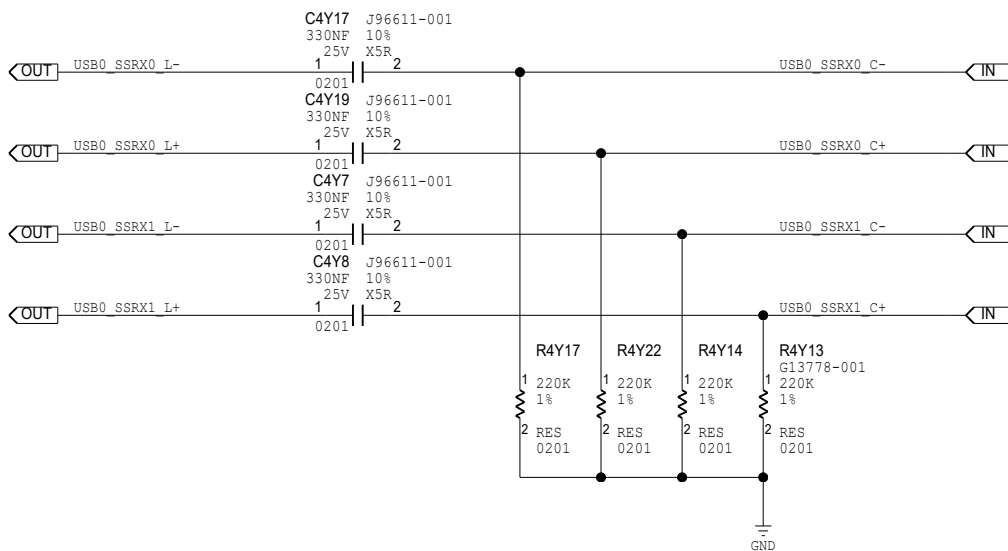
Figure 34: USB 3.2 Gen 2x2 Type-C (5 of 6): Type-C Power Delivery Controller

ADCIN1 decoded value [2]	ADCIN2 decoded value [2]	I²C address Index [1]	Dead Battery Configuration
7	0	#1	SafeMode: The device does not enable the sink path. USB PD is disabled until configuration is loaded. Note that the configuration could put the device into a source-only mode. This is recommended when the application loads the patch from EEPROM.
0	0	#2	
6	0	#3	
5	7	#4	
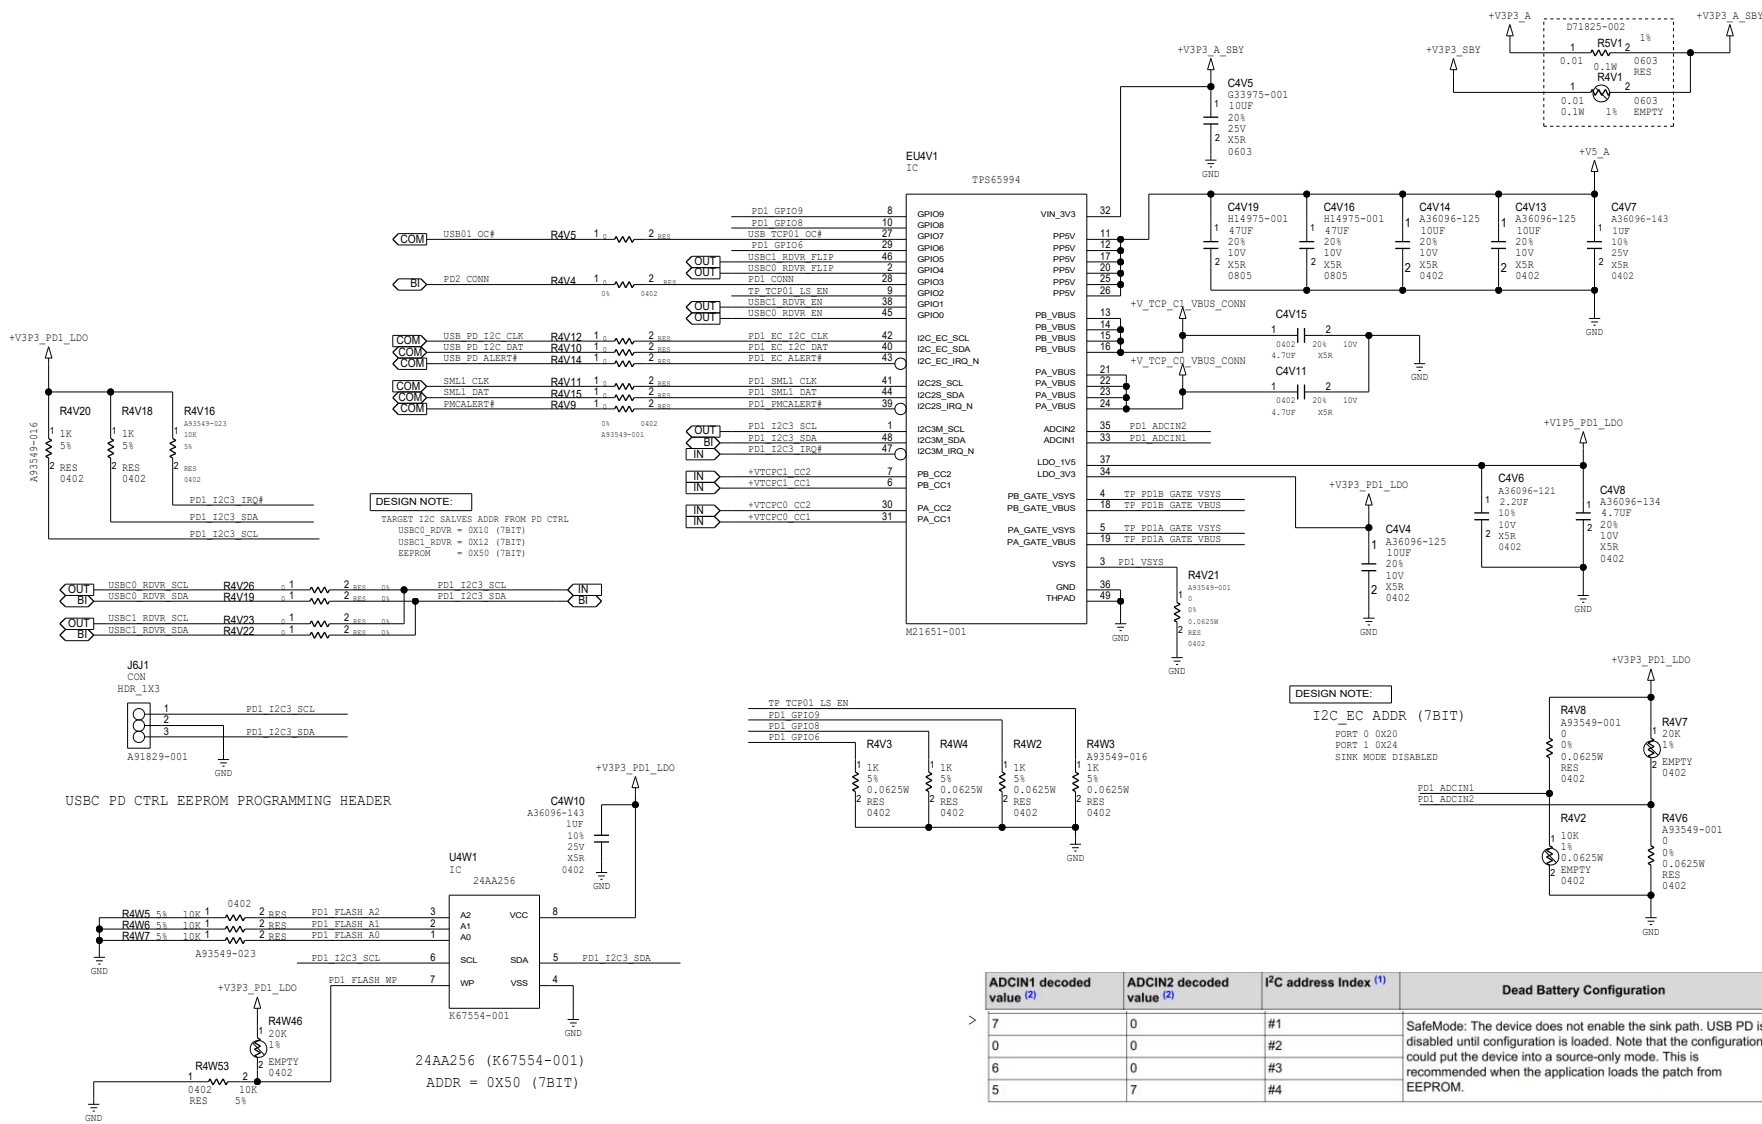
Figure 35: USB 3.2 Gen 2x2 Type-C (6 of 6): Type-C Connector and Port Protection
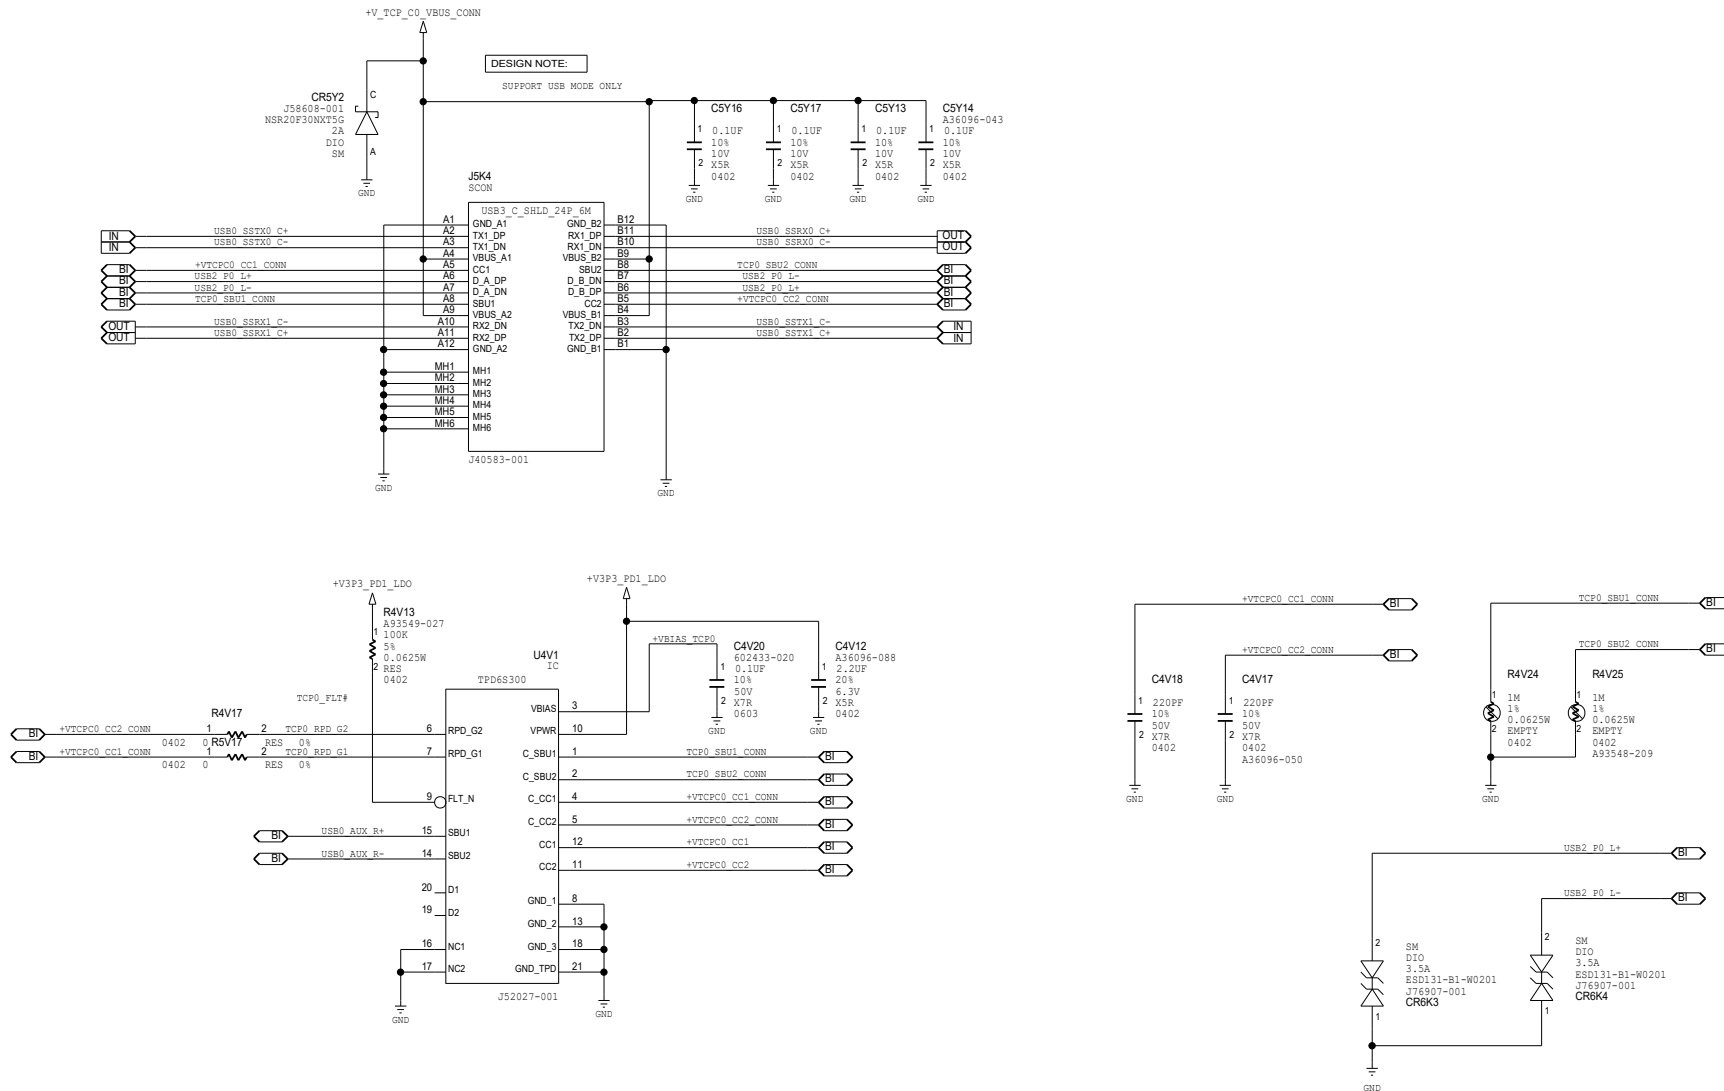
3.7.11. USB4

USB4 – COM-HPC Client and Server

An example COM-HPC USB4 implementation, that supports all USB modes up to USB4 Gen 3x2 and the USB Type-C Alternate Modes is shown in Figure 36 below. Of course this support requires that the COM-HPC Module used supports these modes as well. See Table 10 above for a summary of all the USB modes. The Intel JHL9040R "Hayden Bridge" USB4 retimer is used. Intel considers further details to be NDA protected material. Contact Intel for further details on the Intel part and implementation details.

This example includes a USB Power Delivery controller, Texas Instruments TPS65994. In this example, the power delivery is **out** of the COM-HPC Carrier, at 5V and at up to 3A. See Section 3.7.7. above more some discussion on Power Delivery controllers.

Additionally, a Texas Instruments TPD6S300A "Port Protector" is used to protect the USB Type-C port from adverse conditions.

USB4 – COM-HPC Mini

An example COM-HPC USB4 implementation for the COM-HPC Mini Module, that supports all USB modes up to USB4 Gen 3x2 and the USB Type-C Alternate Modes is shown in Figure 37 below. Of course this support requires that the COM-HPC Module used supports these modes as well. See Table 10 above for a summary of all the USB modes. The Intel JHL9040R "Hayden Bridge" USB4 retimer is used. Intel considers further details to be NDA protected material. Contact Intel for further details on the Intel part and implementation details.

This example includes a USB Power Delivery controller, Texas Instruments TPS65994. In this example, the power delivery is out of the COM-HPC Carrier, at 5V and at up to 3A. See Section 3.7.7. above more some discussion on Power Delivery controllers.

Additionally, a Texas Instruments TPD6S300A "Port Protector" is used to protect the USB Type-C port from adverse conditions.

Figure 36: USB4 on COM-HPC Client
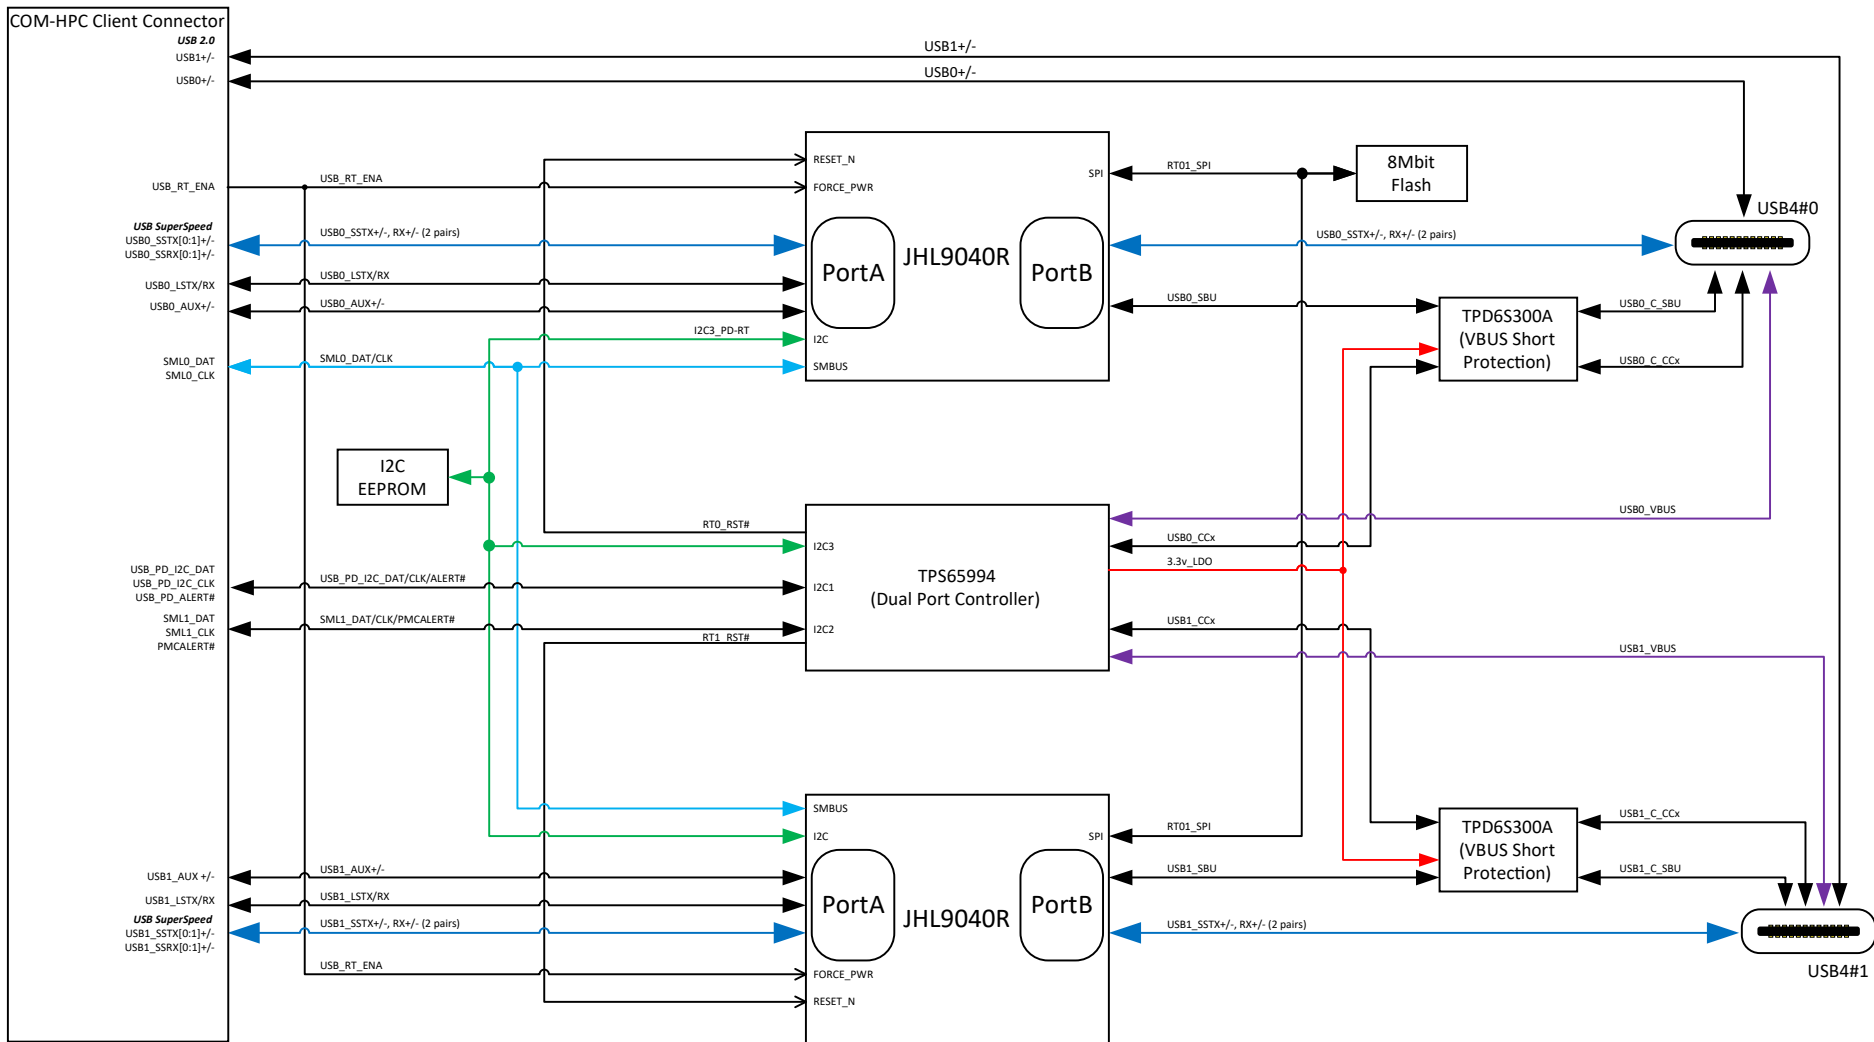
Figure 37:   USB4 on COM-HPC Mini
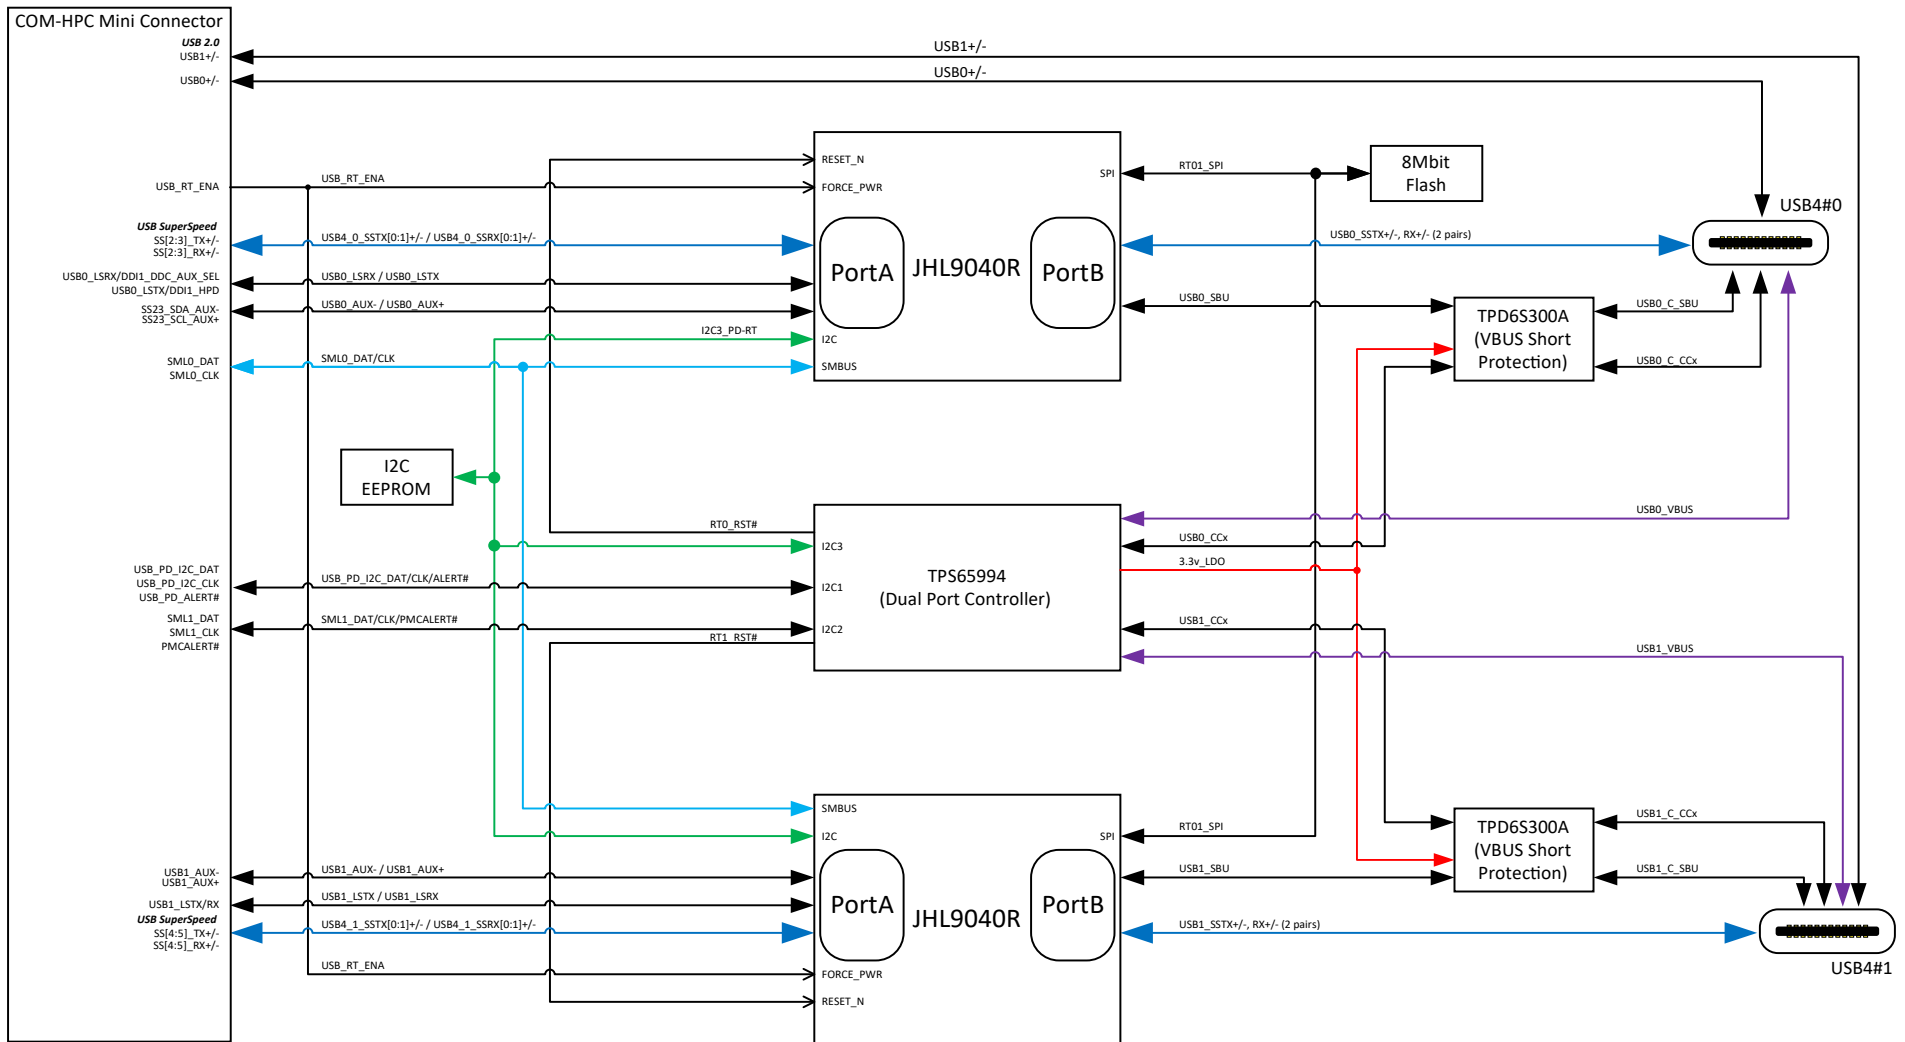
3.8. Boot SPI on Carrier

The COM-HPC Base Specification V1.0 describes various boot SPI options at some length in Sections 4.3.9 and 4.3.10. The layout topology for the BOOT_SPI bus is given in Section 6.11.1 of the **COM-HPC Base Specification V1.0.** Please refer to those Base document sections in addition to the materials presented here to get a bigger picture.

Contemporary x86 chipsets typically have a SPI boot bus with three chip-selects: two for up to two SPI Flash devices to hold various pieces of boot firmware including the BIOS, and possibly a backup BIOS, and a 3[rd] chip-select dedicated to an on-Module TPM. device. The two chip-set chip-selects for boot SPI flash devices can be routed to either two on-Module SPI Flash devices, or to one on-Module SPI Flash and to one Carrier based SPI Flash device. The various possible permutations are selected by a set of three Module strap pins named BSEL0, 1, and 2. See COM-HPC Base Specification V1.0 Section 4.3.10 Table 10 for the decoding of the BSEL[0:2] pins. It is possible to have the entire boot firmware image reside in a Carrier based SPI Flash device. It is also possible of course to have the entire boot image on the module, and it is possible to split the boot image to have some parts on the Module and some on the Carrier. Some Module designs implement multiplexers to allow even more options.

A typical Carrier Boot SPI Flash implementation is shown in Figure 38 below. Some points about this Figure are given on the following page.

Figure 38:   Boot SPI on Carrier (Example 1)
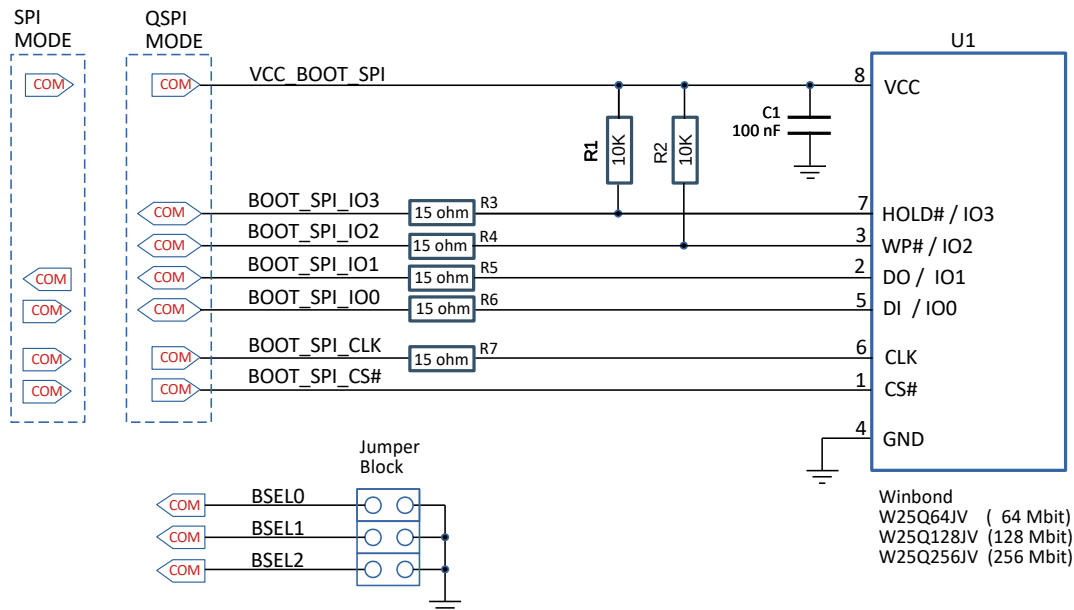
The information shown here applies to all COM-HPC versions. The COM-HPC Rev. 1.20 Specification document allows the VCC_BOOT_SPI power net to be either 3.3V or 1.8V. COM-HPC Mini Module vendors are likely to choose 1.8V as the SPI I/O, but this should be verified with the Module vendor. The Winbond part numbers in the Figure 38 above are 3.3V parts, not 1.8V capable.

It is also wise to ask the Module vendor what specific Carrier boot SPI devices are supported – for all COM-HPC variants (Client, Server, Mini). Some of the internal register details in these boot SPI devices vary and the Module BIOS has to be aware of the device it is dealing with.

- *   ***Notes on Figure 38 above:***

- *   The Carrier SPI Flash device power is provided by a special COM-HPC Module pin named VCC_BOOT_SPI. This should be the only power source for the Carrier SPI Flash device and any related pull-ups and bypass capacitors., as shown in the Figure 38.

- *   The VCC_BOOT_SPI voltage level may be 3.3V or 1.8V.

  - ○   This is Module vendor specific.

  - ○   It is not common or expected that a given SPI Flash device will be able to operate at both 3.3V and 1.8V. A few devices might be able to do so.

  - ○   The SPI Flash devices listed in the Figure above are 3.3V devices, and are not rated at 1.8V.

- *   The VCC_BOOT_SPI power net may be in the S5 (suspend) or S0 (on) power domains.

  - ○   This is Module vendor specific.

- *   QSPI devices from Microchip / SST are shown in the Figure above.

  - ○   Windbond is a very popular selection for QSPI devices: W25Q16JV is a sample Winbond base part number for their 16 Mbit part. There are 16, 32, 64 and 128 Mbit offerings from Winbond.

- *   There are many packaging options available from the QSPI vendors

  - ○   There are package size differences between vendors even for package names that at first glance sound the same (like SOIC8 etc) … so care must be taken.

  - ○   The Winbond SOIC8 packages are smaller than the Microchip devices.

- *   There are register differences between various SPI Flash vendor offerings. The Module firmware / BIOS may not be compatible with some devices. Check with the Module vendor.

  - ○   Carrier designers should use parts from the same SPI Flash vendor(s) and family as the Module vendor uses.

  - ○   There may be reasons to use different package types on the Carrier:

    - ▪   The Module vendor likely uses the smallest possible package size.

    - ▪   Carrier designs may want to implement a removable (socketed) SPI Flash device.

    - ▪   Carrier designers may elect to use a SPI package that is easier to rework.

- *   Contemporary SPI Flash devices may operate in one of several modes:

  - ○   Traditional SPI mode (noted at left side of Figure 38 above).

    - ▪   This mode has one data line into the SPI device and one out.

  - ○   QSPI ("Quad SPI") mode:

    - ▪   This mode has 4 bidirectional data lines, offering a higher net data bandwidth.

    - ▪   The SPI Flash devices typically power up in the traditional SPI Flash mode and must be put into the QSPI mode by software.

- *   The HOLD# and WP# inputs of a traditional SPI device are disabled by pull-ups R1 and R2 in the Figure above. For QSPI mode operation, the PCB trace stubs from the QSPI data lines to these pull-ups should be minimized.

  - ○   If the SPI device is to immediately be put into QSPI mode, it is likely possible that R1 and R2 can be omitted.

- *   There are specific routing rules for the BOOT_SPI_xx nets. See Section 4.4. of this document and Section 6.11.1 of the COM-HPC Base Specification V1.0.

Removable / Reprogrammable SPI Flash Devices

In some situations it is desirable or even required to have a socketed or removable Boot SPI Flash device. This is the case, for example, in some casino gaming jurisdictions, to allow the BIOS device to be removed and inspected by a regulatory technician. A socketed or removable BIOS can also be useful in product development situations, allowing easy replacement of a corrupted BIOS device. Some possible socket solutions are listed in Table 15 below:

Table 15: Boot SPI Socket Suggestions

Vendor	Vendor P/N	Notes
Enplas		Enplas offers a variety of sockets that accept several 8 and 16 pin SOIC sizes. Winbond and other vendors offer some of their SPI Flash devices in a 16 pin SOIC along with a variety of smaller form factors. It may be easier to find a socket for an SOIC16 device. The Carrier Boot SPI Flash device should be from the same flash vendor and family as the part used on the Module. The package details may be different.
Lotes	ACA-SPI-004-K ACA-SPI-006-T01	Should be suitable for the Microchip SST26VFxxxB SOIJ8 parts. Suitable for Macronix MX77U25650F (32 MB 1.8V QSPI) or MX77L25650F (32 MB 3.3V QSPI) 16 pin 300 mil SOIC parts shown in Figure 39 below. Other Lotes socket parts may be relevant here.
Generic		Winbond offers some of their Flash devices in 300 mil DIP format, for which there are many generic sockets.

Some gaming firms design their own removable BIOS assembles. These are sometimes referred to as "cartridges". This allows the use of any SPI Flash device desired, and it can ensure easy removal and replacement of the device..

Some Carrier designers add features that multiplex signals and power to the Carrier SPI Flash device allowing the device to be used as usual in the system or cut the device off from the system and allow the device to be reprogrammed by a cable to an external piece of programming equipment.

An example of such implementations (SPI device in a socket and a multiplexer to allow the SPI flash device to be programmed by an external programming tool) is shown in Figure 39 below. The programming tool in this case is from a company called Dediprog.

Figure 39: Boot SPI on Carrier – Socketed Flash and Multiplexer to External Programmer
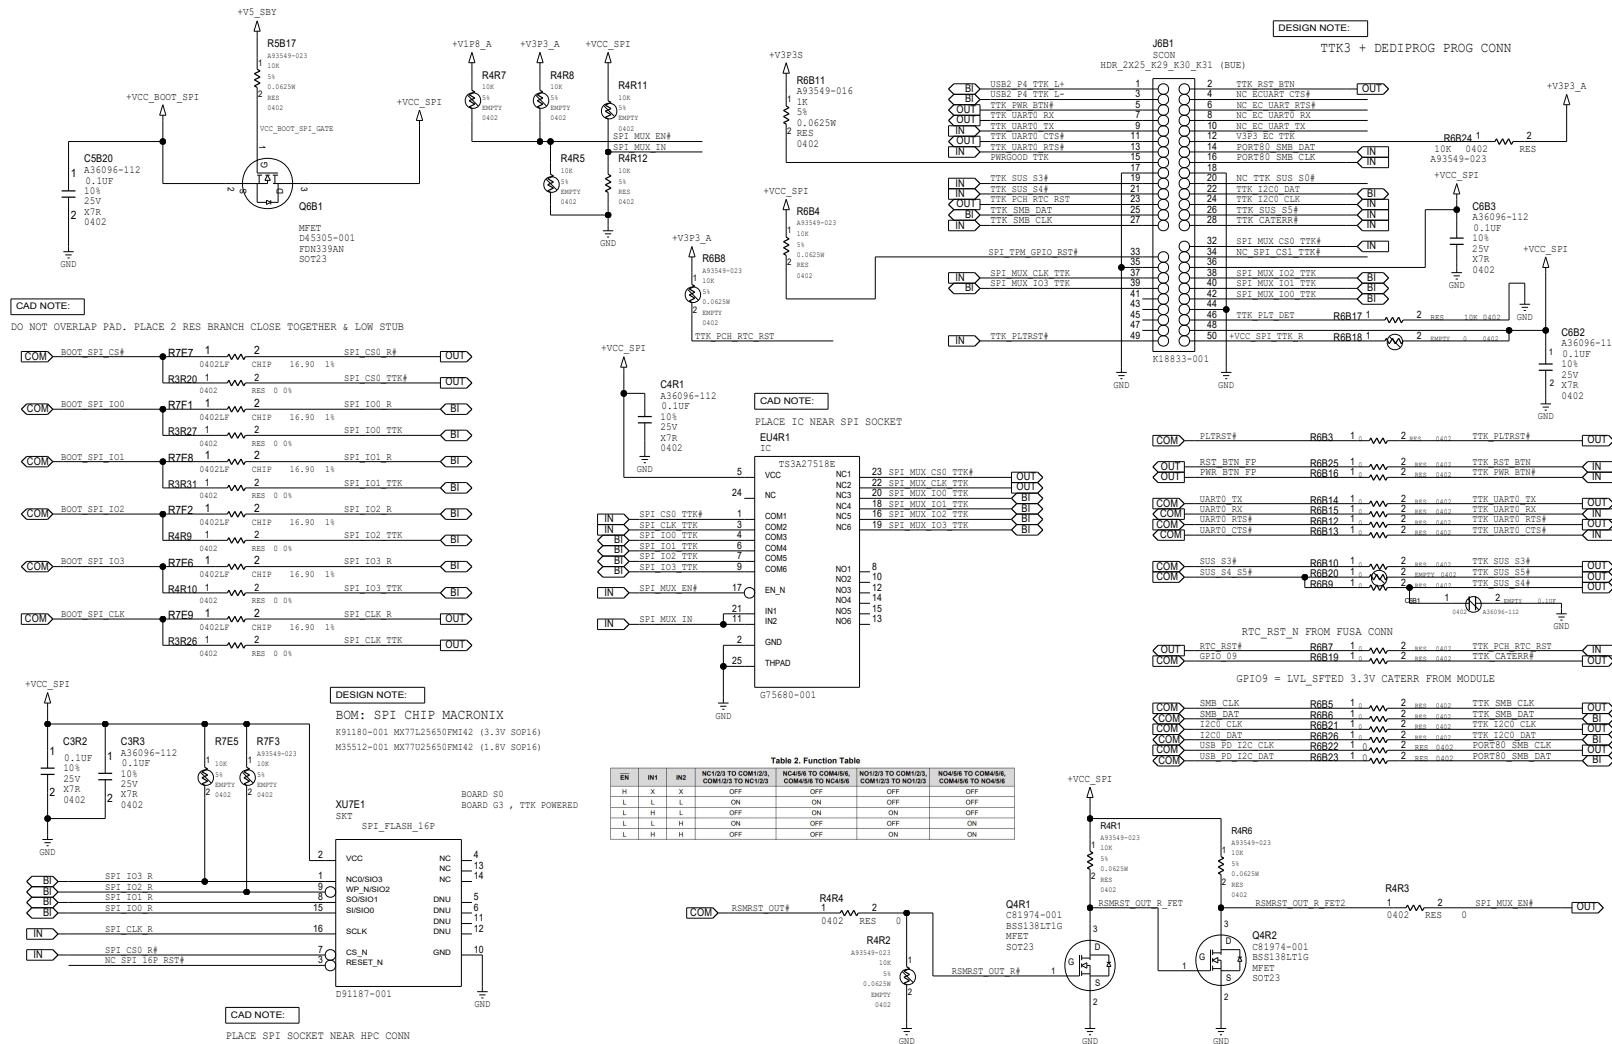
3.9.   eSPI

The COM-HPC Client and Server pinouts support an eSPI (Enhanced Serial Peripheral Interface) port. There may be up to two eSPI devices on the Module and up to two eSPI devices on the Carrier. The eSPI interface is promoted as the successor to the LPC (Low Pin Count) general x86 I/O interface.

The eSPI data and clock signals run at about 50 MHz. The COM-HPC Base Specification Section 6.11.2 recommends a "balanced tree" routing topology. This is also referenced in Section 4.4. of this document.

Figure 40 below illustrates a "generic" eSPI implementation example for one branch of the tree, that might apply to a Carrier Super I/O, FPGA, CPLD, eSPI to LPC bridge, or other eSPI peripheral.

The COM-HPC eSPI interface is a 1.8V level interface that operates in all power states, S5 through S0. The Figure 40 shows some additional signals that are 3.3V level signals, also active in all power domains, that may be needed for some eSPI peripheral implementations.

Some Carrier situations may require legacy Intel LPC (Low Pin Count) compatibility. The Microchip ECE1200 is a suitable eSPI to LPC bridge device that is referenced in some Intel literature for this task.

Microchip is also a popular vendor for Carrier based management micro-controllers with an eSPI interface.

Figure 40:   eSPI Generic Interface Example: SIO, FPGA, LPC Bridge, or Other Peripheral eSPI Device
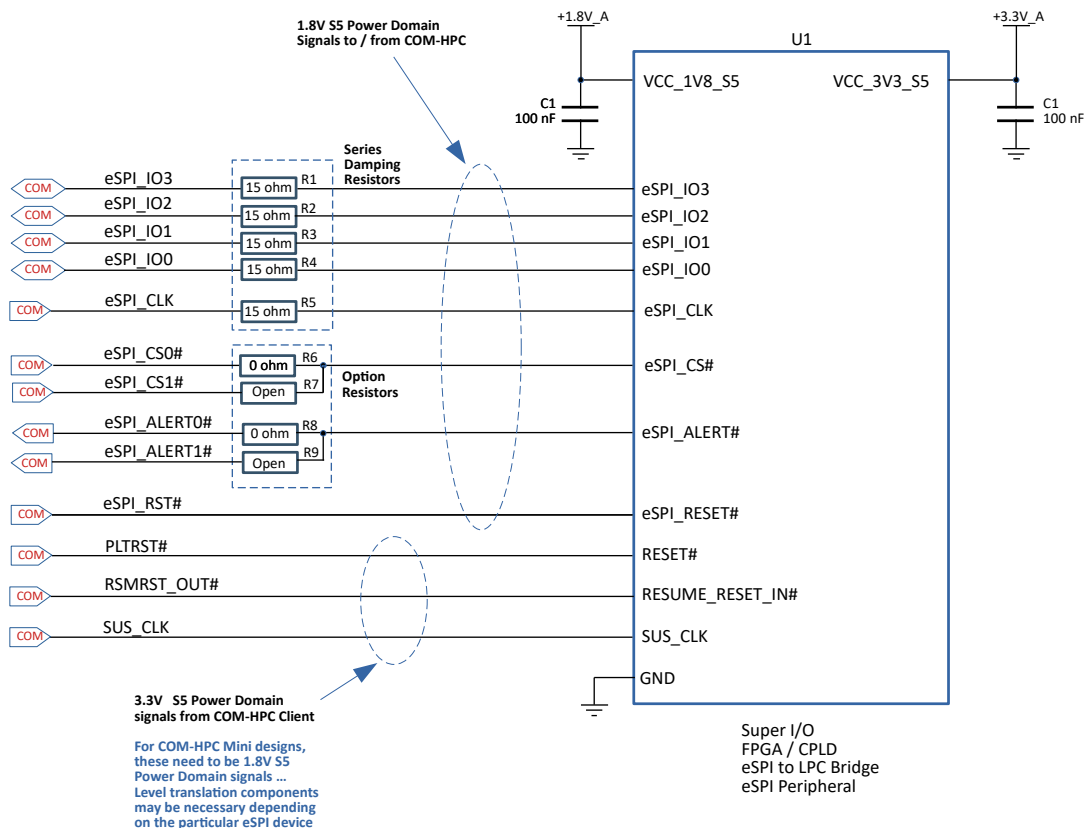
3.10. DisplayPort Over DDI

Figure 41: DisplayPort Over DDI
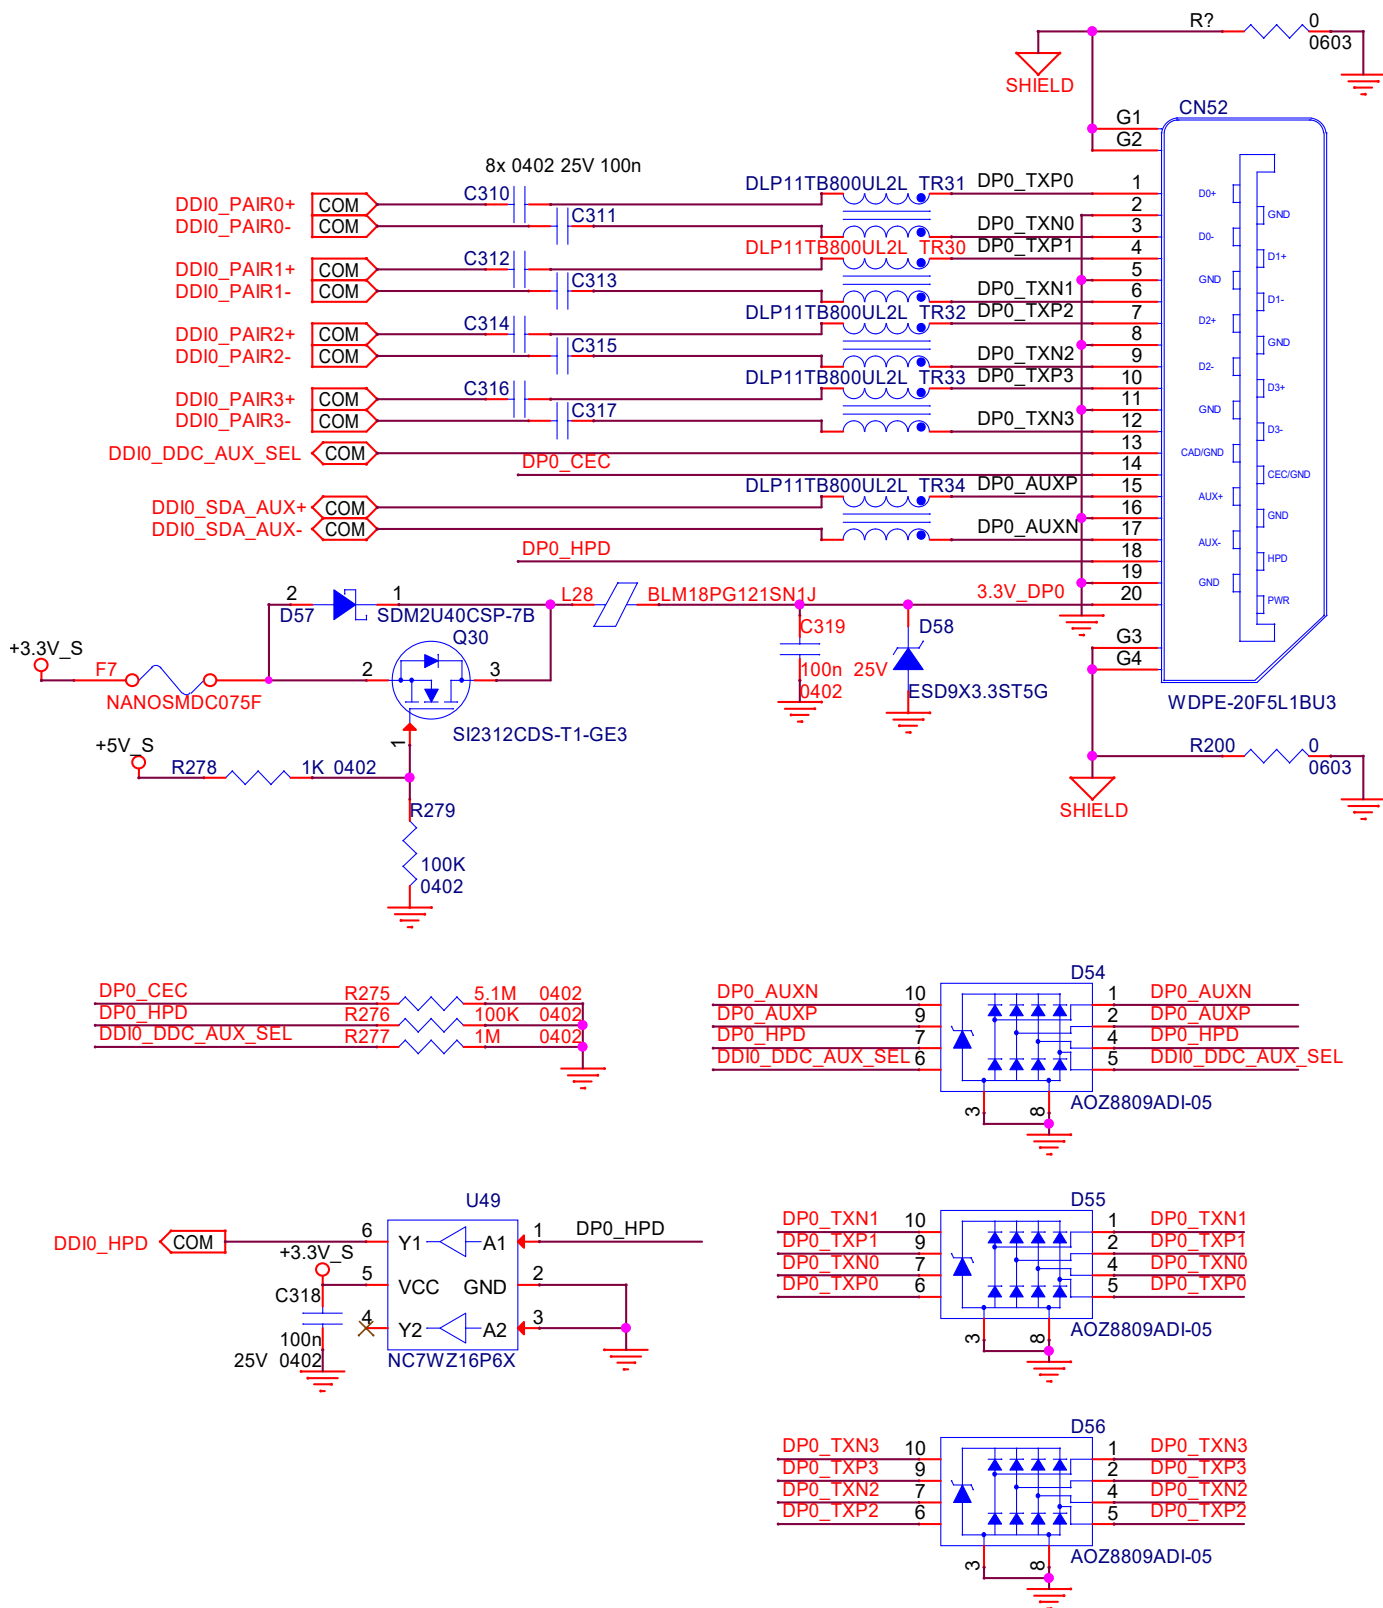
Notes on DisplayPort over DDI:

- COM-HPC supports three DDI channels, designated DDI0, DDI1 and DDI2. The Figure 41 above uses DDI0 as an example.

- DisplayPort data pairs are capacitively coupled, near the DisplayPort cable connector, as seen in the Figure 41 above (C310 through C317).
    - This is in contrast to HDMI data pairs which are typically DC coupled.
    - However, some HDMI buffers / level shifters / redrivers use AC coupling at the buffer inputs and DC coupling at the buffer outputs.

- The DisplayPort AUX channel data pair (net names DDI0_SDA_AUX+ and – in this example) are AC coupled on the COM-HPC Module, when the DDI channel is used in DisplayPort mode.
    - When the DDI channel is used in HDMI mode, the DDIx_SDA_AUX+ and – pair (where 'x' is 0,1 or 2) are DC coupled on the Module, for the HDMI SDA and SCL I2C setup channel.

- The COM-HPC signal DDIx_DDC_AUX_SEL signals are Module input signals that are used to select either DisplayPort or HDMI mode.
    - If the signal is pulled or driven low, or left NC, then the Module invokes DisplayPort mode.
    - If the signal is driven to a logic high, then the Module invokes HDMI mode.
    - In this schematic example, the DDI0_AUX_SEL Module input signal is pulled low by R277 and ESD protected by part of ESD diode array D54.

- Almost all connections to the DP connector CN52 in the Figure 41 above are provided with EMI suppression components (common mode choke elements TR30 through TR34) and ESD protection arrays (D54 through D56).
    - The EMI and ESD mitigation components used must be appropriate for the high data rates used by the DisplayPort data pairs.
        - For the ESD diode arrays, this means selecting parts with a sufficiently low pin capacitance.
        - For the EMI chokes, the selected parts should have a low differential impedance but a relatively high common mode impedance.
    - It is extremely important that all the nets the DisplayPort data path be routed as differential pairs, preferably against an unbroken GND plane and without any stubs, or with minimal stubs.
        - Note that the ESD protection arrays used in the example have 2 lands for each net being protected. This is to facilitate no-stub "flow through" routing.
        - The ESD diode arrays should be positioned next to the DP connector pins.
        - ESD diode array pins can be pin-swapped if needed to provide a cleaner PCB layout.
    - DP connector pin 18 is used as a "Hot Plug Detect" signal. The external display drives this signal to a logic high to signal a display hot plug event. This signal is ESD protected by an element of D54 and buffered and level translated by U49 before being passed on to the COM-HPC module. The buffer input is pulled down by R276 in the example, ensuring that the COM-HPC HPD input signal is low if no DP display is present.

Most of the information presented above for DisplayPort over DDI applies to the COM-HPC Mini as well, although the net names for the DDI related signals differ between the COM-HPC Client and the Mini and the AUX channel coupling caps are not on the Mini Module but rather on the Carrier.

The DDI net name correlations between the Mini and the Client are shown in the COM-HPC Base Specification Revision 1.20 Section 10.5.3 Tables 117 and 118.

The Mini AUX channel implementation with Carrier coupling caps is shown in the COM-HPC Base Specification Revision 1.20 Section 10.5.4 Figure 71.

Finally the Hot Plug Detect signal level to the Mini has to be level shifted to be a 1.8V signal by altering or replacing the U49 circuit in the Figure 41 above.

Some COM-HPC DisplayPort implementations may require a Carrier based redriver. A few industry offerings are listed in Table 16 below. There are of course more parts available on the market.

Table 16: DisplayPort Redrivers and Retimers

Vendor	P/N	Notes
Diodes Inc	PI3DPX1203B	4 lane DisplayPort 1.4 redriver; up to 8.1 Gbps link rate
	PI3DPX8121	DisplayPort 1.4 and 2.0 compatible 2:1 mux and redriver, 2 sets of 4 lane inputs and a 4 lane output, with up to a 10 Gbps link rate.
Parade Semiconductor	PS8463	DisplayPort 1.4 redriver (8.1 Gbps) HDMI 2.0 redriver (6 Gbps) 4 lanes
Texas Instruments	DS160PR410	This part is primarily a 4 lane PCIe Gen 4 capable redriver. However, the TI literature states that the part can be used for DisplayPort 2.0 redriver purposes, by setting a certain strap to disable the "PCIe Detect" mode. This is a very high bandwidth part and may work well with all DisplayPort modes.

There are quite a few USB Type-C and a few USB4 port multiplexers that incorporate redriver and in some cases retimer circuits. Such products come from Diodes Inc., Texas Instruments, and others. The Intel JHL9040R, also known as the "Hayden Bridge", does DisplayPort, USB and PCIe retiming along with other USB Type-C and Thunderbolt functions.

3.11. HDMI Over DDI

Figure 42: HDMI Over DDI
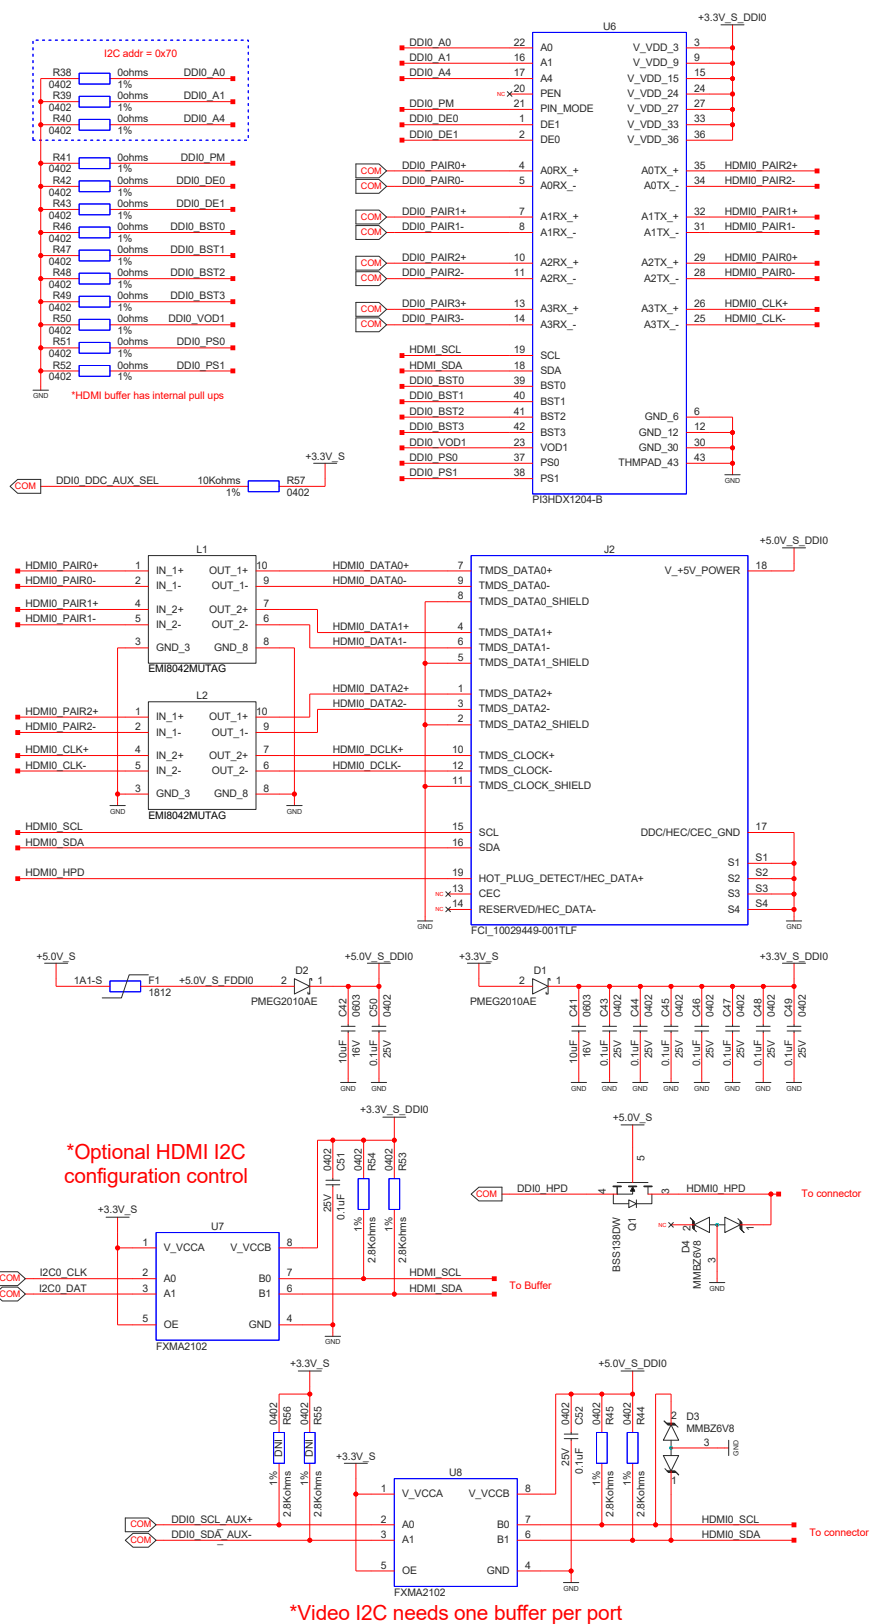
Notes on Figure 42: HDMI Over DDI Above

- COM-HPC DDI signals can generally operate in DP mode or HDMI mode.
- COM-HPC input signal DDIx_DDC_AUX_SEL (where x is 0, 1 or 2) selects between DP mode and HDMI mode.
  - DDIx_DDC_AUX_SEL left open or pulled low selects DP mode for DDIx.
  - DDIx_DDC_AUX_SEL pulled or driven high to +3.3V_S selects HDMI mode for DDIx.
- DP signals are AC coupled
  - DP data pairs are AC coupled on the Carrier near the DP connectors, as can be seen in Figure 41: DisplayPort Over DDI above.
  - DP AUX_SEL pairs are AC coupled on the COM-HPC Module (nets DDIx_SDA_AUX+ and – in Figure 41).
- HDMI signals are generally DC coupled – at least from the HDMI / TMDS driver outputs, across the HDMI cable and on to the HDMI / TMDS receiver. DC coupling is shown for the HDMI / TMDS data pairs and the SDA / SCL setup lines in Figure 42 above.
- Figure 42 above a Diodes Inc / Pericom PI3HDX1204B combination HDMI level translator and redriver (U6 in the Figure 42 above).
  - This part can be configured by resistor straps or over I2C. Both options are shown in the Figure 42 above.
  - Components L1 and L2 are On Semiconductor EMI8042MUT offering combined ESD protection and EMI suppression.
  - Note that the +3.3V level DDIx_SDA_AUX+ and – HDMI setup signals are translated to a +5V level with component U8. ESD protection is included for all signals facing the outside world.
- There are many alternative HDMI level translators on the market.
  - Texas Instruments, Analog Devices, Silicon Labs, Diodes Inc. and others offer HDMI level translators, redrivers and retimers.
  - Many devices have built in ESD protection and level translation for the HDMI data pairs and the SDA / SCL setup channel.
  - See Texas Instruments TPD12S016 for a basic HDMI level translator with integrated ESD protection.
  - Some HDMI redrivers / retimers use AC coupling at their inputs, and DC coupling to the cable at their outputs. See, for example, Texas Instruments TDP158.
- There may be licensing fees involved if HDMI implementations are used, and there are strict rules about logo use. Check with the HDMI organization (www.hdmi.org).

Notes for HDMI on the COM-HPC Mini:

- *The net names for the DDI related signals differ between the COM-HPC Client and the Mini.*

- *The correlation between the Mini "Super Speed" pairs and the DDI functions is found in the COM-HPC Base Specification Version 1.20 Section 10.5.3 Tables 117 and 118.*

- *The correlation between DDI and HDMI data pairs may be found in upper right of Figure 42 above, and in the COM-HPC Base Specification Version 1.20 Table 24.*

- *The HDMI DDC lines into the Mini need level translation from 5V (from the HDMI display) to 1.8V, and this must happen off Module.*

  - *This could be done by changing the supply voltage for the left side of U8 in Figure 42 above from 3.3V to 1.8V (the U8 FXMA2102 can handle this).*

  - *Many other I2C level translators are available.*

- *The Hot Plug Detect signal level to the Mini has to be level shifted to be a 1.8V signal into the Mini by altering or replacing the Q1 circuit in Figure 42 above.*

  - *One way to do this is to insert a 3.3V to 1.8V level shifter (such as Texas Instruments SN74LV1T125) in between Q1 and the COM-HPC Mini HPD pin. A 100K pull-down resistor in between the level shifter input and GND is needed, in case no HDMI display is present.*

    - *The Q1 "output" (the source pin at the left side of Q1) is only approximately 3.3V depending on the FET Vgs.*

    - *The SN74LV1T125 will be fine with an input above 3.3V.*

- *There are HDMI buffer / re-driver parts available now which can handle the 1.8V sideband signals to and from the COM-HPC Mini directly. See, for example, the Texas Instruments TDP0604. The DDC and HPC level translations are integrated into the part, and sideband signals can be 1.8V or 3.3V.*

3.12. eDP

Figure 43: eDP Schematic Example
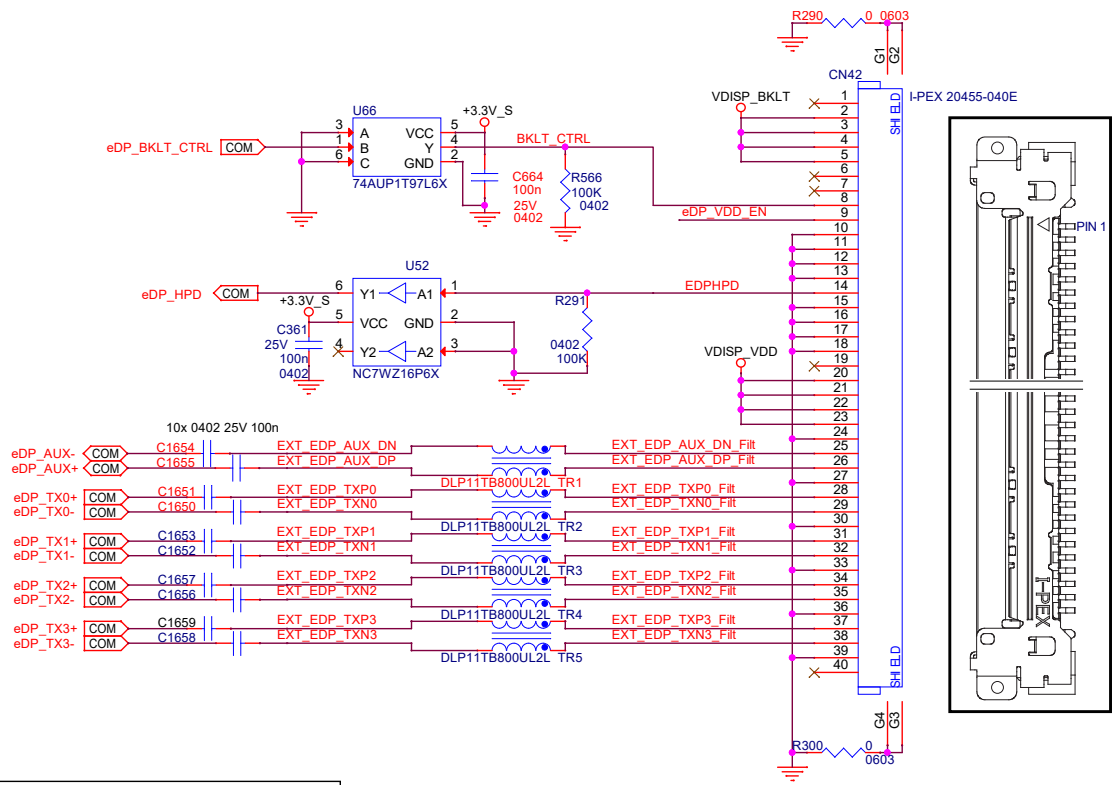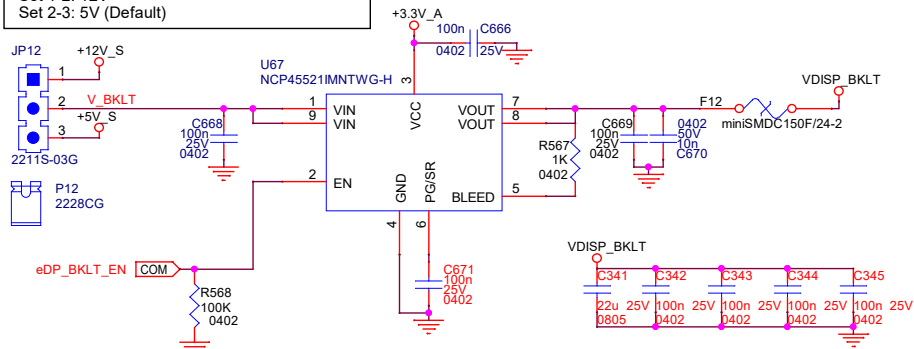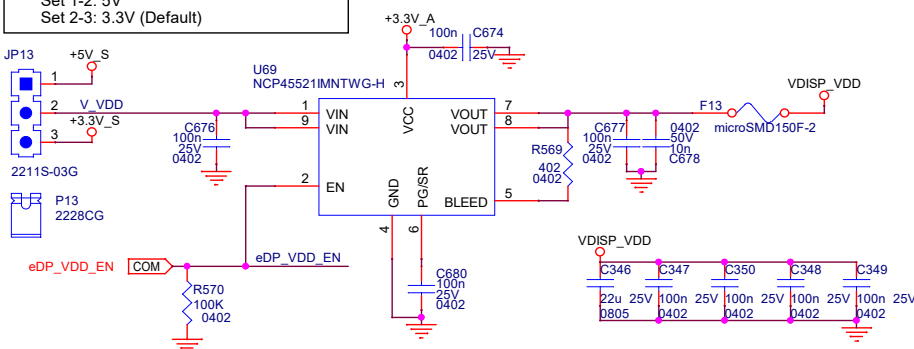
Figure 44:   eDP Connector Pin Numbering

The connector used for most eDP implementations is the I-PEX 20455-040E or equivalent. There can be some confusion about the pin numbering used in eDP systems. The connector vendor defines pin 1 at the left, as shown in the upper part of Figure 44 above. For reasons perhaps better lost to history, VESA and hence the eDP display vendors put pin 1 at the right side of the connector, as illustrated in the lower portion of the Figure, in spite of the datum mark at the left end of the connector.

The eDP schematic sample in Figure 43 above uses the I-PEX pin numbering (as PCB designers generally prefer to follow the component vendor's data sheet when making up PCB footprints). The net result here is that I-PEX pin 1 needs to map to VESA / Display pin 40, I-PEX pin 2 to VESA / Display pin 39 and so on. This happens with the straight through cable shown in Figure 43 above (which uses the I-PEX pin numbering on both ends of the cable). This cable works between a COM-HPC Carrier and a VESA eDP display (which uses the VESA pin order).

Display cables for eDP typically use micro-coax wiring. The + and – conductors of an eDP data pair travel in separate but adjacent coax lines. Hence they are not electromagnetically coupled within the cable assembly, but since each conductor is completely shielded and are equal length, the differential transmission properties are preserved and this works very well even at the highest eDP data rates. The PCB traces on the Carrier and within the display assembly should be edge coupled differential pairs, as per usual.

Additional eDP Example Material

An alternative eDP example implementation is presented in Section 6.2. , Appendix B: Alternative eDP Example near the end of this document.

The information in this Section applies to the COM-HPC Mini as well as the Client. However, the single ended side-band signals (eDP_BKLT_CTRL, eDP_HPD, eDP_BKLT_EN and eDP_VDD_EN) shown here are 3.3V signals on the COM-HPC Client Module but are 1.8V signals on the Mini. COM-HPC Carrier designs for the Mini have to accommodate this, if these lines are used.

3.12.1.   eDP / DP Conversions to Other Video Formats

There are a number of video format conversion bridges available from NXP (www.nxp.com), Chrontel (www.chrontel.com) and others. These products allow conversion from eDP or DP to LVDS, analog VGA, HDMI, DVI and a host of older video formats such as CVBS, S-Video, BT656, BT1120, YPbPr etc.

LVDS displays are not directly supported by COM-HPC. However conversion from an eDP or DP source (from COM-HPC) to LVDS input format displays is easily achieved using either the NXP PTN3460 (PTN3460I for the industrial temperature version) or the Chrontel CH7515. All common LVDS formats (single channel, dual channel, 16 / 18 / 24 bit color depths) are supported by these NXP and Chrontel parts.

COM-Express Modules from several vendors routinely use the NXP PTN3460I behind the scenes to produce the COM-Express LVDS outputs from the chipset eDP channel.

It may be wise to check with your Module vendor before selecting an eDP / DP conversion part, as the vendor may have a preference and have software / firmware support favoring a particular part. Some subtleties such as VESA EDID support, backlight control etc. may be easier using the video conversion part(s) supported by the COM-HPC Module vendor.

Analog VGA support is still important in some limited markets. The NXP PTN3355 and the Chrontel CH7517 are popular parts for this task.

This advice applies to COM-HPC Mini system designs as well.

3.13. MIPI-CSI Camera Interface

Figure 45: MIPI-CSI
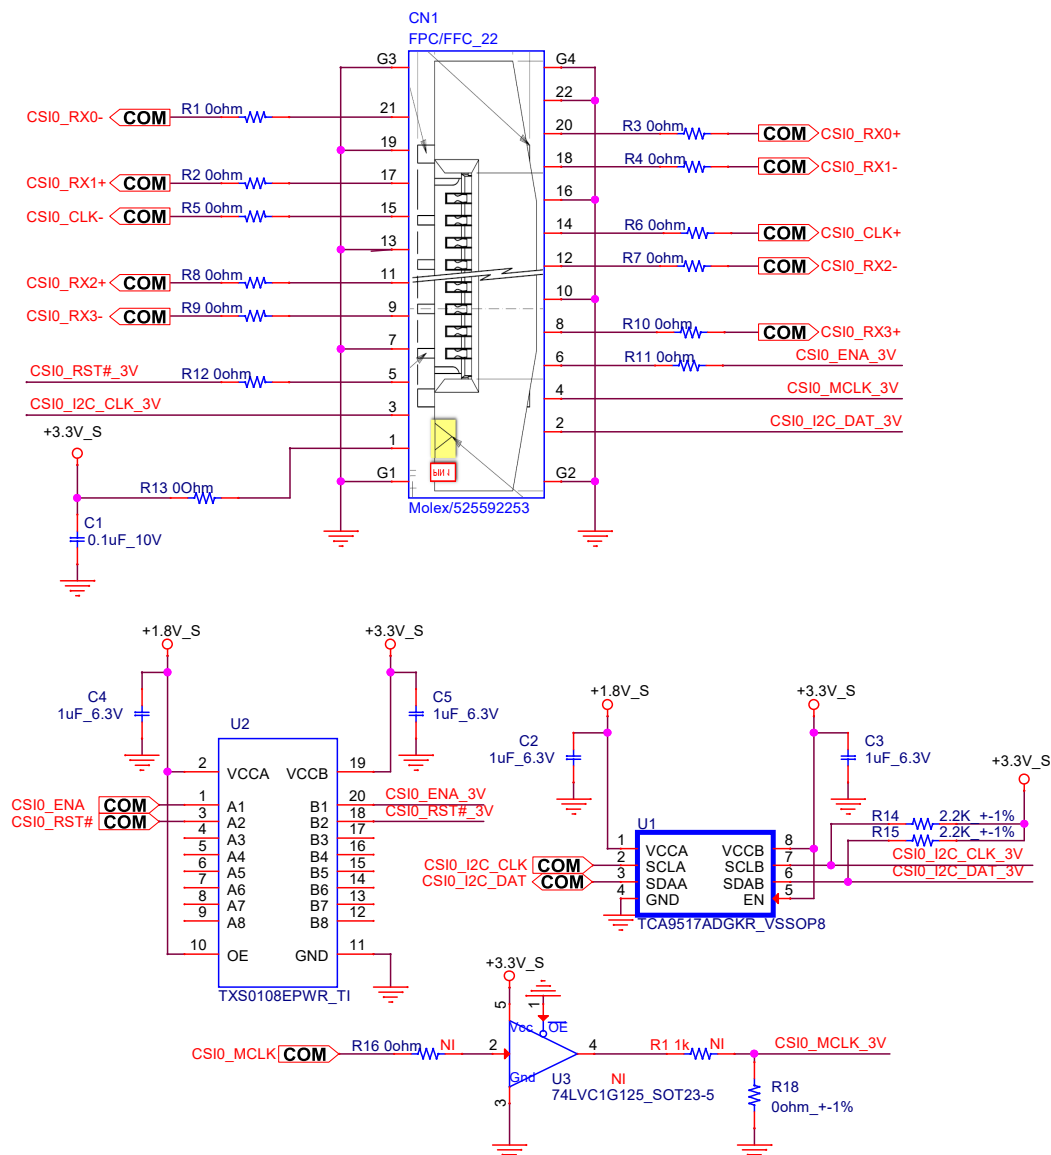
A typical Carrier board MIPI-CSI implementation is shown in Figure 45 above. There is no standard connector for MIPI-CSI use. The Molex part shown in the Figure 45 above is a reasonable choice but many others are used in various situations. The example above distributes +3.3V_S power to the camera, appropriate for many camera assemblies. However, many MIPI cameras are 1.8V devices, and the COM-HPC MIPI-CSI is defined as a 1.8V interface in the COM-HPC Base Specification document.

Camera and support software selection is an important part of implementing a MIPI-CSI system. The cameras have particular data formats and non-linear data compensation requirements to account for camera characteristics. It is important to have a software driver plan that aligns with the camera choice and the Module chipset or SOC choice. There may well be NRE charges from the Module vendor to get a MIPI-CSI camera solution working, unless the vendor has a "canned" solution to offer.

This information just above does NOT apply to the COM-HPC Mini. The Mini has separate FFC (Flat Foil Connector) connectors – separate from Module connector – see the Base Specification Revision 1.2 Section 10 for information on the Mini MIPI CSI implementation.

3.14. Audio Interfaces

3.14.1. General Discussion

The COM-HPC Client Module pinout allows for up to four SoundWire audio ports and one I2S audio port. No audio support at all is offered on the COM-HPC Server Module pinout. The first two COM-HPC Client Mode SoundWire ports, numbered as 0 and 1, are free and clear and are not shared. Port 0 and Port 1 are two pins each, with CLK and DAT lines. The 3rd and 4th COM-HPC SoundWire ports, numbered as 2 and 3, are pin shared with an I2S audio port.

No Intel HD Audio support at all is offered with COM-HPC revision 1.0. However it is to be offered in COM-HPC Base Specification revision 1.1 due to some delays in the industry SoundWire rollout.

SoundWire is expected to be the mainstream x86 system audio interface going forward. I2S audio interfaces are also available on many contemporary x86 chipsets. I2S is the most popular audio interface on ARM designs at the time of this writing. This may shift to SoundWire over time.

3.14.2. MIPI SoundWire Summary

- A 2 wire interface (CLK and DAT) is used.
  - For most implementations there is a single Master and there may be multiple Slaves
  - The CLK is an output from the Master
  - DAT is bidirectional signal, with data to and from Slaves
  - The Master controls the DAT line direction, per MIPI SoundWire protocol
- The SoundWire CLK and DAT lines may be run at 1.8V or 1.2V (per the MIPI specification)
  - COM-HPC uses 1.8V SoundWire signaling
  - This signaling should be available in all system states, S5 through S0
- There may be up to 11 Slave devices on a SoundWire bus
  - It is more common to have up to 4 Slave devices on a single SoundWire bus
  - There is a MIPI defined enumeration process to identify the Slaves
    - It involves a bit of trial and error but in time all Slaves are identified
- Some details on SoundWire clocking and signaling include:
  - The CLK frequency used is set by the Master, and may be as high as 12.288 MHz
    - The lowest appropriate frequency is used
  - DDR (Double Data Rate) signaling is used (meaning that data is clocked on the rising CLK edge and the next bit on the falling CLK edge)
  - The CLK frequency may be slowed or completely stopped by the Master, as required
    - These clocking / data features allow lower power operation
  - SoundWire uses a "modified" NRZI (Non Return to Zero Inverted) protocol on the data line
    - This allows the enumeration capability and other features described in the MIPI specification
  - Audio data may be encoded in several formats:
    - PCM (Pulse Code Modulation) – the most common format
    - PDM (Pulse Density Modulation) – has low hardware implementation overhead and is useful for simple devices such as digital microphones
    - Bulk Mode – for large data blocks
  - Slaves may initiate in-band interrupts and wake events
  - I2C support for SoundWire devices is generally not needed (unlike for I2S)

MIPI Slave devices may be wired together either in daisy-chain fashion or in a branched – tree topology, as shown in Figure 46 below. In either case, the SoundWire CLK and DAT lines should be routed together – not as a differential pair, but as a signal pair following approximately the same route paths, with approximate length matching all along the paths and for each branch, so that the signal flight time from Master to Slave for both the SoundWire CLK and DAT are about the same. If the balanced tree topology is used, the length of branches of the tree should be about the same.

Chipset design guide examples tend to show a point to point SoundWire implementations with signal integrity measures. These include series damping resistors and snubbing capacitors, as depicted in Figure 47 below, The component values are design and layout dependent and may range from 0 to about 22 ohms for the series resistors and from 0 (i.e. not loaded) to about 22 pf for the capacitors.

The MIPI Master – to Slave implementation is straightforward, as it only involves the CLK and DAT lines. There are sure to be many more CODEC or MIPI Slave device implementation details – such as filtered analog power supplies, decoupling and other component recommendations, analog layout recommendations etc., not covered here. This information is available from the CODEC and Slave device vendors.

Table 17: SoundWire Audio CODECs

Vendor	Vendor P/N	Notes
Cirrus Logic	CS42L42	SoundWire and I2S Audio CODEC – data freely available on the web
Realtek	ALC711-VD	SoundWire and I2S Audio CODEC – data restricted at time of this writing

There are quite a few SoundWire Slave devices available, such as microphones and amplifiers, that are simpler than full CODECs. Vendors include Analog Devices, Maxim Integrated Products, TDK, Texas Instruments and more. From a hardware compatibility view, these low end devices may be tied directly to one of the COM-HPC SoundWire ports – but be sure to check out the software support situation before putting hardware together.

The reference designs from some x86 SOC and chipset vendors show SoundWire device implementations grouped into functions. For example, the first SOC or chipset SoundWire bus may host two or more output amplifiers, the second SoundWire bus an audio CODEC, and the third SoundWire bus hosts an array of SoundWire microphones.

Check with your COM-HPC Module vendor to see if they have any specific SoundWire device recommendations and port mapping recommendations.

Intel SoundWire Sample Schematics and Design Guide

Sample SoundWire implementations may be found in some Intel reference schematics. See, for example, NDA protected Intel document numbers 627073 and 627205.

Intel NDA protected document number 627205 devotes several pages to SoundWire design They basically show a "balanced tree" with two branches of approximately equal length, and with two SoundWire loads. An alternative daisy chain arrangement with up to four loads is described as well. Also recommended are some optional series damping resistors. In the COM-HPC case, the optional damping resistors would be placed in the SoundWire clock and data lines near the COM-HPC connector.

Figure 46: MIPI SoundWire Routing Topologies
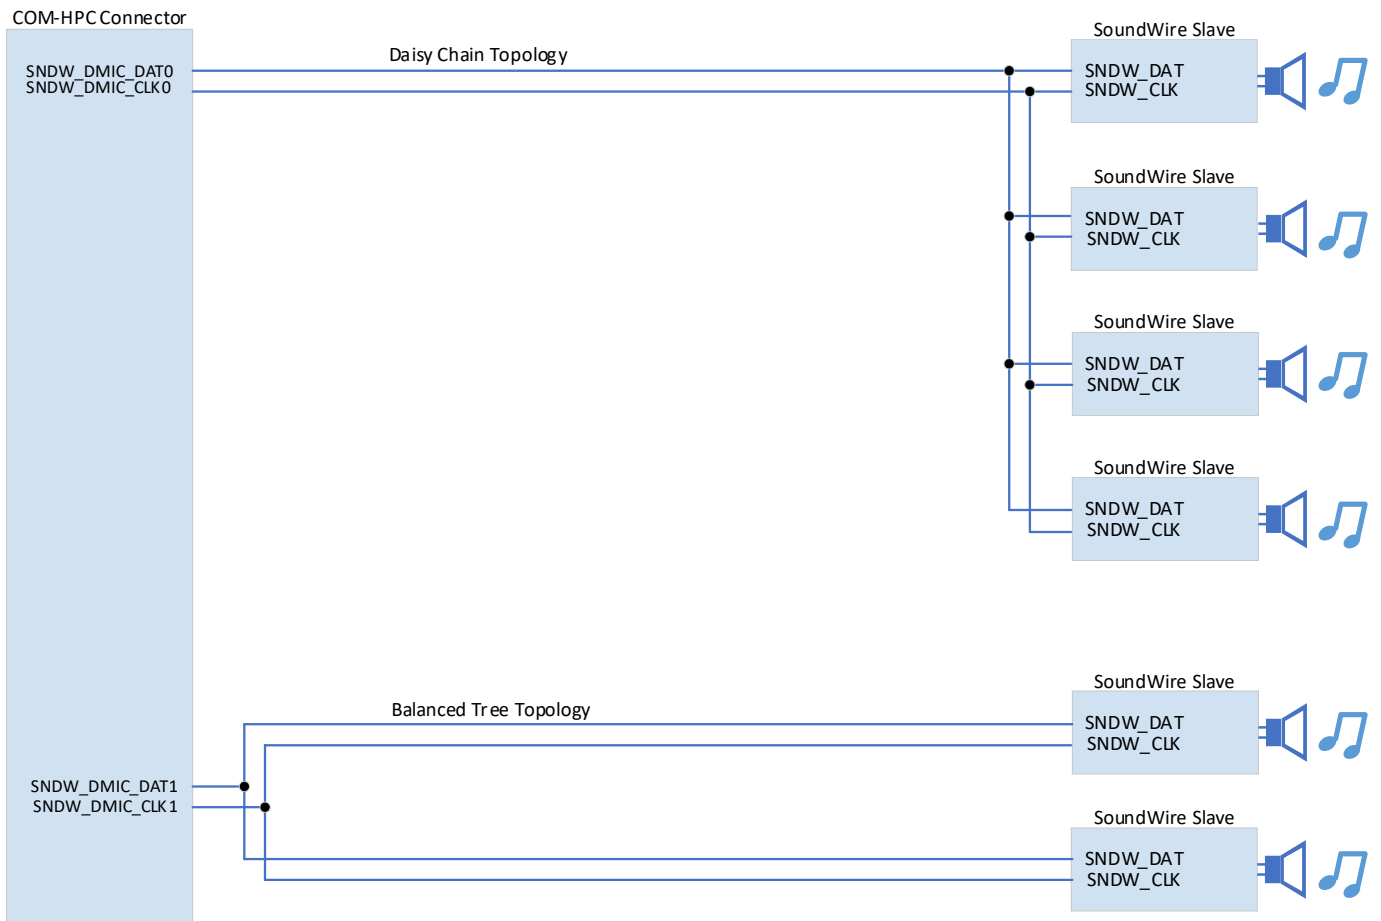
Figure 47: MIPI SoundWire Point to Point Connection With SI Components

3.14.3.   I2S Implementations on COM-HPC

An I2S audio implementation example specifically for COM-HPC is not available at the time of this writing. There are two I2S audio CODEC examples in the ***SMARC Design Guide*** that may be useful for reference:

- Cirrus Logic WM8904 Ultra Low Power CODEC
- Texas Instruments TLV320AIC3105 Low Power CODEC

SMARC defines two I2S ports versus a single I2S port defined for COM-HPC. Apart from that, the definitions are very close:

- 1.8V logic level signaling.
- S0 power domain operation.
- Same signal definitions (although pin names do not quite match):
    - An I2S clock out pin from the Module to a Carrier Slave.
    - An I2S data out pin defined.
    - An I2S data in pin defined.
    - An I2S Left – Right audio channel clock output pin defined.
    - An I2S audio master clock output defined.

I2S implementations generally require a companion I2C interface to setup I2S device registers. This is evident in the SMARC sample drawings. The digital I/O levels for I2S and I2C on an I2S CODEC are generally the same. The COM-HPC I2S interface is a 1.8V interface, and hence the I2C interface used would need to be at 1.8V. The COM-HPC I2C1 interface is defined to be a 1.8V interface; the COM-HPC I2C0 is a 3.3V interface. Of course level translation can be implemented.

Intel NDA protected document 616553, a schematic for an Elkhart Lake validation platform, shows an I2S CODEC implemented in a x86 based system. Elkhart Lake is an Atom class SOC and is not likely to be implemented on COM-HPC Client but may be implemented on a COM-HPC design. In any case the Intel eI2S example may be useful to designers looking to implement I2S audio on a COM-HPC Carrier.

SoundWire does not require a companion I2C interface. Just the SoundWire Clock and Data lines are sufficient for both audio data and SoundWire slave register setup.

The information in this Section applies to the COM-HPC Mini as well. The SGET SMARC Design guide is a useful document to find I2S CODEC and amplifier implementation examples, and many other 1.8V I/O parts.

3.14.4.   High Definition Audio (HDA) on COM-HPC Client and Mini

An Intel HD Audio option was introduced to the COM-HPC Client on the COM-HPC Base Specification Rev. 1.10, due to the delayed rollout of SoundWire. The COM-HPC Mini supports HD Audio as well, as introduced in Base Specification Rev. 1.20

There are good Carrier Board HDA example designs available in the PICMG COM Express and the SGET SMARC Design Guides, although the CODECs shown are quite old. The net names between the COM-HPC Module and the target Carrier CODEC do not agree with the COM Express and SMARC Design Guides, but there are only 5 signals involved and the correlation is straight forward.

3.15.   Asynchronous Serial Port Interfaces

3.15.1.   COM-HPC UART Interfaces

Two 3.3V logic level asynchronous serial ports, designated UART0 and UART1 are defined by COM-HPC. Each port has TX and RX signals for data use and RTS# and CTS# signals for optional handshake / flow control use. For logic level use, the TX and RX signals are active high and the RTS# and CTS# signals are active low. Some data sheets omit the trailing '#' signal but the logic level handshake signals are active low nonetheless. The idle state, or 'mark' state, of the logic level TX line is high, or 3.3V in the COM-HPC case.

These ports may be used directly as logic level asynchronous serial connections between COM-HPC Module and Carrier based devices, or between COM-HPC Module and Carrier based mezzanine devices such as certain Mini-PCIe or M.2 cards. Care has to be taken that the logic I/O levels match up. Note, for example, the (unused) UART connections on the left side of the M.2 E-Key card shown in Figure 14 above. The PCI-SIG M.2 specification defines the E-Key UART pins to be 1.8V signals so some non-inverting level translation would be needed in this case: 3.3V to 1.8V on the TX and RTS# lines leaving the COM-HPC Module, and 1.8V to 3.3V translation for the RX and CTS# lines coming into the COM-HPC Module. Dozens of suitable logic level translation products are available on the market. One such product is the Texas Instruments SN74LV1T125.

For off-board, cabled connections, the logic level UART signals are usually translated into one of three common formats: RS-232, RS-422 or RS-485. RS-232 is a single ended format in which the 'mark' or 'idle' or 'logic 1' state is a negative voltage between -3V and -25V, and the 'space' or 'logic 0' state is a positive voltage between +3V and +25V. An RS-232 level translation implementation for UART0 and UART1 is shown in Figure 48 below. This example uses a pair of Maxim (Texas Instruments) MAX3243E level translators. These parts have built in capacitor based charge pumps that create RS-232 compliant voltage levels and avoid the need to distribute a negative voltage on the Carrier. Note that the signal polarities are inverted by the device. This particular Maxim device has built-in +/-15 kV air gap and +/-8 kV contact ESD discharge survivability. There are many similar devices from Texas Instruments, Analog Devices / Linear Technology, Diodes Inc. and others. The cable length that can be achieved with RS-232 interfaces depends on the data rate used. Generally, RS-232 cables lengths are limited to about 50 feet or less.

RS-232 signals are most often used with D subminiature DB-9 or DB-25connectors. The RS-232 standard defines DTE (Data Terminal Equipment) and DCE (Data Communications Equipment) connector pinouts. A DTE chassis connector is a male connector, and a DCE chassis connector is female. The DTE and DCE pinouts are defined such that a straight cable (pin 1 to pin 1, pin 2 to pin2 etc. on the cable) may be used. With the straight cable, the DTE TX pin (or TX# pin if using that notation) lands on the DCE RX(or RX#) pin, and so on.

For longer cable lengths, on the order of 1000 feet or more, differential signaling formats such as RS-422 or RS-485 are often used. These implementations are usually terminated in the twisted pair cable impedance at the receiving endpoints. Many suitable parts are available from Analog Devices / Linear Technology, Texas Instruments and others. These vendors offer very informative Application Notes. They also offer "multi-protocol" devices – devices that can handle RS-232, RS-422 and / or RS-485 hardware protocols. Some of these devices have switchable internal cable termination. Some products from these vendors offer galvanic isolation.

The RTS# handshake line is often used in RS-485 implementations as a transceiver enable line. This of course needs appropriate software support.

The COM-HPC Mini asynchronous serial ports are 1.8V level ports and the RS-232 level translators shown in the Figures below are not suitable as they are 3.3V parts. Suggestions for alternative parts are given on the page below.

Figure 48:   UART0 and UART1 RS-232 Level Translated Asynchronous Serial Ports
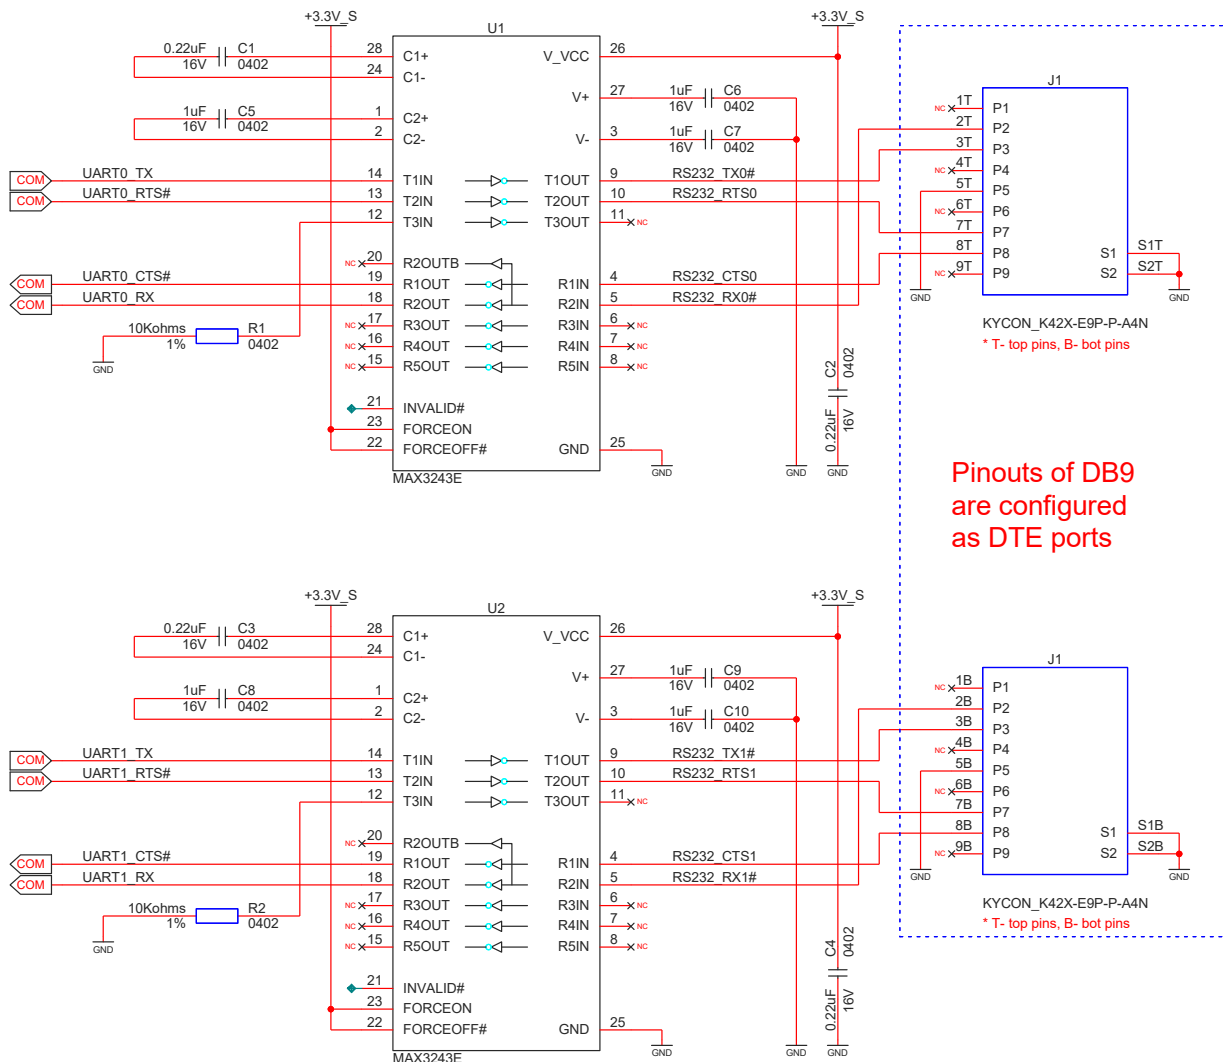
The example presented just above for RS-232 level translation is NOT suitable for the COM-HPC Mini implementations as the Texas Instruments MAX3243 does not operate with a 1.8V supply voltage and chipset I/O voltage. However, there are dozens of parts available that do operate at 1.8V levels.

One example is the Texas Instruments TRS3122E. There are many other possibilities from TI.

Another useful vendor for these types of parts is Analog Devices. The AD LTC2802 and LTC803 are some of the many possibilities.

The SGET SMARC Design Guide shows additional 1.8V I/O possibilities.

3.15.2. Legacy Compatibility With 16C550 UART Register Set

The I/O mapped UART that was the de-facto standard defined at the dawn of the personal computer age is the National Semiconductor (now Texas Instruments) 16550 or 16C550. Many BIOSes support 16C550 operations early in the BIOS boot (before USB devices are enumerated). Console redirect and Port 80 debug codes are often directed to a 16C550 compatible I/O register set. Windows, Linux and other popular operating systems used in embedded system almost universally support the 16C550 UARTs. The COM-HPC Base Specification encourages but does not require 16C550 register compatibility for the UART0 and UART1 ports. Check with your Module vendor.

Once the operating system is running and drivers are loaded, 16C550 compatibility is a non-issue, but for early boot support it is valuable.

3.15.3. Alternative / Additional Carrier Board UART Implementations

If additional or perhaps higher performance UARTs beyond what the COM-HPC Module provides are needed, there are a number of excellent options available. A few of these are summarized in the Table below..

Table 18: Alternative / Additional Carrier Board UART Implementations

Vendor	Interface	Sample Vendor P/Ns	16C550 Compatible	Features / Notes
FTDI	USB 2.0 FS	FT232RUSB	No	Future Technology Devices Inc Web www.FTDIchip.com Several similar parts available Win 10 and Linux drivers
MaxLinear (Exar)	USB 2.0 FS	XR21V1410 XR21B1420	No	Web www.maxlinear.com
MaxLinear (Exar)	PCIe x1 Gen 2	XR17V352 XR17V354 XR17V358	Yes	Web www.maxlinear.com Dual, Quad and Octal parts Very deep FIFOs, high bit rates Native Windows and Linux support Vendor drivers also available RS485 support
Microchip	USB 2.0 FS	MCP2220	No	Web www.microchip.com

3.16. I2C / I3C Ports

The COM-HPC pinout definition supports a traditional I2C port, designated I2C0, and a second port designated I2C1, targeting (optional) MIPI I3C use along with backward compatibility to traditional I2C.The I2C0 port runs at 3.3V and I2C1 at 1.8V. Both are active in suspend and full-on power states.

I2C is an abbreviation for "Inter Integrated Circuit". It was defined by Philips (and later inherited by NXP) as an easy to use two wire method for a Master device to set and read back Slave peripheral IC registers and data values. It uses, in it's basic form, open-drain drivers and passive pull-ups. The current NXP specification document is freely available (see reference in Section 1.9.) and defines several modes of operation, summarized in the following Table:

Table 19: I2C Operating Modes

I2C Mode	Operating Frequency	Max Rise Time	Max Bus Capacitance	Notes
Standard	100 KHz	1000 nsec	400 pF	
Fast	400 KHz	300 nsec	200 pF (passive pull-up) 400 pF (active pull-up)	See NXP UM10204 Section 5.2
Fast Plus	1 MHz	120 nsec	550 pF (active pull-up)	
High Speed	3.4 MHz			3 Mbps throughput Not described in NXP UM10204 Referenced in some literature including Intel

The COM-HPC Base Specification V1.0 document recommends a 2.2K ohm on-Module pull-up on the I2C0 and I2C1 Clock and Data lines. This value is sufficient for all Standard Mode (100 Khz) situations as the RC time constant is 880 nsec (= 2.2K * 400 pF) in the worst case, under 1000 nsec. In most situations, the I2C bus capacitance is much lower than 400 pF:

- Typical IC pin capacitances are 6 to 8 pF

- A typical PCB trace capacitance is 4 pF / inch – this varies with stackup details

- If, for example, there are 10 devices on the bus and there is a 20 inch total trace length, the bus capacitance would be about 160 pF (= 8 pF * 10 + 4 pF / inch * 20 inch)

- This rough calculation includes both the Module and Carrier I2C devices and trace lengths, with about 5 inches assumed on the Module.

The 2.2K Module pull-up is not sufficient for the worst case Fast Mode (400 Khz) passive mode bus capacitance of 200 pF. The 2.2K value handles up to about 100 pF of bus capacitance. A 2nd, parallel set of 2.2K pull-ups on the Carrier I2C Clock and Data lines would be advisable if the bus loading is over 100 pF. The Carrier pull-ups can always be left unpopulated if they are not needed.

There are various application notes on this subject available on-line. See, for example, Texas Instruments document SLVA689 titled "*I2C Bus Pullup Resistor Calculation*".

The Module design may include active circuitry to better support I2C Fast Mode and to support Fast Mode Plus. Check with the Module documentation or with the Module vendor.

3.16.1. I2C Addressing

I2C uses a 7 bit addressing scheme to differentiate I2C resources. The address lines are designated A6 … A0, but they are part of an 8 bit frame, in bit positions 7 through 1. Bit position 0 is used in the I2C device address to designate whether the current operation is a Read (a logic '1') or a Write (a logic '0'). Thus I2C addresses can be described in either 7 bit or 8 bit formats. If using the 8 bit format to describe I2C addresses, the R/$\overline{W}$ bit is always assumed to be '0'. The 8 bit frame is transmitted MS bit first.

The 4 most significant bits in the 7 bit I2C address field are used to define I2C device categories that are fixed by the I2C specification. The 3 least significant bits allow up to 8 devices within an I2C category to be identified by the I2C silicon vendor and / or the user. There are usually pin straps or specific product SKUs that define the 3 bit LS bit addresses.

It is useful to put together a table of I2C devices and addresses used in a design on both the Module and the Carrier, to ensure that there are no conflicts and to provide the information to software engineers. The Module vendor should provide such a list for I2C devices on the Module that are exposed / accessible on the COM-HPC I2C0 and I2C1 buses.

Two I2C memory devices for COM-HPC use, at specific I2C0 addresses, are designated in the COM-HPC Base Specification, Version 1.0, Section 5.2:

- A Module EEPROM device at hexadecimal address 0x50 (7 bit I2C addressing) or 0xA0 (8 bit addressing)
- A Carrier EEPROM device at hex address 0x57 (7 bit addressing) or 0xAE (8 bit addressing)

The Module and Carrier EEPROM data structures and contents are described in the PICMG documents *EeeP for COM-HPC (Embedded EEPROM for COM-HPC)* and the *PICMG COM-HPC Platform Management Interface Specification*.

Figure 49: I2C0 Example: Carrier EEPROM in S5 Power Domain
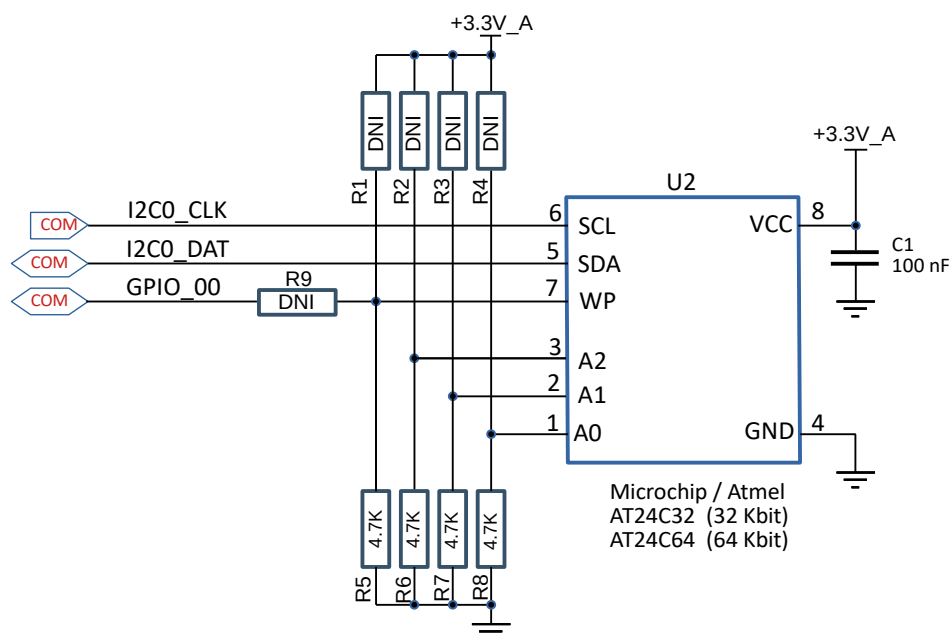
For COM-HPC Mini use, this Figure 49 needs to be adapted to 1.8V I/O levels. The parts listed for U2 above can be used with a 1.8V supply. The resistor pull-ups need to be pulled to 1.8V as well.

3.16.1.   I2C0 Example: Carrier I2C Device in S0 Power Domain

The COM-HPC I2C0 port is a basic 3.3V I2C port active in all power states (S5 through S0). It is a three wire port with CLK, DAT and ALERT# pins. The ALERT# pin, if supported by the COM-HPC Module, can serve as an interrupt input to the Module. An implementation example is given in Figure 50 below. In this example, a temperature sensor that is in the S0 power domain (+3.3V_S power net) is connected to the S5 domain COM-HPC Module I2C0 port through isolation FETs T2, T3 and T4. These FETs serve to prevent the S5 power domain I2C0 port from being dragged down when the +3.3V_S power rail is absent. I2C0 bus pull-up resistors to the +3.3V_A (Always on) power rail are on the COM-HPC Module. Secondary pull-ups to the +3.3V_S (Switched) power rail are shown in the Figure 50 below. The ALERT# pin on the IC shown in the Figure is asserted by the LM75B device when a temperature threshold that had been previously set using the I2C0 interface is crossed.

The I2C address of the LM75B device is set at 0x90 (8 bit addressing). The lower 3 bits of the LM75B I2C address are set by the A2, A1,A0 pin straps. Up to 7 additional devices could be deployed by setting different addresses options on these pin straps.

Figure 50: I2C0 Example: Carrier Temperature Sensor in S0 Power Domain
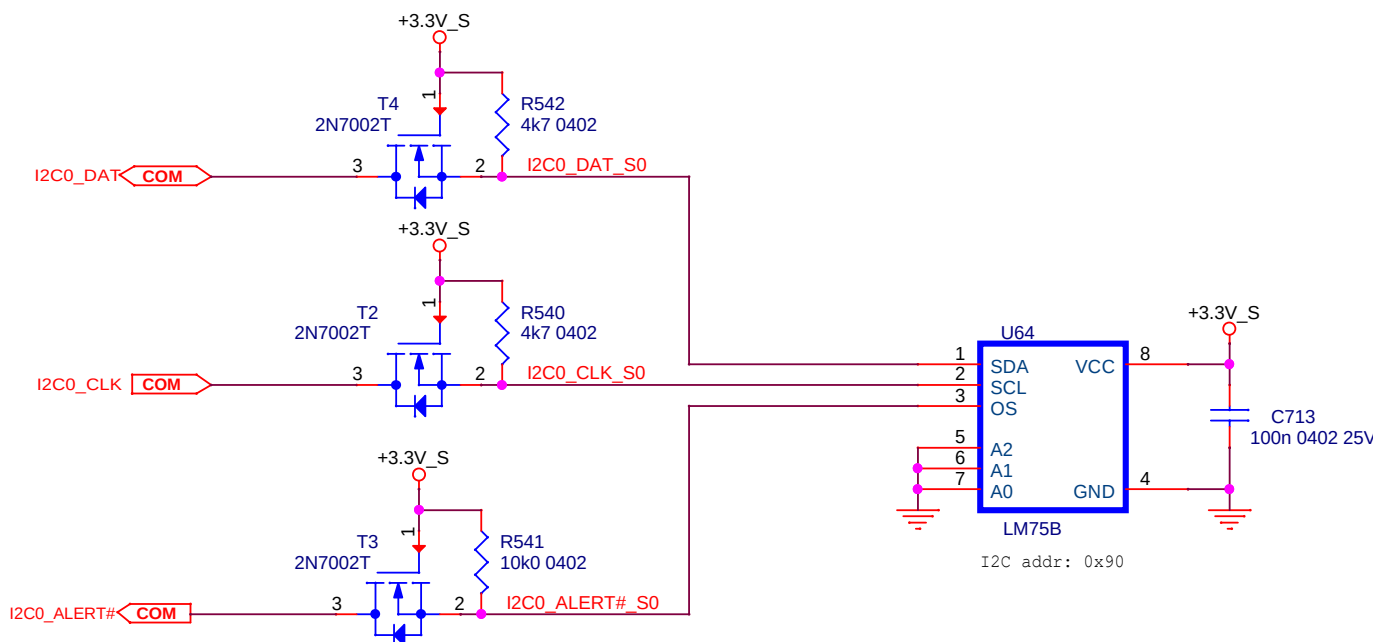
Note: the COM-HPC Mini I2C0, I2C1 and I2C2 ports are defined as 1.8V ports, so the Figure 50 above, using 3.3V supplies and parts does not directly apply to the Mini.

Newer parts such as the Texas Instruments TMP1075 I2C temperature sensor can operate with a 1.8V supply.

N channel FETs with a lower Vgs than the 2N7002 would need to used if this circuit is used for the Mini. The Rohm RA1C030LD is one of many possibilities. This part has a Vgs ranging from 0.5V to 1.5V.

Alternatively, the N FETs in the Figure 50 above may be replaced with I2C level shifters for the Mini. There are dozens of such parts. The TI / NXP PCA9306 is an example. This is addressed in Section 3.16.2. just below.

The SMARC Design Guide may be a valuable resource as well for identifying 1.8V capable parts.

3.16.2. I2C Bus Buffers / Level Translators

There are a number of bus buffers and level translators that target I2C and SMBus situations. Some suggested vendors and parts to consider are listed in Table 20 below. These parts allow for power domain isolation as the I/O pins go into a high impedance mode if one or both of power rails collapse. There are many other possibilities.

Table 20: I2C Bus Buffers / Level Translators / Power Domain Isolation

Vendor	Part	Notes
Texas Instruments	TCA9517	An upgrade and replacement for the popular NXP PCA9517
On Semiconductor	FXMA2102	
Texas Instruments / NXP	PCA9306	

3.16.3.   I2C1 (COM-HPC) and Optional I3C Support

The COM-HPC Base Specification defines a second I2C port, designated I2C1. This port is a 2 wire port (clock and data; no ALERT#) that operates from the 1.8V S5 and S0 power rails. The COM-HPC specification states that this port supports I2C and optionally supports I3C operation.

I3C is mostly backward compatible with I2C but there are some caveats and differences, summarized in the section just below. If the user "just" wants an additional I2C port, than the I2C / I3C differences are not important and the user may proceed with a traditional I2C implementation, bearing in mind the 1.8V operating voltage and the S5 / S0 power domain. If combined I3C / I2C operation is expected, then the material presented just below is important.

MIPI I3C Discussion

The MIPI Alliance has defined a significant enhancement to traditional I2C, known as I3C, an acronym for "Improved Inter Integrated Circuit" communication. I3C is significantly faster than I2C, and has some notable feature improvements, summarized below. It is largely, but not completely, backward compatible with I2C, also noted below.

Recall that I2C is a 2 wire interface (with an optional 3rd wire for an ALERT# input) that operates in most cases as a 100 KHz or 400 KHz interface, with some 1 MHz and 3.4 MHz implementations.

I3C enhancements beyond I2C include:

- 12.5 Mbps SDR (Standard Data Rate) operation using a 12.5 MHz clock
- 25 Mbps DDR (Double Data Rate) operation (12.5 MHz clock, using rising and falling clock edges)
- 33.3 Mbps Ternary Encoding operation (too complicated to explain here; see the MIPI documentation)
- Higher bandwidth and lower power operation
    - Active pull-ups rather than passive pull-ups increase speed and lower power consumption
- In – band interrupts (ALERT# pin not needed) – 2 wire operation only
- Error detection
- Error correction, in the Ternary mode
- Dynamic addressing
- Hot – Join operation (devices may be powered down and rejoin at power up; not same as Hot-Plug)

However, there are some issues with complete backward compatibility between I3C and legacy I2C:

- I2C devices on an active I3C bus need "50 nsec spike filters" in series with their Clock and Data pins to prevent the legacy I2C devices from getting confused by some fast short signals ("spikes") present in I3C traffic.
    - The spike filters can be simple RC circuits (series resistor in front of I2C device pin, capacitor between IC pin and GND).
- I2C clock stretching is not allowed on an I3C bus
- I2C bus 10 bit addressing mode is not allowed in I3C

Check with your COM-HPC Module vendor for details about their possible I3C support on the COM-HPC I2C1 port.

3.16.4. I2C on COM-HPC Mini

The COM-HPC Mini supports three general purpose I2C ports, designated I2C0, I2C1 and I2C2. This is one more general purpose I2C port than the COM-HPC Clients and Servers have. On the Client and Server Modules, I2C0 is a 3.3V interface and I2C1 is a 1.8V port.

On the COM-HPC Mini, I2C0, I2C1 and I2C2 are all 1.8V ports.

The I2C2 port is targeted to support SGMII PHYs in COM-HPC Mini systems. The actual net names are I2C2_CLK/ETH_MDIO_CLK and I2C2_CLK/ETH_MDIO_DAT.

COM-HPC Mini SGMII implementation example in Section 3.25. below. It shows the I2C2 / MDIO pins used for SGMII PHY support.

3.17. Port 80h Debug Display Over COM-HPC USB_PD_I2C

Figure 51: Port 80h Debug Display Over COM-HPC USB_PD_I2C

The COM-HPC hardware specification allows for BIOS Port 80h debug codes to be serialized and transmitted over the USB Power Delivery I2C bus (COM-HPC pins USB_PD_I2C_DAT and _CLK). Figure 50 above illustrates how the codes can be de-serialized and displayed on a pair of 7-segment LED displays. This feature is optional.

There are other methods for Port 80h codes to be conveyed for debug use. The Port 80h I/O writes can be picked off of the eSPI bus or even a PCIe x1 link by appropriate hardware such as a CPLD, FPGA or some Super I/O devices. Some BIOSes provide 4 digit codes as opposed to 2 digit codes.

It is also possible for special debug versions of a BIOS to transmit ASCII versions of the Port 80h debug codes over one of the asynchronous serial ports … check with your Module vendor.

This circuit may be used on a COM-HPC Mini system, with a small modification: the PCA6416A device can be configured for 1.8V I2C operation: U58 pin 23 and it's bypass capacitor C702 in the Figure 50 above need to be moved off the 3.3V power rail and onto a 1.8V rail.

3.18. Carrier BMC with IPMB Link to Module

A high end Carrier BMC (Board Management Controller) using an Aspeed AST2500 / AST2520 is shown in Figures 52, 53 and 54 below. The three Figures do not show the complete implementation – the DDR4 memory devices and the Aspeed power section are omitted. What is shown are the features relevant to COM-HPC operation – the interfaces to the COM-HPC and to the user. Refer to the Aspeed documentation for complete design information.

The BMC interfaces to the Module include IPMB, eSPI, UART1, I2C0 and a collection of status and control signals such as power state status, reset, power button etc. BMC operator interfaces include a USB port for keyboard and mouse use, a VGA port, and a 1000BASE-T management network interface.

The primary management interface to the COM-HPC Module is over the IPMB. The Module, if it supports management functions, includes a small satellite controller known as the MMC (Module Management Controller). The MMC has at minimum an IPMB slave interface to the BMC.

The design shown includes two SPI Flash devices attached to the BMC, and an eSPI interface to the COM-HPC. The COM-HPC BIOS image can reside in the BMC attached SPI Flash. This allows the BMC to manage Out of Band Flash BIOS updates, if so desired.

This (just above) does not apply to the COM-HPC Mini.

Figure 52:   Carrier BMC with IPMB Link to Module – Sheet 1
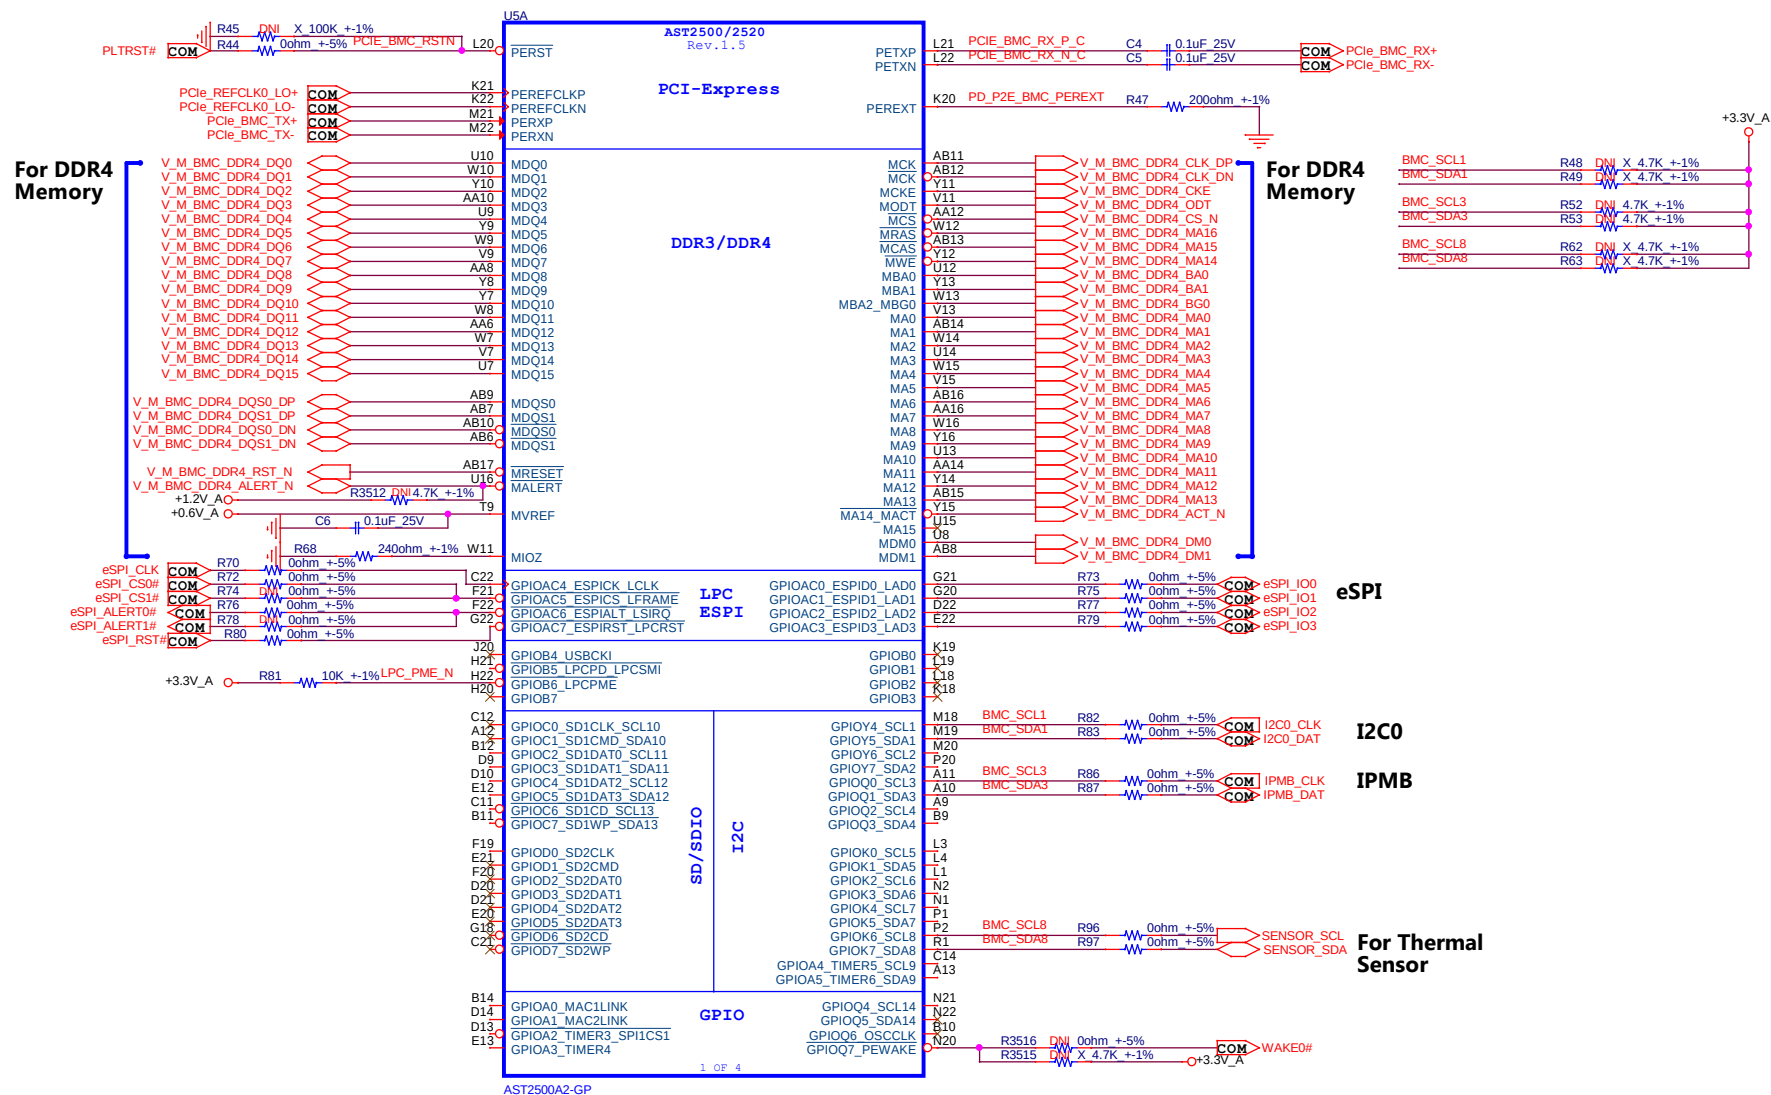
Figure 53:   Carrier BMC with IPMB Link to Module – Sheet 2
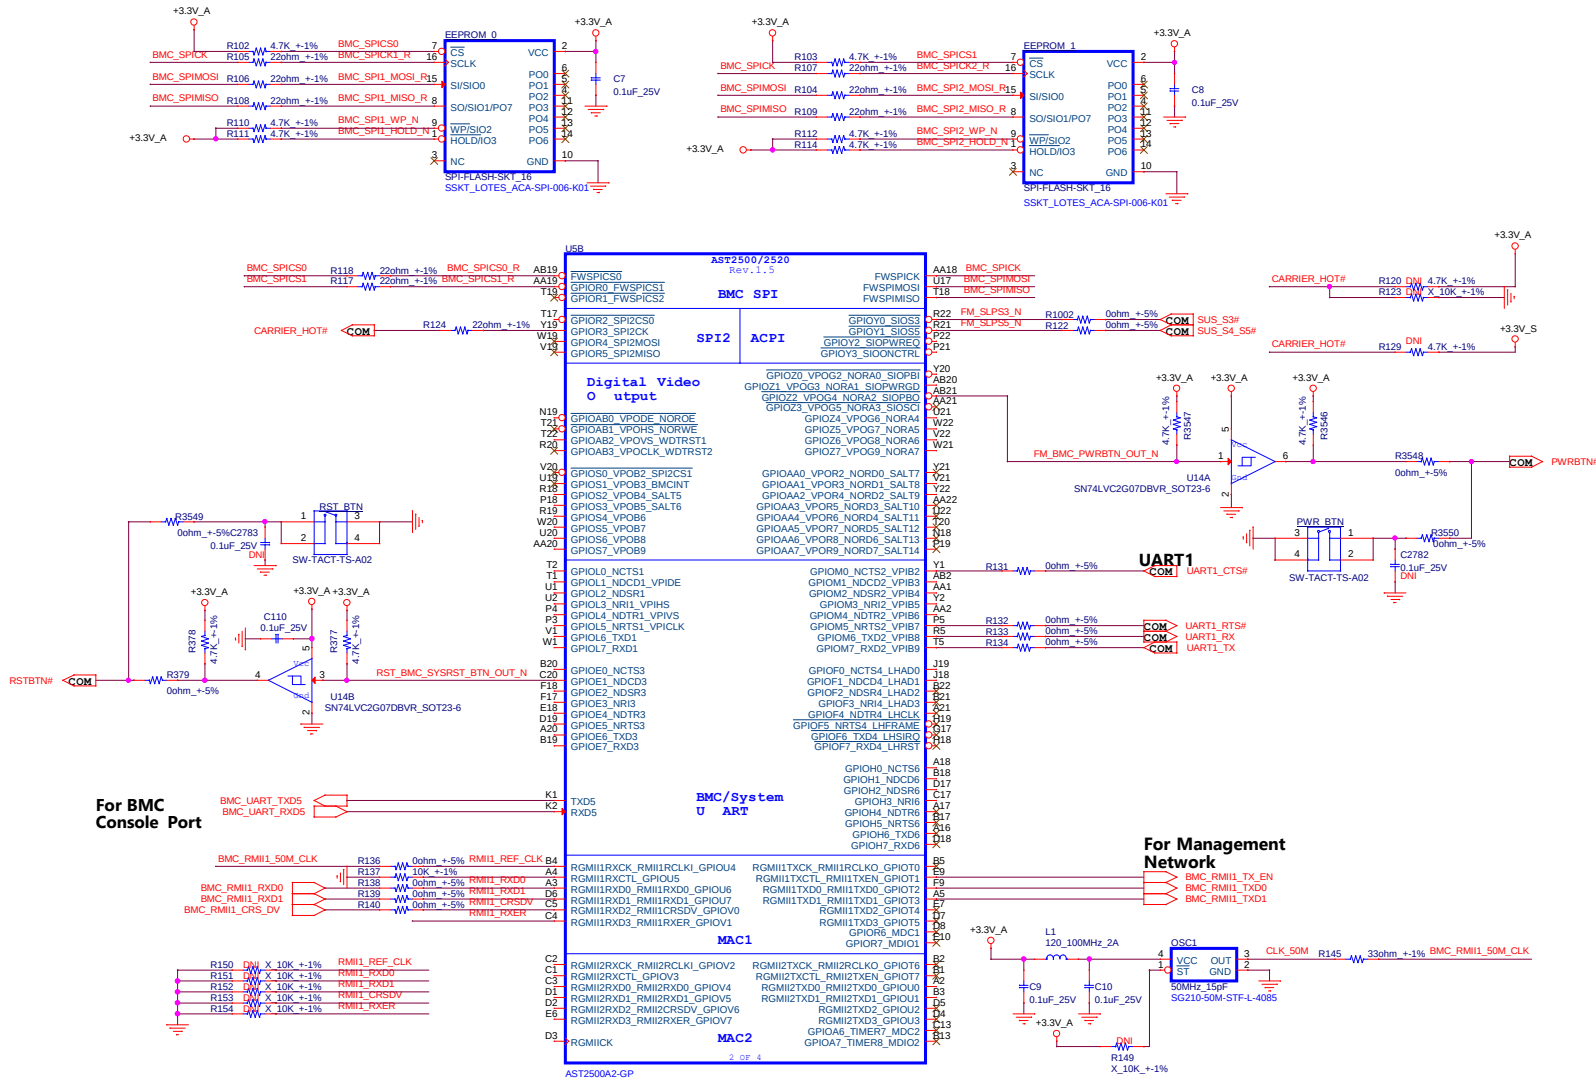
Figure 54: Carrier BMC with IPMB Link to Module – Sheet 3
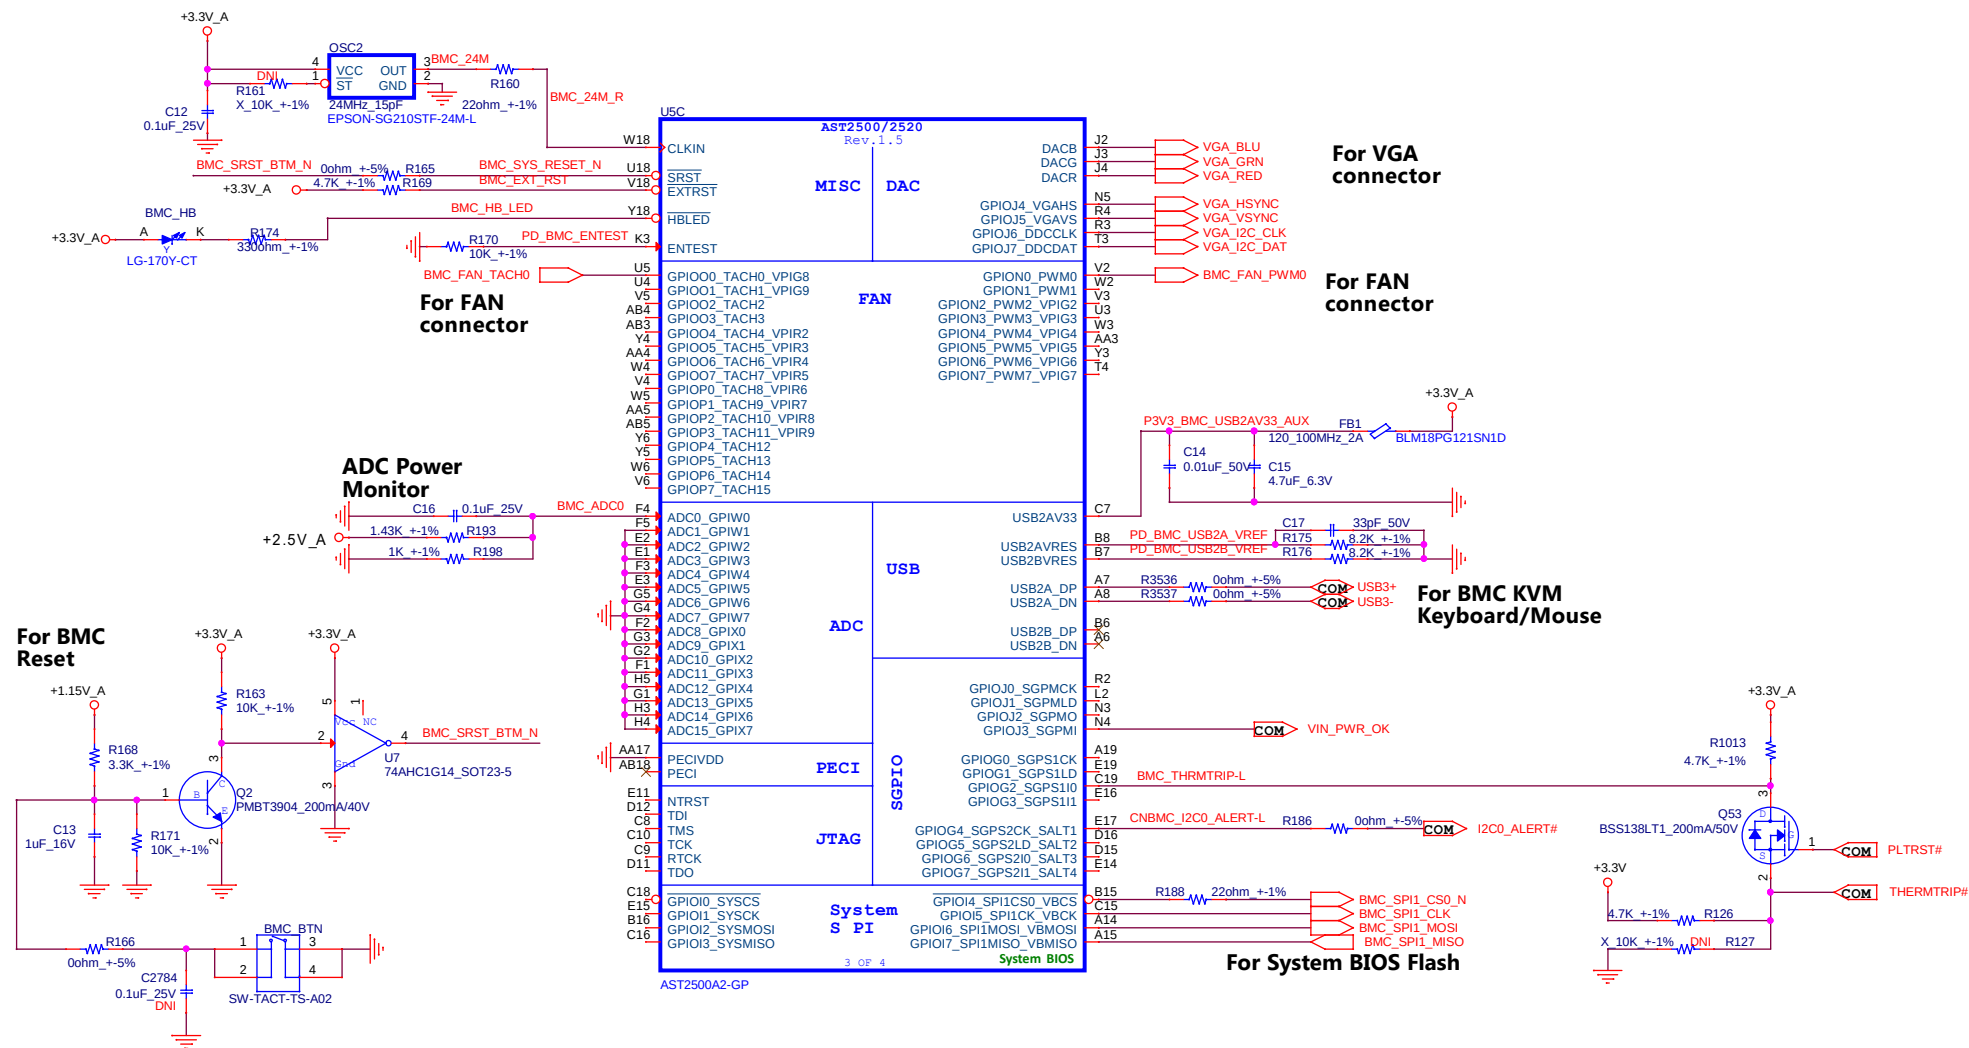
3.19.   General Purpose SPI

The COM-HPC General Purpose SPI port is, from a hardware view, an easy to use interface for Carrier peripherals that requires only four signal pins on the peripheral device (data in, data out, clock and chip select), plus an optional interrupt output. It may be used to implement a wide variety of low to medium speed (circa 4 to 20 MHz signaling rate) peripheral devices such as A/D and D/A converters, touch controllers, CPLDs, FPGAs, Flash memories and many more. The GP SPI interface is significantly faster than traditional I2C ports (400 KHz max for most I2C implementations) but much slower than PCI Express (2.5 GHz signaling and up). A SPI interface is easier to implement within a peripheral device than PCIe, resulting in lower costs.

The COM-HPC pinout definitions allow for four General Purpose SPI chip selects, allowing up to four Carrier GP SPI devices. The COM-HPC GP SPI interface uses 3.3V signal levels, active in the S0 (full on) power state. Carrier GP SPI devices may be daisy-chained or routed in a branch topology with the root at the COM-HPC connector. The data in, data out and clock lines for a particular target device on the Carrier should be loosely kept together and have approximately the same length from the COM-HPC connector to the particular target device. The chip-select lines should be routed directly from the COM-HPC connector to the target device.

Note: the COM-HPC Mini General Purpose SPI signals are 1.8V level signals.

3.20.   Rapid Shutdown

Rapid Shutdown is a rarely used feature but one that is important to some defense industry segment customers. It's purpose is to rapidly collapse all Module and Carrier power rails and remove all bias voltages to prevent damage to the electronics in certain extreme wartime situations. It is purely a hardware feature, without any consideration for an orderly software shutdown.

It is expected that some COM-HPC Module designs will incorporate Rapid Shutdown capability, but that the feature be depopulated unless needed by certain customers. The feature is both fairly simple in concept but potentially tricky in implementation: if the Module Rapid Shutdown pin is asserted by a 5V logic level signal, all power Module and Carrier rails are collapsed by a N-channel FET and appropriately sized drain resistor on each power rail. The +12V or Wide Range power source to the system must be immediately cut as well, and isolated from any bulk capacitance that might provide hold-up power. This usually is achieved by using hot-swap controller devices to gate the system power input, with the bulk hold-up capacitance located at the input side of the hot-swap controller circuitry. All power rails on the Carrier must be collapsed as well.

Design drawings are not shown here. If your Module vendor supports Rapid Shutdown, then they should be able to provide implementation design support.

3.21.   Thermal Protection

COM-HPC defines two pins related to thermal protection of the system:

- CARRIER_HOT#
    - This is a 3.3V *(1.8V on the Mini)* level S0 power domain input signal, with an on-Module pull-up
    - This signal may be left open, or it may be driven low by Carrier hardware if a system over-temperature situation is detected.
    - Module support for this signal is required, per the **COM-HPC Base Specification.**
    - There is no definition in the COM-HPC Base Specification as to how long CARRIER_HOT# should stay low in a system over-temperature situation
- THERMTRIP#
    - This is a 3.3V *(1.8V on the Mini)* level S0 power domain output
    - If driven low, it indicates that the CPU is in an over-temperature situation
    - There is no definition in the COM-HPC Base Specification as to how long THERMTRIP# should stay low in an over-temperature situation
    - Carriers may leave this signal open, or they may act on it
        - Ideally, a Carrier circuit removes system S0 power if the situation persists and is not a short term glitch, and sets a bit in a non-volatile memory that can be read by Module firmware on the next S0 power up.
        - A Carrier BMC may also track / process this COM-HPC output signal

3.22.   System Management Bus (SMBus)

SMBus Introduction

The SMBus is primarily used to manage system peripherals on the COM-HPC Module and on the Carrier. SMBus devices such as the Serial Presence Detect (SPD) EEPROM(s) for the system RAM, thermal sensors, PCIe devices, clock buffers, Smart Battery, etc. are managed over the SMBus. Designers need to take note of several implementation issues to ensure reliable SMBus interface operation. The SMBus is derived from I2C. However, I2C devices have the potential to lock up the data line while sending information and require a power cycle to clear the fault condition. SMBus devices contain a timeout to monitor for and correct this condition. Designers are urged to use SMBus devices when possible over standard I2C devices. COM-HPC Modules are required to power SMBus devices from the suspend power rail in order to have control in all system power states.

The COM-HPC Module may not function correctly or at all if Carrier SMBus devices interfere with proper Module SMBus device operation.

Note: the SMBus signals on the Mini are all 1.8V level signals.

SMBus Power Domain Isolation

The devices on the Carrier Board using the SMBus are usually powered by the main 3.3V (S0) power rail. To avoid current leakage between the suspend (S5) and the main (S0) power rails, the SMBus devices in the S5 power domain must be separated by a bus switch from S0 domain SMBus devices. FET devices, as shown in Figure 50 above, or I2C / SMBus isolation devices, as shown in Table 20 above may be used to achieve the S5 / S0 power domain isolation.

SMBus Addresses

Since the SMBus is used by the Module and Carrier, care must be taken to ensure that Carrier based devices do not overlap the address space of Module based devices. Typical Module SMBus devices and their binary I2C / SMBus addresses include memory SPD (Serial Presence Detect) addresses 1010 000x, 1010 001x, up to 1010 111x for 8 DIMMs, programmable clock synthesizers (1101 001x), clock buffers (1101 110x), thermal sensors (1001 000x), and management controllers (vendor defined address). The 'x' in the binary addresses is the SMBus / I2C R/W bit. Contact your Module vendor for information on the SMBus addresses used on the Module.

3.23.   General Purpose Inputs / Outputs

COM-HPC defines 12 General Purpose I/O pins. It is expected that the 12 pins can be individually configured as either inputs or outputs, that they be configured as inputs on power up, that they be powered by the Module 3.3V suspend (S5) power rail, that there be a 100K pull-up on the Module, and that the Module GPIO pins be able to generate interrupts to the Module CPU.

As the COM-HPC GPIO may be inputs, outputs or bidirectional signals, there are a variety of ways to use and protect them. If the target I/O devices are in a different power domain from (I.e targets are in the S0 domain), that needs to be taken into account by an appropriate logic buffer or FET arrangement, similar to the S5 – S0 power domain isolation shown in this document for I2C and the SMBus.

If any of the GPIO signals are exposed to the outside world and exposed to human contact and ESD events, then there needs to be appropriate ESD protection, EMI mitigation, and hardening against accidents such as short circuiting or exposure to power rails. The details of the protection implemented depend on the factors such as:

- What level of ESD protection is expected?
- What is the GPIO signal bandwidth?
  - Low bandwidth GPIO signals may be protected with simple measures including:
    - Dual Schottky diodes:
      - 1^{st} diode with anode (A) at GND and cathode (K) at the GPIO signal level
      - 2^{nd} diode with anode at GPIO signal level and cathode at the GPIO VCC level
    - Alternatives to the dual Schottky diodes proposed above may be specialty diodes or diode arrays designed for ESD mitigation such as those shown in the NBASE-T, Ethernet, USB, DP and HDMI sections of this document.
    - A series resistor between the Schottky diode K – A node and the COM-HPC GPIO pinout
    - Possibly a ferrite in between the COM-HPC GPIO pin and the external connector
  - If the GPIO signal bandwidth is somewhat higher, then adjustments have to be made:
    - The ESD diode pin capacitance needs to be lower, and appropriately sized for the bandwidth at hand
    - The series resistor value needs to be lowered
    - The ferrite inductance value may need adjustment
- If user abuse is expected (hot plugging, sudden removal etc) then protection measures may include:
  - Some or all of the protection measures listed just above
  - Robust buffer ICs that stand between the COM-HPC pins and the protection measures
    - If the GPIO is to be used as a single direction input or output, then robust buffering is easy
    - Many bidirectional buffers are available, including buffers that auto-sense the signal direction

Note: the GPIO signals on the Mini are all 1.8V level signals.

3.24. Module Type Detection and Protection

There are three TYPE pins defined in the COM-HPC pinouts allowing up to eight Types to be defined. At present, four Types are defined, per Table 21 below:

Table 21: COM-HPC Type Definitions

	Module Connections			Meaning
Ref	TYPE2	TYPE1	TYPE0	
7	NC	NC	NC	Mini Module – Wide Range 8V to 20V input
6	NC	NC	GND	Reserved
5	NC	GND	NC	Reserved
4	NC	GND	GND	Server Module – Fixed 12V input
3	GND	NC	NC	Reserved
2	GND	NC	GND	Reserved
1	GND	GND	NC	Client Module – Wide Range 8V to 20V input
0	GND	GND	GND	Client Module – Fixed 12V input

COM-HPC Carrier hardware may optionally implement hardware to hold off the application of power to the main Carrier circuits and to the Module if the Module and Carrier Types do not match up. The COM-HPC Client and Server pinouts are different (the differences are noted in Table 2 earlier in this document) and it is not desirable to power up a system in which the Carrier and Module Types do not match.

The Carrier hardware shown in Figure 55 below holds off power distribution if the Module Type is not a fixed input voltage Client or a wide-range input Client. The example uses and ATX style power supply. The 5V Standby power to the Carrier and Module is cut off by power switch U60 in the Figure 55 below. The open drain FET T1 along with pull-down resistor R539 ensure that the ATX power control line (ATX_PSON#) is floating and not pulled low. This prevents the main ATX power rails from coming on. FET T1 should not be replaced by a logic gate as the gate's internal ESD protection diodes might provide a path for the ATX_PSON# signal to be pulled low unintentionally.

Figure 55: Module Type Detection / Protection – ATX Power Supply and Client Type Module / Carrier
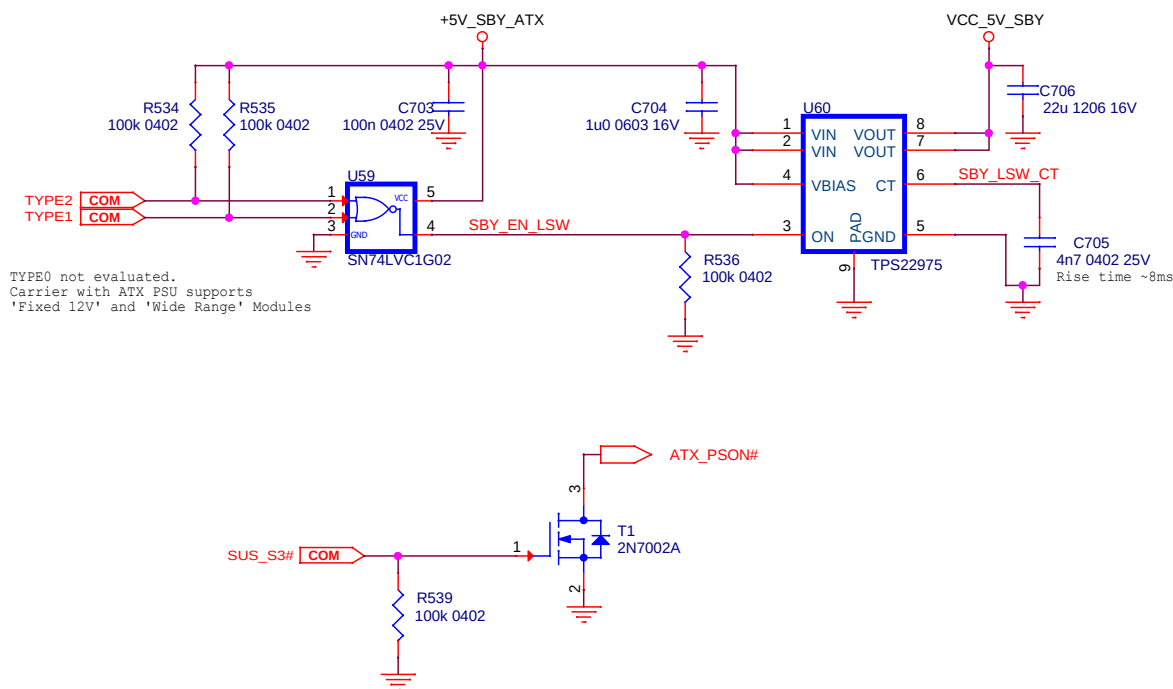
The Carrier hardware shown in Figure 56 below holds off power distribution if the Module Type is not a Server Type Module. The example uses and AT style supply.

Figure 56: Module Type Detection / Protection – AT Power Supply and Server Type Module / Carrier
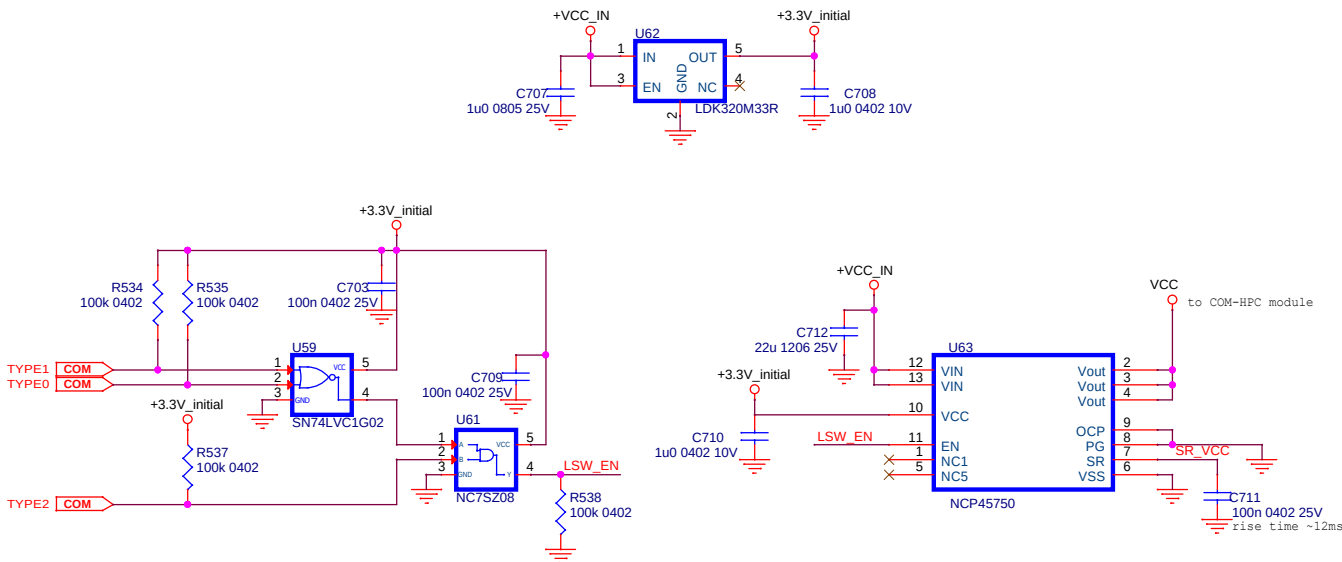
Server Modules may use either AT or ATX or other style power supplies. An AT supply is used here just as an example.

The circuit example below prevents the application of power to the COM-HPC Mini if the Module and Carrier Type pins do not match up.

Figure 57: Module Type Detection / Protection – COM-HPC Mini Module
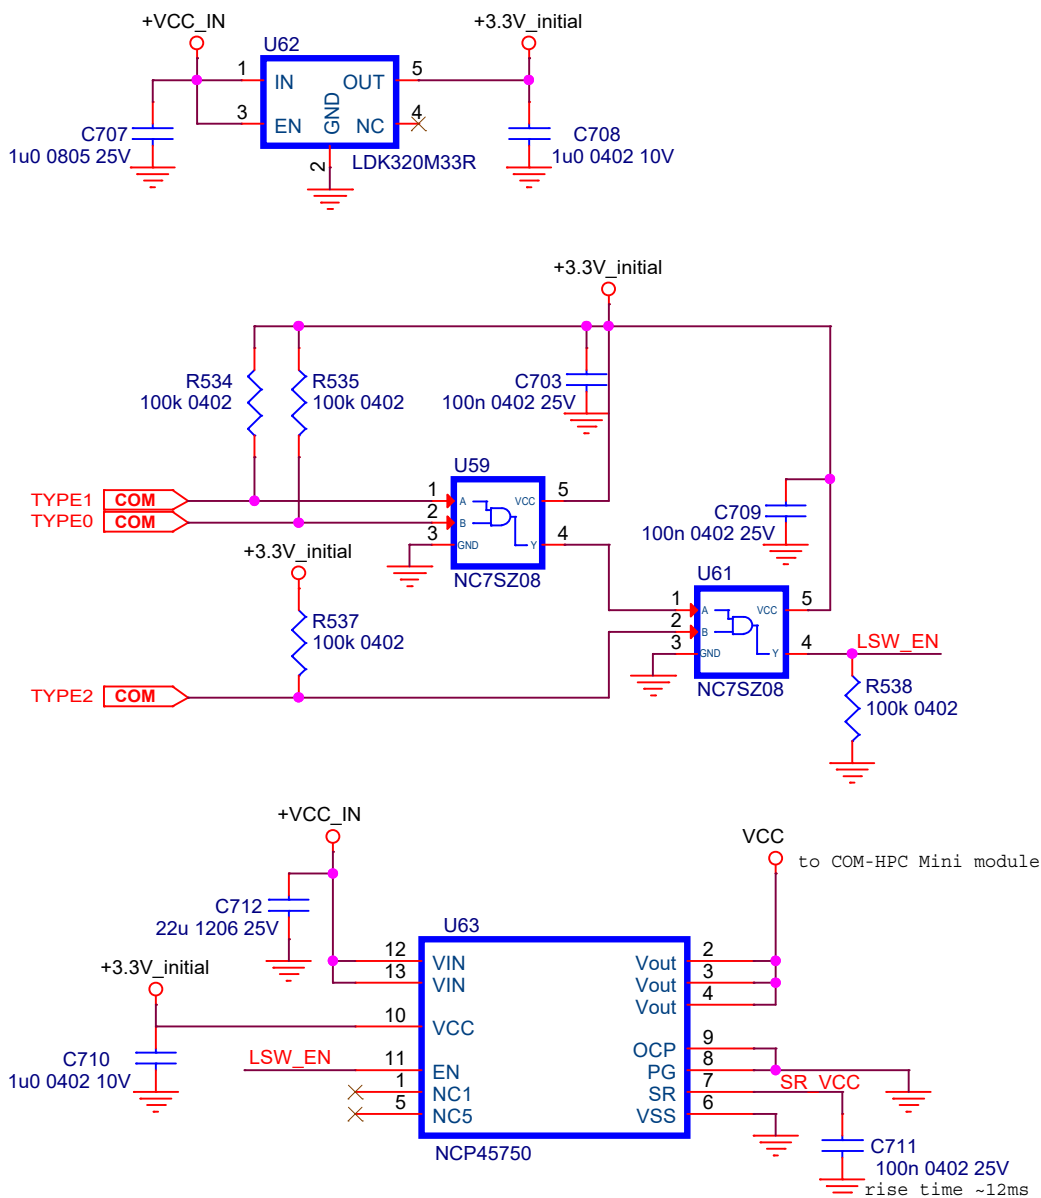
3.25.   COM-HPC Mini SGMII Ethernet Interface Example

Revision 1.20 of the COM-HPC Base Specification document allows for two SGMII port options. They are pin shared with COM-HPC Mini PCIe lanes 3 and 2. See the Base Specification document Section 10 for details. A circuit example for a 10/100/1000 TSN aware Ethernet implementation is shown here. Note that the SGMII TX and RX pairs require series coupling capacitors on the Carrier; they are not present on the COM-HPC Mini Module. The Mini also has an MDIO pin pair (essentially an I2C bus) available to set up the PHY.

Figure 58:   COM-HPC Mini SGMII Gbps Ethernet Example
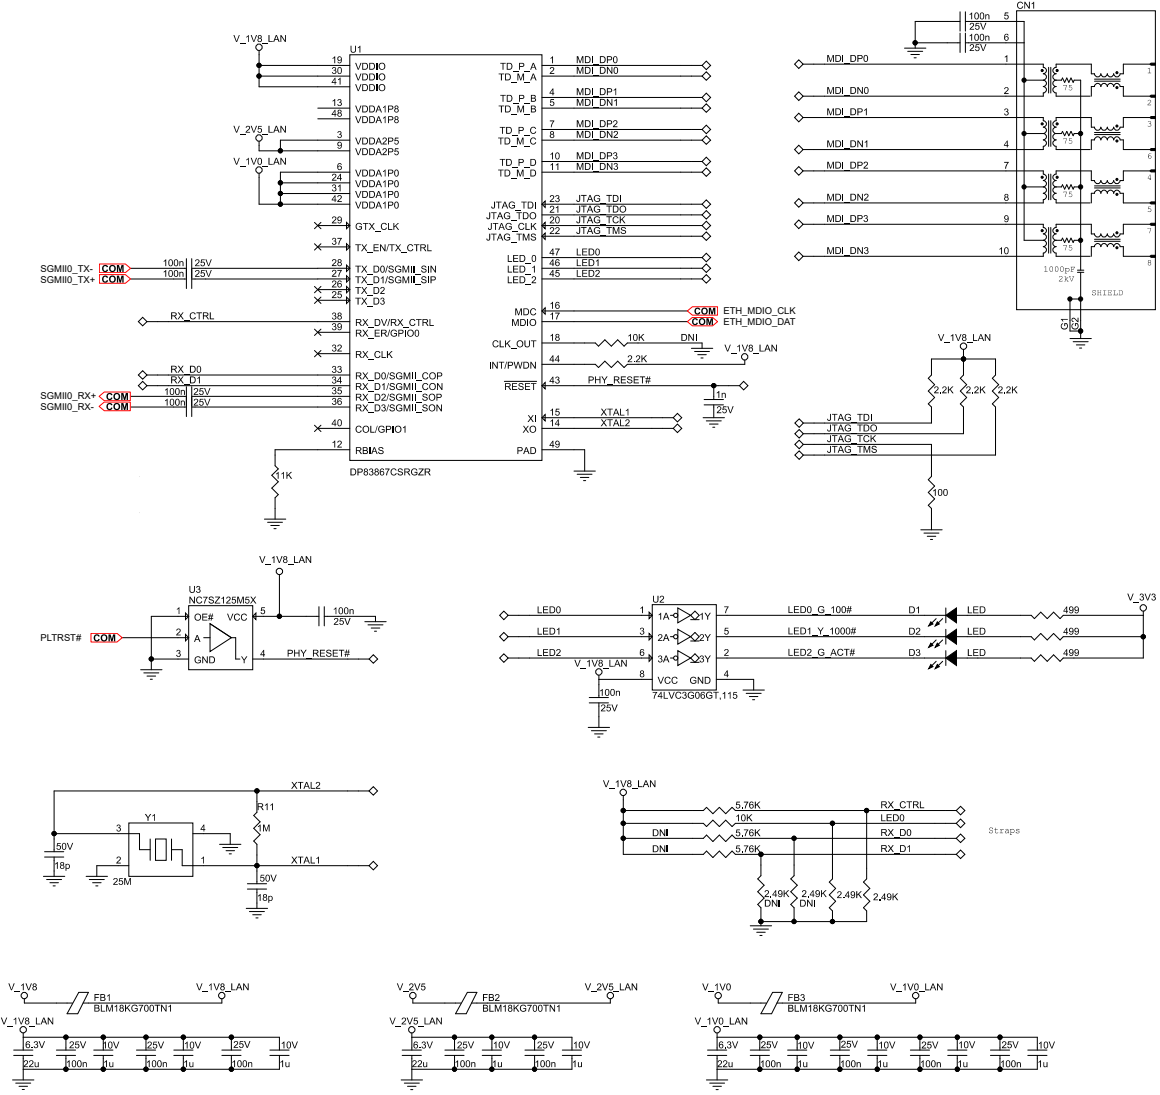
Note: this schematic example is borrowed and adapted from the SGET SMARC Design Guide, with permission.

4. PCB Design Rule Summaries

4.1. High Speed PCB Design Information – Design Guides and Books

4.1.1. Intel and AMD Design Guides

Intel and AMD have a wealth of design guide material available, although most current materials are NDA (Non Disclosure Agreement) protected and Carrier designers must obtain their own NDAs with these vendors to access these documents.

A few useful documents are listed in Table 22 below. Although these guides are centered around CPU board development there is also much general high speed design information, often in graphical format, about topics such as how to keep differential pairs length matched, about stackups, about via stubs, about voiding planes under certain components and features, and so on. There is also information about peripheral interfaces such as PCIe, USB 3 and 4 etc.

Note for the COM-HPC Mini: the high speed rules shown here for USB Super Speed, PCIe, SATA etc apply to the Mini as well the COM-HPC Client and Server.

The differences between the COM-HPC Client, Server and the Mini have mainly to do with the I/O voltage levels of the single ended sideband signals such as the SM buses, I2C buses etc. and the pinout choices.

Table 22: Intel and AMD Design Guides

Vendor	Doc #	Description / Title	Notes
AMD	5515	Socket SP3 Processor Mother Board DG	Some general high speed PCB design info Fiber weave effect information PCIe Gen 3 and 4 information
Intel	576513	Intel Confidential	Some general high speed PCB design info PCB differential pair length matching techniques PCIe Gen 3 and 4 information USB 3.1 Ethernet KR 10G/25G information
Intel	607872	Tiger Lake UP3 UP4 Platform DG	Fiber weave effect information PCIe Gen 3 and 4 length matching information USB4 routing information
Intel	618429	Tiger Lake H Platform DG	
Intel	627205	Intel Confidential	Fiber weave effect information No stub routing techniques Voiding advice PCIe Gen 4 and Gen 5 advice
Intel	406926	Fiberweave Effect White Paper	

4.1.2. Books on High Speed PCB Design Principles

The publications listed below are much more academic than the design guides listed in Section Error: Reference source not found above. These books may be useful to designers interested in the engineering and physics details of what is going on with very fast signal propagation.

The book titled *High Speed Digital Design: Design Of High Speed Interconnects And Signaling* is the newest and perhaps most relevant book in this list. It was written by a trio of Intel engineers and covers contemporary high speed serial interface topics quite thoroughly.

- ***Advanced Signal Integrity For High-Speed Digital Designs***
  Stephen H.; Heck and Howard L Hall
  ISBN 13: 9780470192351
  ISBN 10: 0470192356
  © 2009 Wiley-IEEE Press

- ***High Speed Digital Design: Design Of High Speed Interconnects And Signaling***
  Hanqiao Zhang, Steven Krooswyk and Jeffrey Ou

  © 2015 Morgan Kaufman, Elsevier Inc.
  ISBN: 978-0-12-418663-7

- ***High-Speed Signal Propagation – Advanced Black Magic***
  Howard Johnson and Martin Graham
  © 2003 Pearson Education, Prentice Hall Professional Technical Reference

- ***Right the First Time – A Practical Handbook on High Speed PCB and System Design Volumes 1 and 2***
  Lee W. Ritchey
  ©2006 Speeding Edge

- Signal Integrity Issues and Printed Circuit Board Design
  Douglas Brooks
  ©2003 Pearson Education, Prentice Hall Professional Technical Reference

4.2. High Speed Serial Interfaces – General PCB Design Rules

Figure 59: PCB Cross Section Terms and Notations

Some of the terms and notations in the Figure 59 above are used in the Tables and text on the following pages. The long copper colored thin rectangles represent the GND or PWR planes and the smaller rectangles, for the most part, the edge coupled differential pairs. The upper and lower signal layers within a Dual Stripline structure should be routed orthogonally to each other to minimize coupling and thereby crosstalk.

Table 23: General Design Rules for High Speed Serial interfaces

Ref	Rule / Recommendation
1	High speed pairs should be routed as edge-coupled differential pairs referenced to and closely coupled to an unbroken GND plane.
2	High speed pairs with Nyquist frequencies at 4 GHz or more (PCIe Gen 3,4,5, USB 3.2 Gen2, USB4 Gen 3, DisplayPort, eDP, HDMI, Ethernet KR) require the most PCB routing care.
3	The preferred routing environments for high speed pairs are ranked here, from most desirable to least: 1. Symmetric Stripline routing with clean GND planes above and below gives the best signal integrity, but it is often an impractical luxury. The two GND planes should be periodically tied together with stitching vias, every inch or so in both X and Y. 2. Asymmetric Stripline routing with the differential pair traces close the primary reference plane, an unbroken GND plane, and further from the secondary plane. The secondary plane can be a GND plane (preferred) or a power plane, possibly with plane splits. 3. Asymmetric Dual Stripline routing with high speed pairs close to the primary reference plane, a GND plane. The traces on the "other" routing layer should be as far away as possible and be routed orthogonally to the GND referenced high speed pairs. The "other" traces can be high speed pairs as well, if their primary reference plane is also GND and if the two signal layers are orthogonal. If the routes on the Asymmetric Dual Stripline routing layers are not be truly orthogonal (90 degrees) they should be angled at least 30 degrees relative to each other. 4. Microstrip routing.
4	Use as few vias as possible. What few vias there are should be symmetrically placed, such that the + and – lines in the pair "see" the same obstacles and impedance discontinuities. If there is a reference plane change, there must be a stitching via close to the signal via. If the planes are at the same potential (e.g. both GND), direct (DC coupled) stitching vias are used. If the reference planes are at different potentials (not desirable for high speed pairs) then a stitching capacitor is used near the signal vias. These concepts are illustrated in some of the Design Guides referenced in Table 22 above.
5	The higher speed interfaces may need no-stub vias or very short stub vias. This may mean backdrilling the vias with controlled depth drills to hollow out the unused portion of the via barrel. Alternatively, via structures that are built up or are laser drilled and only transit a limited number of layers (say from Layer 1 to Layer 3, with Layer 2 being a GND plane) may be used. Another strategy to avoid via stubs is to arrange that the vias connect layers on opposite sides of the PCB. Then there is no stub (for outer layer to outer layer) or perhaps a shorter stub. Yet another strategy is to use sequential lamination PCB construction. For example, a 12 layer PCB can be built as two 6 layer PCBs and then laminated together to form a 12 layer PCB, with short vias spanning layers1-6 and layers 7-12 and longer vias spanning layers 1-12.
6	If layer transitions must be done, having the high speed signals in question straddle a common GND plane is beneficial as there is no change in the reference layer. For example, if signal pairs are on Layer 1 and 3 and Layer 2 is GND, then there is no change in GND reference plane for the Layer 1 – 3 transitions. If the layer changes result in a change in GND reference planes, then there need to be GND stitching vias close to the trace vias. The stitching vias tie the GND planes together in the vicinity of the signal pair layer transition.
7	It is critically important that the + and – signal lines in a differential pair are closely length matched. The matching is on the order of a few mils for the faster interfaces. This is sometimes called **intra-pair length matching**. In this document, this is referred to as **differential pair + and – length matching.** The different pairs in a group (e.g. the four TX+ and – pairs in x4 PCIe link) do not need to be matched very closely at all for many modern interfaces. This is sometimes called **inter-pair length matching.** In this document, it is referred to as **pair to pair length matching** (or similar). For some interfaces, this mismatch can be on the order of inches.
8	Coupling capacitors should be discrete 0402 or 0201 package size parts. Do not use capacitor arrays as these can have internal cross coupling that can severely attenuate the differential signal.
9	The plane under the coupling capacitors for the higher speed interfaces should be voided (meaning that rectangular holes about the same size as the capacitor lands, or slightly larger, should be created in the plane) – whether it is a GND or PWR plane. See the Intel Document 627205 referenced in Table 22 above for details.
10	Plane layers that do NOT connect to a particular via should be voided with a circular void around the via barrel. This is done anyway so that the plated via hole does not connect to the plane, but the void should be expanded somewhat to avoid signal coupling to the plane. See the Intel Document 627205 referenced in Table 22 above for details.

Ref	Rule / Recommendation
11	High speed traces should not be run close to the board edges, especially for long runs parallel to the edge. If they are run this way, they may be creating an EMI hazard.
12	Sometimes the differential pairs are serpentined to adjust the pair length to match another pair. There is usually a "minimum distance to self", listed in some of the Tables below.
13	The fastest interfaces need to account for the "Fiberweave Effect". This effect is due to the periodic variations in the PCB material dielectric constant caused by the fiberglass weave pattern within a PCB layer. The mitigation strategies are to either arrange that the PCB routes are not parallel (in x or y) to the glass fibers in the weave, or to use a PCB material that does not show this effect. Some of the references in Table 22 have details and illustrations on this effect.

4.3. PCB Design Rule Summaries – High Speed Differential Pair Serial Interfaces

The COM-HPC high speed serial interfaces and two of the fastest single ended interfaces were extensively simulated by a Signal Integrity subgroup during the development of the COM-HPC specification. These efforts resulted in a set of loss budgets, maximum trace length values and other related recommendations. This is documented in Section 6 of the COM-HPC Base Specification. The loss-budget approach allows the findings to be adapted to various different PCB materials (e.g. Standard Loss, Mid Loss, Low Loss and Very Low Loss).

The Base Specification recommendations as they apply to COM-HPC Carrier designs are summarized in the Sections below . Some recommendations such as trace length matching requirements are not found in the Base Specification document; rather they are compiled from industry sources.

4.3.1. NBASE-T Design Rule Summary

Table 24: NBASE-T Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Signaling Rate / Nyquist Frequency	1000BASE-T: 250 Mbps / ~80 MHz 10GBASE-T: 2.5 Gbps / ~450 MHz
2	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference Microstrip routing may be used Microstrip is necessary near connectors Quiet unbroken well bypassed power plane may be used as a reference plane.
3	Differential Trace Impedance	100 ohm +/- 10%
4	Single Ended Trace Impedance	55 ohm +/- 15%
5	Max Module Trace Length	1GBASE-T STD Loss PCB Material: $\leq$ 3000 mils 10GBASE-T STD Loss PCB Material: $\leq$ 1500 mils 10GBASE-T MID Loss PCB Material: $\leq$ 1500 mils
6	Max Carrier Trace Length	1GBASE-T STD Loss PCB Material: $\leq$ 5000 mils 10GBASE-T STD Loss PCB Material: $\leq$ 2500 mils 10GBASE-T MID Loss PCB Material: $\leq$ 4500 mils
7	Differential Pair +/- Length Matching (Carrier / Module)	5 mil / 5 mil
8	TX Pair to TX Pair Length Matching (Carrier / Module)	500 mil / 500 mil
9	RX Pair to RX Pair Length Matching (Carrier / Module)	500 mil / 500 mil
10	TX Pair to RX Pair Length Matching (Carrier / Module)	500 mil / 500 mil
11	TX Pair to RX Pair Spacing (Carrier / Module)	D $\geq$ 5*H (Asymmetric Stripline)
12	TX or RX pair Spacing to Other Signals	DX $\geq$ 8*H (Asymmetric Stripline)
13	Max Via Stub Length	80 mil

See Figure 59 above for definitions of D, DX and H.

4.3.2. Ethernet KR Design Rule Summary

Table 25: Ethernet KR Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Signaling Rate / Nyquist Frequency	10G KR: 10.3125 Gtps / ~5.1 GHz 25G KR: 25.78125 Gtps / ~12.9 GHz 40G KR4: 10.3125 Gtps / ~5.1 GHz 100G KR4: 25.78125 Gtps / ~12.9 GHz
2	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference
3	Differential Trace Impedance	93 ohm +/- 10%
4	Single Ended Trace Impedance	50 ohm +/- 15%
5	Maximum Trace Lengths on Carrier (adapted from COM-HPC Base Specification V1.0 Tables 79, 81, 83)	

PHY Down on Carrier Carrier Trace Lengths	Budget (dB)	Standard Loss (SL) PCB Material (inches)	Mid Loss (ML) PCB Material (inches)	Low Loss (LL) PCB Material (inches)	Very Low Loss (VLL) PCB Material (inches)
10GBASE-KR	16.0	13.2	21.0	27.4	37.3
25GBASE-KR RS-FEC	20.0	7.9	12.9	16.5	23.1
25GBASE-KR BASE-R FEC	16.0	6.3	10.3	13.2	18.5
25GBASE-KR No FEC	12.0	4.7	7.8	9.9	13.8

Module1 MAC to Module2 MAC Carrier Trace Lengths	Budget (dB)	Standard Loss (SL) PCB Material (inches)	Mid Loss (ML) PCB Material (inches)	Low Loss (LL) PCB Material (inches)	Very Low Loss (VLL) PCB Material (inches)
10GBASE-KR	12.0	9.9	15.7	20.5	28.0
25GBASE-KR RS-FEC	12.0	4.7	7.8	9.9	13.8
25GBASE-KR BASE-R FEC	8.0	3.2	5.2	6.6	9.2
25GBASE-KR No FEC	3.0	1.2	1.9	2.5	3.5

SFP Connector on Carrier Carrier Trace Lengths	Budget (dB)	Standard Loss (SL) PCB Material (inches)	Mid Loss (ML) PCB Material (inches)	Low Loss (LL) PCB Material (inches)	Very Low Loss (VLL) PCB Material (inches)
SFP+ Max Carrier Trace	1.50	1.24	1.97	2.56	3.50
SFP28 Max Carrier Trace	2.00	0.79	1.29	1.65	2.31

Ref	Parameter Description	Parameter Value
6	Differential Pair +/- Length Matching (Carrier / Module) Note the very tight matching for 25GBASE-KR This is actually relaxed from some Intel recommendations, per PICMG consultation with Intel	2.5 mil / 2.5 mil for 10GBASE-KR 1.5 mil / 1.5 mil for 25G BASE-KR
7	TX Pair to Pair Length Matching (Carrier / Module)	500 mil / 500 mil (KR4 only; N/A for KR)
8	RX Pair to Pair Length Matching (Carrier / Module)	500 mil / 500 mil (KR4 only; N/A for KR)
9	TX Pair to RX Pair Length Matching (Carrier / Module)	500 mil / 500 mil
10	TX Pair to RX Pair Spacing (Carrier / Module)	$D \geq 5*H$ (Asymmetric Stripline)
11	TX or RX pair Spacing to Other Signals	$DX \geq 8*H$ (Asymmetric Stripline)
12	Max Via Stub Length	10 mil (KR and KR4)

4.3.3. SATA Design Rule Summary

Table 26: SATA Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Signaling Rate / Nyquist Frequency	Gen 1: 1.5 Gtps / 0.75 GHz Gen 2: 3 Gtps / 1.5 GHz Gen 3: 6 Gtps / 3 GHz
2	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference
3	Differential Trace Impedance	85 ohm +/- 10%
4	Single Ended Trace Impedance	45 ohm +/- 15%
5	Maximum Trace Lengths (from COM-HPC Base Specification V1.0 Tables 58 and 60)	

Device Up on M.2 or mSATA Card	Budget dB	Standard Loss (SL) PCB Material Inches	Mid Loss (ML) PCB Material Inches	Low Loss (LL) PCB Material Inches	Very Low Loss (VLL) PCB Material Inches
SATA Gen1 Carrier Trace	1.1	4.07	6.88	8.46	9.17
SATA Gen2 Carrier Trace	1.8	3.91	6.55	8.14	10.11
SATA Gen3 Carrier Trace	2.9	3.74	6.03	7.80	10.21

Cabled Interface	Budget dB	Standard Loss (SL) PCB Material Inches	Mid Loss (ML) PCB Material Inches	Low Loss (LL) PCB Material Inches	Very Low Loss (VLL) PCB Material Inches
SATA Gen1 Carrier Trace	0.7	2.59	4.38	5.38	5.83
SATA Gen2 Carrier Trace	1.1	2.39	4.00	4.98	6.18
SATA Gen3 Carrier Trace	1.8	2.32	3.74	4.84	6.34

Ref	Parameter Description	Parameter Value
6	Differential Pair +/- Length Matching (Carrier / Module)	2.5 mil / 2.5 mil to support SATA Gen 3
7	TX Pair to RX Pair Length Matching (Carrier / Module)	No requirement
8	TX Pair to RX Pair Spacing (Carrier / Module)	D $\geq$ 5*H (Asymmetric Stripline)
9	TX or RX pair Spacing to Other Signals	DX $\geq$ 8*H (Asymmetric Stripline)
10	Max Via Stub Length	80 mil

4.3.4. PCIe Design Rule Summary

Table 27: PCIe Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Signaling Rate / Nyquist Frequency	Gen 3: 8 Gtps / 4 GHz Gen 4: 16 Gtps / 8 GHz Gen 5: 32 Gtps / 16 GHz
2	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference
3	Differential Trace Impedance (PCIe data and Ref CLK pairs)	85 ohm +/- 10%
4	Single Ended Trace Impedance	45 ohm +/- 15%
5	Maximum Trace Lengths (from COM-HPC Base Specification V1.0 Tables 52 and 54)	

Device Down	Budget (dB)	Standard Loss (SL) PCB Material (inches)	Mid Loss (ML) PCB Material (inches)	Low Loss (LL) PCB Material (inches)	Very Low Loss (VLL) PCB Material (inches)
Gen 3 Max Carrier Trace	11.00	10.37	15.65	20.16	26.68
Gen 4 Max Carrier Trace	12.50	7.11	10.77	13.44	17.64
Gen 5 Max Carrier Trace	13.00	4.23	6.86	8.97	12.65

Device Up	Budget (dB)	Standard Loss (SL) PCB Material (inches)	Mid Loss (ML) PCB Material (inches)	Low Loss (LL) PCB Material (inches)	Very Low Loss (VLL) PCB Material (inches)
Gen 3 Max Carrier Trace	6.50	6.13	9.25	11.91	15.77
Gen 4 Max Carrier Trace	7.50	4.26	6.46	8.06	10.58
Gen 5 Max Carrier Trace	8.00	2.60	4.22	5.52	7.78

Ref	Parameter Description	Parameter Value
6	Differential Pair +/- Length Matching (Carrier / Module)	Gen 3, 4, 5: $\leq$ 2.5 mil / 2.5 mil
7	TX Pair to TX Pair Length Matching (Carrier / Module) RX Pair to RX Pair Length Matching (Carrier / Module)	Gen 3, 4, 5: $\leq$ 500 mil / 500 mil
8	TX Pair to RX Pair Length Matching (Carrier / Module)	No TX to RX matching required
9	TX Pair to RX Pair Spacing	Gen 3, 4, 5: D $\geq$ 5H
10	TX or RX pair Spacing to Other Signals	Gen 3, 4, 5: DX $\geq$ 8H
11	PCIe Data Pair Distance to Self	Gen 3, 4, 5: D $\geq$ 3W
12	PCIe RX or TX Data Pair Length relative to PCIe Reference Clock Pair Length	No matching required. The Reference Clock Pairs should be routed as directly as possible.
13	Maximum Via Stub Lengths:	Gen 3: $\leq$ 80 mil Gen 4: $\leq$ 30 mil Gen 5: $\leq$ 10 mil
14	Land Pattern and Via Voiding Recommendations:	Gen 3, 4, 5: planes adjacent to component lands should be voided. All layers that a PCIe coupling via passes through should be voided, unless the via connects on that layer. See the Intel Document 627205 referenced in Table 22 above for illustrations on voiding and length matching.
15	Fiberweave Effect Mitigation	See Fiberweave Effect references in Table 22 above. Alternatively, use a PCB material that does not exhibit this effect.

4.3.5. USB 2.0 Design Rule Summary

Table 28: USB 2.0 Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Signaling Rate / Nyquist Frequency	480 Mbps / 240 MHz (USB 2.0 High Speed)
2	PCB Routing Environment	Asymmetric Stripline is best Microstrip may be used Unbroken GND plane primary reference is best Quiet PWR plane may be used as a reference Plane splits should be avoided If plane splits are unavoidable, stitching capacitors should be used to tie the plane regions together, for AC signals
3	Differential Trace Impedance (USB 2.0 data pair)	90 ohm +/- 10%
4	Single Ended Trace Impedance	Circa 45 to 50 ohm
5	Max Carrier Trace Length	Cabled Interface: 14 inches Device Down on Carrier: 28 inches
6	Differential Pair +/- Length Matching (Carrier / Module)	20 mil / 20 mil
7	Pair Spacing to other USB 2.0 Pairs (Carrier / Module)	$D \geq 5*H$ (Asymmetric Stripline)
8	TX or RX pair Spacing to Other Signals	$DX \geq 8*H$ (Asymmetric Stripline)
9	Max Via Stub Length	80 mil

4.3.6. USB 3.2 and USB4 Design Rule Summaries

Table 29: USB 3.2 and USB4 Design Rule Summaries

Ref	Parameter Description	Parameter Value
1	Signaling Rate / Nyquist Frequency	USB 3.2 Gen 1: 5 Gtps / 2.5 GHz USB 3.2 Gen 2: 10 Gtps / 5 GHz USB4 Gen 2: 10 Gtps / 5 GHz USB4 Gen 3: 20 Gtps / 10 GHz
2	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference
3	Differential Trace Impedance (USB SuperSpeed Pairs)	USB 3.2 Gen 1: Historically was 90 ohm Going forward, 85 ohm is OK USB 3.2 Gen 2: 85 ohm +/- 10% USB4 Gen 2: 85 ohm +/- 10% USB4 Gen 3: 85 ohm +/- 10%
4	Single Ended Trace Impedance	45 ohm +/- 15%
5	Maximum **Carrier** Trace Lengths (from COM-HPC Base Specification V1.0 Tables 67 and 69)	

USB SuperSpeed Device Down	Budget (dB)	Standard Loss (SL) PCB Material (Inches)	Mid Loss (ML) PCB Material (Inches)	Low Loss (LL) PCB Material (Inches)	Very Low Loss (VLL) PCB Material (Inches)
USB 3.2 Gen 1	3.2	4.7	7.6	10	12.8
USB 3.2 Gen2	5.5	4.7	7.6	9.8	13.4
USB4 Gen 2	5.5	4.7	7.6	9.8	13.4
USB4 Gen 3	10.5	5.2	7.6	10.7	15

USB SuperSpeed Cabled Interface	Budget (dB)	Standard Loss (SL) PCB Material (Inches)	Mid Loss (ML) PCB Material (Inches)	Low Loss (LL) PCB Material (Inches)	Very Low Loss (VLL) PCB Material (Inches)
USB 3.2 Gen 1	3.2	4.7	7.6	10	12.8
USB 3.2 Gen 2	5.5	**4.7**	**7.5**	**9.8**	**13.4**
USB4 Gen 2	5.5	**4.7**	**7.5**	**9.8**	**13.4**
USB4 Gen 3	10.5	**NA**	**NA**	**NA**	**NA**

The values marked in red in the Table just above indicate that there is insufficient overall margin for a direct cabled interface with USB 3.2 Gen 2 or USB4 Gen 2 or USB4 Gen 3. Redrivers or retimers placed close to the cable connectors are advised. The Carrier maximum trace lengths in the Device Down Table above may be used for the run between the COM-HPC Module and the redrivers or retimers.

Ref	Parameter Description	Parameter Value
6	Differential Pair +/- Length Matching (Carrier / Module)	≤ 2.5 mil / 2.5 mil
7	TX Pair to TX Pair Length Matching (Carrier / Module) RX Pair to RX Pair Length Matching (Carrier / Module)	≤ 100 mil / 100 mil Applies only to x2 configurations (2 TX pairs and 2 RX pairs)
8	TX Pair to RX Pair Length Matching (Carrier / Module)	≤ 100 mil / 100 mil
9	TX Pair to RX Pair Spacing	D ≥ 5H
10	TX or RX pair Spacing to Other Signals	DX ≥ 8H
11	TX or RX Data Pair Distance to Self	
13	Maximum Via Stub Lengths:	USB 3.2 Gen 1 ≤ 80 mil USB 3.2 Gen 2 ≤ 30 mil USB4 Gen 2 ≤ 30 mil USB4 Gen 3 ≤ 10 mil

Ref	Parameter Description	Parameter Value
14	Land Pattern and Via Voiding Recommendations:	Gen 3, 4, 5: planes adjacent to component lands should be voided. All layers that a coupling via passes through should be voided, unless the via connects on that layer. See the Intel Document 627205 referenced in Table 22 above for illustrations on voiding and length matching.
15	Fiberweave Effect Mitigation	See Fiberweave Effect references in Table 22 above. Alternatively, use a PCB material that does not exhibit this effect.

4.3.7. DisplayPort Design Rule Summary

Table 30: DisplayPort Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference Microstrip is necessary in region near DP connector
2	Differential Trace Impedance	85 ohm +/- 10%
3	Single Ended Trace Impedance	45 ohm +/- 15%
4	Maximum Carrier Trace Lengths (adapted from COM-HPC Base Specification V1.0 Tables 86 and 87) *(see tables below)* Red text in upper table above indicates that there is insufficient margin in the overall channel and that particular configuration should not be used.	
5	Differential Pair +/- Length Matching (Carrier / Module)	≤ 2.5 mil / 2.5 mil
6	Data Pair to Pair Length Matching (Carrier / Module)	≤ 100 mil / 100 mil
7	Pair to Pair Spacing	D ≥ 5H
8	TX or RX pair Spacing to Other Signals	DX ≥ 8H
9	TX or RX Data Pair Distance to Self	
10	Maximum Via Stub Lengths:	80 mil (per lane bit rate ≤ 5.4 Gtps) 30 mil (per lane bit rate ≤ 13.5 Gtps) 10 mil (per lane bit rate 20 Gtps)
11	Land Pattern and Via Voiding Recommendations:	For DP modes with bit-rate at or above 8 Gtps per lane: Planes adjacent to component lands should be voided. All layers that a coupling via passes through should be voided, unless the via connects on that layer. See the Intel Document 627205 referenced in Table 22 above for illustrations on voiding and length matching.

Tables for Ref 4:

DisplayPort Cabled Interfaces No Carrier Redriver	Bit Rate / Nyquist (per Lane) Gbps / GHz	Standard Loss PCB Material inches	Mid Loss PCB Material inches	Low Loss PCB Material inches	Very Low Loss PCB Material inches
DP HBR	2.7 / 1.3	3.2	5.4	6.7	8.3
DP HBR2	5.4 / 2.7	2.9	4.6	6.0	7.8
DP HBR3	8.1 / 4.0	**2.3**	3.6	4.7	6.3
DP UHBR10	10 / 5.0	**1.65**	**2.6**	3.4	4.6
DP UHBR13.5	13.5 / 6.7	**1.8**	3.0	3.8	5.3
DP UHBR20	20 / 10	**1.3**	**2.2**	**2.8**	3.8

DisplayPort With Carrier Retimer/Redriver	Bit Rate / Nyquist (per Lane) Gbps / GHz	Standard Loss PCB Material inches	Mid Loss PCB Material inches	Low Loss PCB Material inches	Very Low Loss PCB Material inches
DP2.0 UHBR13.5	13.5 / 6.7	5	8	10.4	14.4
DP2.0 UHBR20	20 / 10	3.7	6	7.6	10.7

4.3.8.   eDP Design Rule Summary

Embedded DisplayPort signal integrity considerations were not explicitly addressed by the COM-HPC Signal Integrity subgroup. As such, it would be reasonable for COM-HPC Carrier designers to use the COM-HPC DisplayPort Design Rule Summary outlined in Section 4.3.7. for eDP layouts. For eDP panels, only the lower bit rate formats (HBR, HBR2, HBR3) are likely to come into play.

Alternatively, Carrier Designers can consult some of the Intel and AMD Design Guides listed in Table 22 above for eDP guidance. The Intel Document 627205 in particular has lots of eDP advice. However, these Design Guides are targeting laptop and motherboard designs and it can be tricky to map these recommendations to the COM-HPC system case. The general rule of thumb is that about half of the motherboard or laptop board budget goes to the COM-HPC Module and half to the COM-HPC Carrier.

4.3.9. HDMI Design Rule Summary

Table 31: HDMI Design Rule Summary

Ref	Parameter Description	Parameter Value
1	Preferred PCB Routing Environment	Asymmetric Stripline Unbroken GND plane primary reference Microstrip may necessary in the region near the HDMI connector
2	Differential Trace Impedance	85 ohm +/- 10% (on Carrier before HDMI buffer) 100 ohm +/- 10% (after buffer)
3	Single Ended Trace Impedance	45 ohm +/- 15% (before buffer) 55 ohm +/- 15% (after buffer)
4	Maximum Carrier Trace Lengths (adapted from COM-HPC Base Specification V1.0 Table 88) <table><tr><th>HDMI Buffer / Driver on Carrier near HDMI Connector</th><th>Bit Rate / Nyquist (per Lane) Gbps / GHz</th><th>Standard Loss PCB Material inches</th><th>Mid Loss PCB Material inches</th><th>Low Loss PCB Material inches</th><th>Very Low Loss PCB Material inches</th></tr><tr><td>HDMI 1.4</td><td>3 / 1.5</td><td></td><td></td><td></td><td></td></tr><tr><td>HDMI 2.1</td><td>6 / 3</td><td></td><td></td><td></td><td></td></tr><tr><td>HDMI 2.1</td><td>12 / 6</td><td>4</td><td>5.75</td><td>6.75</td><td>10</td></tr></table>	
5	Differential Pair +/- Length Matching (Carrier / Module)	$\leq$ 2.5 mil / 2.5 mil
6	Data Pair to Pair Length Matching (Carrier / Module)	$\leq$ 100 mil / 100 mil
7	Pair to Pair Spacing	D $\geq$ 5H
8	TX or RX pair Spacing to Other Signals	DX $\geq$ 8H
9	TX or RX Data Pair Distance to Self	
10	Maximum Via Stub Lengths:	80 mil (per lane bit rate $\leq$ 6 Gtps) 30 mil (per lane bit rate = 12 Gtps)
11	Land Pattern and Via Voiding Recommendations:	For HDMI modes with bit-rate at 12 Gtps per lane: Planes adjacent to component lands should be voided. All layers that a coupling via passes through should be voided, unless the via connects on that layer. See the Intel Document 627205 referenced in Table 22 above for illustrations on voiding and length matching.

4.4. PCB Design Rules for Single Ended (SE) Interfaces

Table 32: Design Rules for Single Ended Interfaces

Ref	Rule / Recommendation
1	Most COM-HPC SE traces may be routed using a 55 ohm +/- 15 %. trace impedance. The BOOT_SPI_xxx, eSPI_xxx and GP_SPI_xxx nets are the exceptions and should be routed as 50 ohm +/- 15%
2	SE nets may be routed as Stripline or Microstrip traces, referenced to a GND plane or to a quiet PWR plane.
3	Crossing plane splits should be avoided for the faster SE interfaces (BOOT_SPI_xxx, eSPI_xxx and GP_SPI_xxx). If these nets do cross a split in the reference plane, then the split should be "stitched" with a small capacitor that bridges the split for AC signals.
4	SE signals with higher bit rates and faster edge rates need more routing care than slower signals. The higher bit rate SE signals include: • BOOT_SPI_xxx Up to circa 100 Mhz in some cases; up to circa 50 MHz is more typical • eSPI_xxx Up to circa 50 MHz • GP_SPI_xxx Up to circa 50 MHz • I3C Up to circa 33 MHz • SoundWire Up to circa 12 MHz • Various I2C signals Up to circa 1 MHz or 400 kHz in some cases but more typically are 100 kHz max • UART_xxx Up to circa 1 MHz in some cases – usually less – 115 kHz max is more common COM-HPC SE signals not listed just above are likely to be very slow, almost static in many cases. "More routing care" can mean: • Signal should be GND referenced • No plane split crossings • Stripline routing preferred, with primary reference to GND • Series damping resistors for the signals listed as 50 MHz or more • BOOT_SPI_xxx, eSPI_xxx and GP_SPI_xxx have specific routing rules (see below)
5	The COM-HPC BOOT_SPI_xxx signals are arranged in a "balanced tree" topology. Full details can be found in the COM-HPC Base Specification Version 1.0 Section 6.11.1. Up to 4 BOOT_SPI_xxx devices are allowed, but 3 are on the Module and only 1 (or 0) are allowed on the Carrier. The trace lengths for the BOOT_SPI Data and Clock between the COM-HPC connector balls and the Carrier device must be at least 2000 mils long and no more than 3000 mils long. This is to "balance" the on-Module and off-Module branches of the tree. The Data and Clock lines for this branch of the tree should be length matched to within 250 mil. A series damping resistor is recommended. Refer to the COM-HPC Base specification for more details and a diagram. The Chip Select line associated with the Carrier BOOT_SPI_xxx signals does not need length matching and should be routed as directly as possible.
6	The COM-HPC eSPI_xxx signals are arranged in a "balanced tree" topology. Full details can be found in the COM-HPC Base Specification Version 1.0 Section 6.11.2. Up to 4 eSPI_xxx devices are allowed: up to 2 on the Module and up to 2 on the Carrier. The trace lengths for the eSPI Data and Clock lines between the COM-HPC connector balls and the Carrier device(s) must be at least 2000 mils long and no more than 3000 mils long. This is to "balance" the on-Module and off-Module branches of the tree. A series damping resistor is recommended. There should be separate branches in the tree if there are 2 Carrier devices. Refer to the COM-HPC Base specification for more details and a diagram. The Chip Select line associated with the Carrier eSPI_SPI_xxx signals does not need length matching and should be routed as directly as possible.
7	GP_SPI_xxx net routing should follow the same rules as the eSPI_xxx nets. If there are 2 GP_SPI devices, there should be 2 separate tree branches.
8	If any SE signals leave the Carrier and are exposed to the outside world and to potential contact with users, there should be both EMI and ESD mitigation measures implemented close to the connectors that face outside.

5. Mechanical Considerations

5.1. Heat Spreader / Module / Carrier Attachment Details

5.1.1. Heat Spreader to Module Attachment Notes

The COM-HPC Base Specification calls out Module PCB mounting holes that are to align with corresponding Heat Spreader Plate, Carrier board and possibly system chassis mounting holes or features to hold the entire assembly together.

However, the COM-HPC Base Specification also recommends that there be a separate set of vendor-specific holes to secure the Heat Spreader Plate (HSP), the Thermal Interface Materials (TIM) and the COM-HPC Module board together as a subsystem that can be shipped as a unit, independent of the larger system that includes the Carrier and other components (chassis, heat sinks, etc.). This is desirable as the TIM stack can be a sensitive, precision assembly that is best handled once and only once by the Module vendor.

The reference to separate, design specific holes in the Module and HSP for this purpose are in the **COM-HPC Base Specification** V1.0 in Section 7.5.4 Table 93 Ref 5, reproduced here:

The implementation specific holes / spacers / standoffs used to secure the HSP to the Module should be different from those used at the COM-HPC defined mounting hole sites.

The x-y positions, the number of the vendor-specific HSP / TIM / Module attachment points and other implementation details are not defined by the COM-HPC specification document. However, a typical vertical cross section diagram of how this can be implemented is shown in Figure 60 below.

Figure 60:   Vendor Specific Heat Spreader to Module Attachment – Bottom Side Module PCB Access

Legend:

Heat Spreader Plate – Vendor Specific Implementation Details

HSP to Module Spacers or Standoffs – Vendor Specific Locations and Implementation Details

TIM – Vendor Specific Implementation Details – Typically, Compliant Foam or Phase Change Material

CPU / SOC Die or Lid

Module PCB

Spacers / Standoffs at COM-HPC Defined X-Y Positions, For Module Mounting – Vendor Specific Implementation Details

Module to Carrier Connectors

5.1.2. Heat Spreader / Module Assembly Attachment to Carrier and Chassis

Figures 61 through 64 illustrate a variety of hardware mechanical component and assembly options to secure the COM-HPC HSP, Module, Carrier and system chassis together,

Figure 61: Heat Spreader Assembly Mounting to Carrier – Bottom Side Screw Access
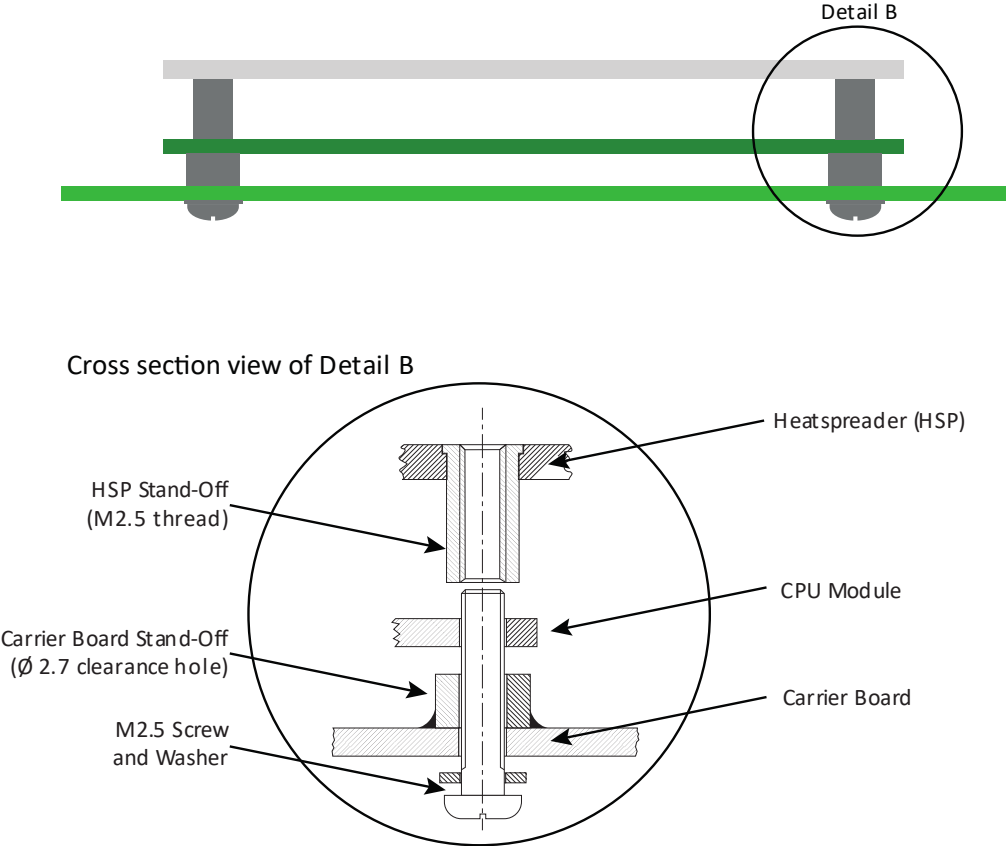
Figure 62:   Heat Spreader Assembly Mounting to Carrier – Top Side Screw Access

Detail A

Cross section view of Detail A

M2.5 Screw
(flat head)

Heatspreader (HSP)

HSP Stand-Off
(Ø 2.7 clearance hole)

CPU Module

Baseboard Stand-Off
(M2.5 thread)

Carrier Board

***Figure 63: Heat Spreader Assembly Mounting to Carrier With Broaching Nut – Top Side Screw Access***

**Figure 64: Heat Spreader Assembly Mounting to Carrier and Chassis – Top Side Screw Access**

Some useful vendors and vendor part numbers for mechanical parts that may be used in Figures 61 through 64 above are listed here:

- PEM TSOS-M25-1500
  - M2.5 threaded blind standoff for sheet metal / plate use – 15 mm overall length (for Client)
- PEM TSOS-M25-1800 (18 mm for Server)
  - M2.5 threaded blind standoff for sheet metal / plate use – 18 mm overall length (for Server)
- www.pemnet.com

- Wurth 9774050951 5 mm Length x 5.1 mm OD x 2.7 mm ID SS SMT Clearance Hole Spacer
- Wurth 9774100951 10 mm Length x 5.1 mm OD x 2.7 mm ID SS SMT Clearance Hole Spacer
  - May be SMT soldered to Carrier Top side as shown in the Figures 61 and 64 above
- www.wuerth.com

- EFCO (Taiwan) has numerous mechanical parts for COM-HPC and other Module standards
- www.efcotec.com
- Or use a search engine, look for "efcotec com accessories"

5.2. Alternative COM-HPC Board Stack Assembly Suggestion

An alternative COM-HPC board stack assembly method and set of mechanical hardware is presented just below. This material has been submitted by Samtec. These assembly mechanics make use of Samtec defined connector hardware components, known as JSOM, for "Jack Screw Stand-off – Micro". These mechanical hardware parts are used in PC-104 and in some VITA assemblies. Samtec JSOM data sheets and drawings are readily available online.

This approach defines an assembly stack allowing a COM-HPC Module and Carrier to be mounted to a metal chassis which is below the Carrier. This assembly method does not include considerations for a Heat Spreader Plate. Thermal management components such as heat sinks or HSP / heat sink combinations would be handled on separate holes.

The 'ASP' references in some of the Figures below are Samtec designations for "Application Specific Parts". There is an ASP summary in Figure 71 several pages below.

5.2.1. Precision Jack Screw Standoffs

Precision jack screw standoff hardware (referred to as JSOM by Samtec) can be used to help mating and unmating procedures in high-normal-force, multi-connector applications. They work like traditional stand-offs but contain an internal machined hex screw that can be turned in a counterclockwise direction to lift the Module Card from the Carrier Board. JSOM based assemblies can mitigate damage to the connector pins, components, boards, and solder joints.

Assembly / Dis-assembly Procedure Overview

Before mating the Module Card to the Carrier Board, use a 1.5mm hex driver to turn the JSOM screw clockwise until the screw is fully seated in the JSOM standoff.

Figure 65: JSOM (Jack Screw Standoff – Micro) Diagram and Application Cutaway

Once all four JSOM screws are fully seated, apply even downward pressure over the J1 and J2 connector regions to mate the Module Card to the Carrier Board. Once the Module Card is fully mated secure the Module Card to the Carrier Board with four hex nuts and lock washers as shown in Figure 66 (a).

Use a torque wrench to tighten the hex nuts to *3.0 (+/- 0.5) in-lbs*. Tighten the nuts in an alternating diagonal pattern shown in Figure 66 (b). For detailed mating recommendations, refer to section 7.5.5 of the **COM-HPC® Module Base Specification, Revision 1.0.**

Figure 66:   (a) Hex Nuts to Torque (b) Diagonal Torque Application / De-application (c) Hex Screw Turns

Unmating the Module Card from the Carrier Board

To unmate the boards remove the locking nuts and washers. Using the diagonal pattern shown in Figure 66 (b) insert the 1.5mm hex key shown in Figure 66 (c) into the JSOM screw labeled 1 and turn counterclockwise a ¼ turn. Repeat this procedure for all JSOM screws labeled 2, 3, and 4 until the connectors unmate. The Module Card can then be removed from the Carrier Board.

Figure 67:   COM-HPC Stack Dis-assembly Procedure Using JSOM Hardware

5.3.   Carrier Board Stiffener

FEM (Finite Element Method) mechanical simulations were conducted to understand the amount of deflection and temporary stress that can occur in the Carrier Board as it is being mated with a Module Card. The simulations assumed that the Carrier Board was fabricated using standard 0.0625" thick FR4 material and fixed to a stiff chassis using metal stand-offs attached to the mounting holes adjacent to both the Carrier P1 and P2 connectors. As shown in Figure 68 below, a downward force was applied evenly over the length of the connector, and the amount of deflection was measured. The results confirmed that 0.0625" Carrier Boards should be supported using some type of stiffening mechanism.

Figure 68:   FEM Simulation Results – 0.0625" FR4 Carrier – No Stiffener

Note however that the stiffness of a piece of sheet material such as a PCB is proportional to the **cube** of the sheet thickness. Hence using a thicker PCB may relax or obviate the need for a Carrier stiffener. PCB thickness of 0.079" (2mm), 0.092" and even 0.125" are common. However, be aware that if the Carrier uses through hole parts (typically for I/O connectors) then increasing the PCB thickness too much will result in a soldering problem as the through hole part leads need protrude beyond the PCB a bit for wave soldering.

A metal simple stiffener design is shown in Figure 69 below with the corresponding keep-out regions shown in Figure 70 below. This Figure shows the Carrier PCB Top side. The Carrier stiffener keep-out regions are on the Bottom side of the Carrier board, as indicated by the dashed lines.

When designing a Carrier Board stiffener there are some points to consider.

- The stiffener should provide uniform support directly underneath the Carrier Board connector and span the entire length of the connector region. This should be done for both the P1 and P2 Carrier connectors.

- The stiffener should be securely anchored to the chassis through mechanical mounting hardware or attached to the bottom side of the Carrier Board using an adhesive.

- The stiffener thickness should be as thick as the application allows.

- Care must be taken when using conductive materials such as steel or alloys.

- This stiffener concept will require a keep-out region where peripheral components cannot be placed.

- It may be necessary to exclude via pads from the PCB Bottom side in the in the keep-out region, or to insulate vias from a metallic stiffener. Kapton tape is the usual remedy for this situation. But such a solution may not be appropriate for high – vibration situations. A thicker, compliant foam material may also be considered.

*Figure 69: Mechanical Carrier Stiffener Possibility*

*Figure 70:   Carrier Board Stiffener Keep-Out Region (Seen Through Carrier)*

Dimensions in the two Figures above are in mm.

*Figure 71: Application Specific Part Number (ASP) Reference Guide*

Non-Metallic Stiffener Possibilities

A simple but effective Carrier board stiffener option is to fabricate a simple non-metallic rectangular bar that is positioned between the Carrier PCB Bottom side and the system chassis. The stiffener bar extent shadows the Carrier connector and the adjacent mounting holes as shown by the dashed lines in Figure 70 above. Nylon is a suitable material. Metal press fit inserts at the mounting hole positions may be beneficial.

6.   Appendices

6.1.   Appendix A: Synchronous Ethernet

Synchronous Ethernet, or SyncE, is an ITU-T standard that allows precision timing information to be embedded into Ethernet physical layer. This signal can be correlated to an external high precision master clock. It is important to telecom providers as the telecom infrastructure moves away from TDM based standards such as SONET and to packet based Ethernet implementations.

Introduction to SyncE

- Synchronous Ethernet (SyncE) distributes a frequency signal through Ethernet
    - Defined in ITU-T G.8261, G.8262, G.8262.1, G.8264
- GbE and above always‡ sends symbols (data or idle)
- SyncE recovers received data rate
    - Ethernet requires ±100 ppm clocking
    - Receivers must handle up to 200 ppm clock delta
    - SyncE saves off a fractional rate to drive DPLL

External DPLLs

- External DPLLs can take in multiple clock sources
    - 1PPS and 10 MHz inputs–GPS/GNSS input(s)
    - often 1PPS as well
    - SyncE recovered clocks
    - IEEE 1588/PTP-driven clocks (also often 1PPS)
    - Local oscillator
    - Long-term oscillator (TCXO or OCXO)
- DPLL sets a priority of inputs
- All outputs driven synchronously off selected input(s)
   - TX side of all PHYs and/or SoCs driven from PLL clock

SyncE on PHYs

- Some SoCs support SyncE on internal PHYs
- Some external PHYs support SyncE clock recovery
- Each RX port adapts to meet incoming data rate
    - Each RX port may be different
    - Fractional clock rate from selected port(s) sent to DPLL
- TX side driven from DPLL
    - All TX ports driven at same rate
- Driver support for SyncE with external DPLLs may vary

Implications for Modules / Pinouts

- If Module SoC and Carrier Board PHY both need SyncE, need SyncE info across connectors
- Carrier sends recovered clock(s), 1PPS input(s)
- Module sends TX clock(s), 1PPS output(s)

Figure 72: Synchronous Ethernet Overview
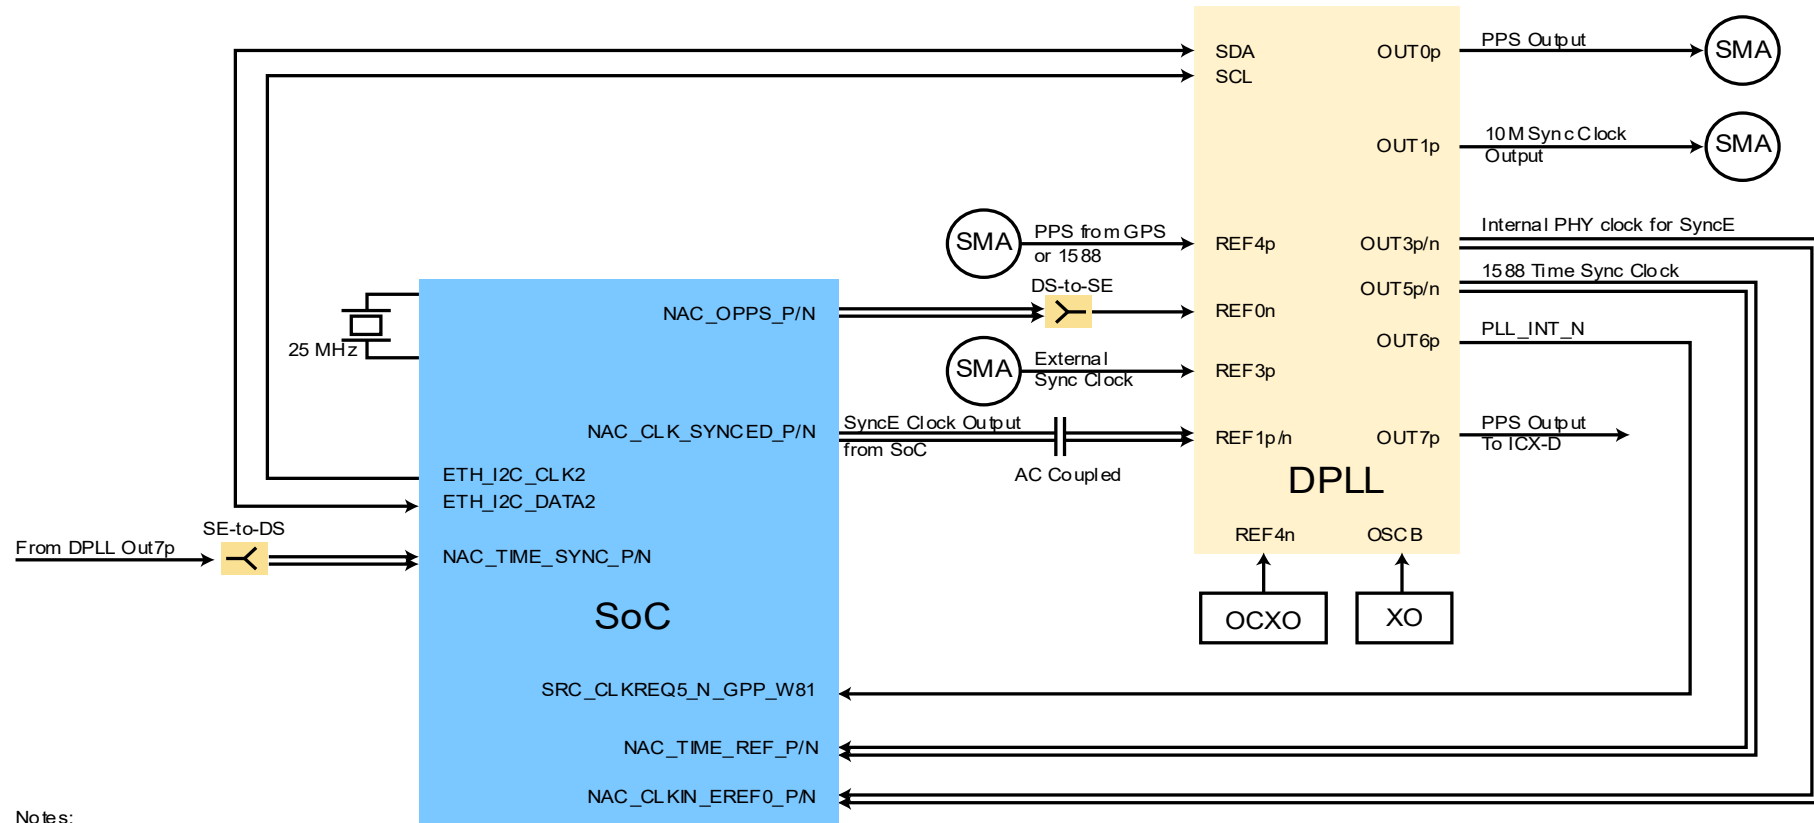
Figure 73: *Synchronous Ethernet Example Implementation*
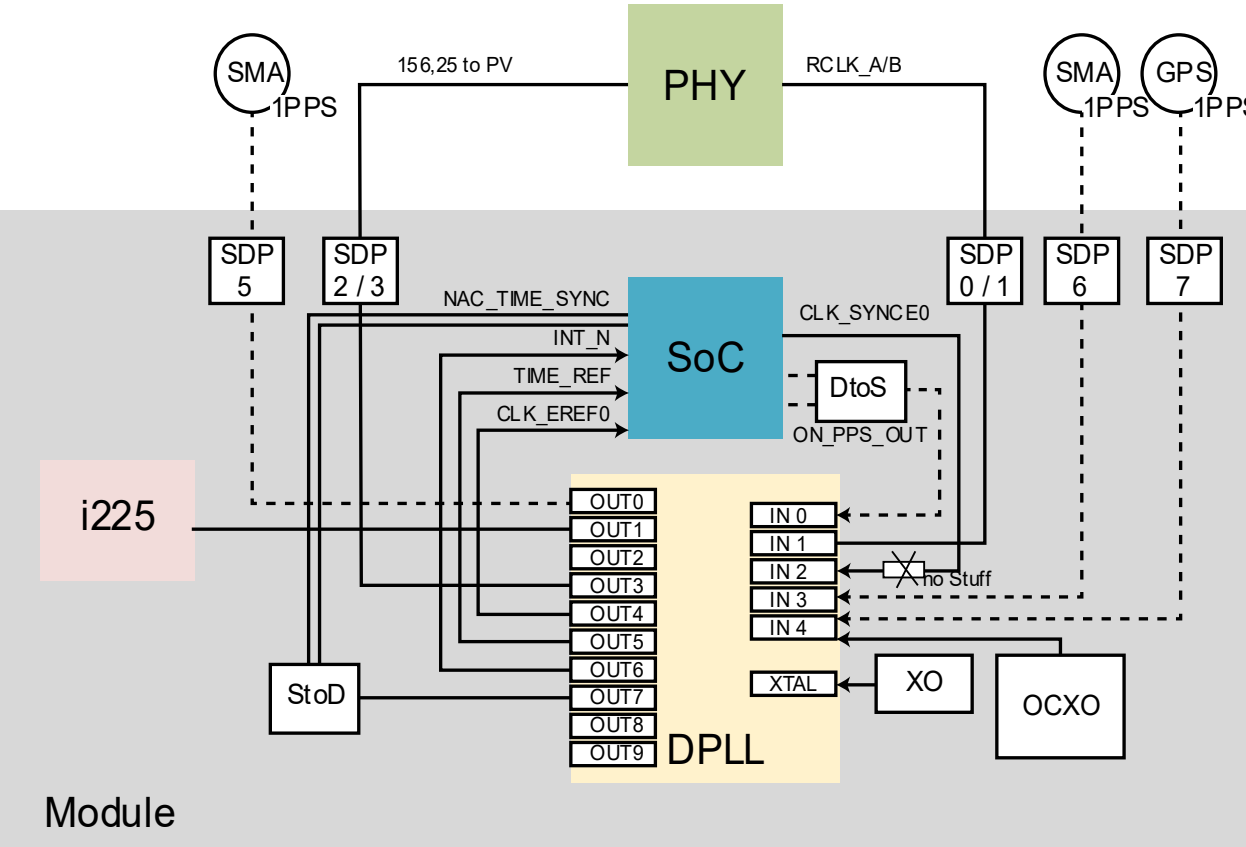
Table 33: *SDP Use in Figure 73 above*

SDP	Meaning	Direction	Notes
0	Recovered Clock A	In	In (to Module)
1	Recovered Clock B	In	
2	Output Clock (+)	Out	Differential Pair on SDP2 and 3
3	Output Clock (-)	Out	
4			
5	1 PPS Out	Out	
6	1 PPS In	In	
7	1 PPS In (GNSS)	In	

SyncE Summary

- Provides physical layer synchronization signal over Ethernet
  - Allows expensive central clock to be shared across the network
- Defined in ITU-T G.8261, G.8262, G.8262.1, and G.8264 specs
  - G.8261 and G.8262 series define physical layer interface
  - G.8264 defines messaging channel used to provide pedigree of clock sources

SyncE can be used alone or in conjunction with PTP:

Table 34:   SyncE / PTP Matrix

Attribute	SyncE Only	PTP Only	SyncE + PTP
Frequency Accuracy	Yes	Yes	Yes
Phase Accuracy	No	Yes	Yes
Time of Day (ToD)	No	Yes	Yes

6.2.   Appendix B: Alternative eDP Example

The alternative eDP example presented in Figures 74 through 79 below comes, with permission, from an Intel reference schematic for a late model CORE series processor. Some parts of the example may not be directly relevant to COM-HPC embedded designs in that they dwell on eDP back-light display power supplies and on a display connector used in certain reference platforms. Nonetheless, the materials may be of interest to some readers and are included in this Appendix.

Figure 74: *Alternative eDP Example (Sheet 1 of 6): Passive Stuffing Options – eDP and DSI*
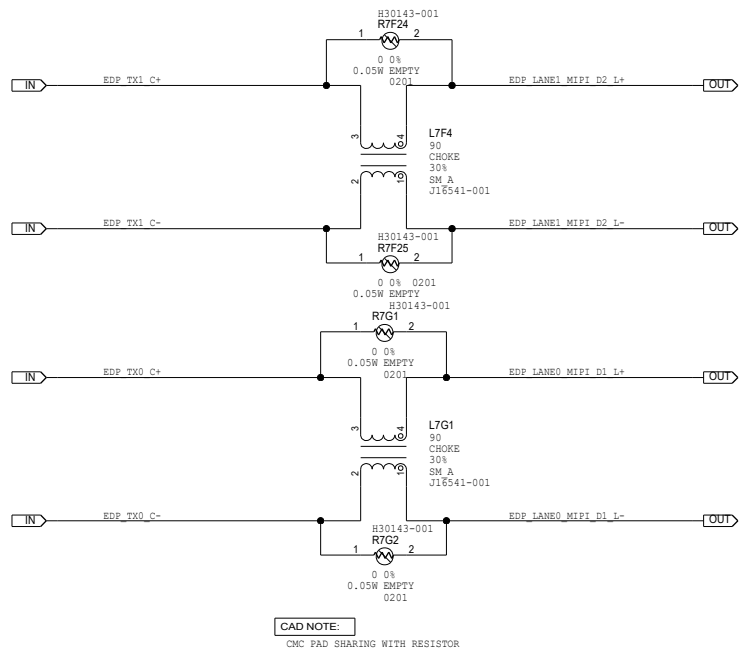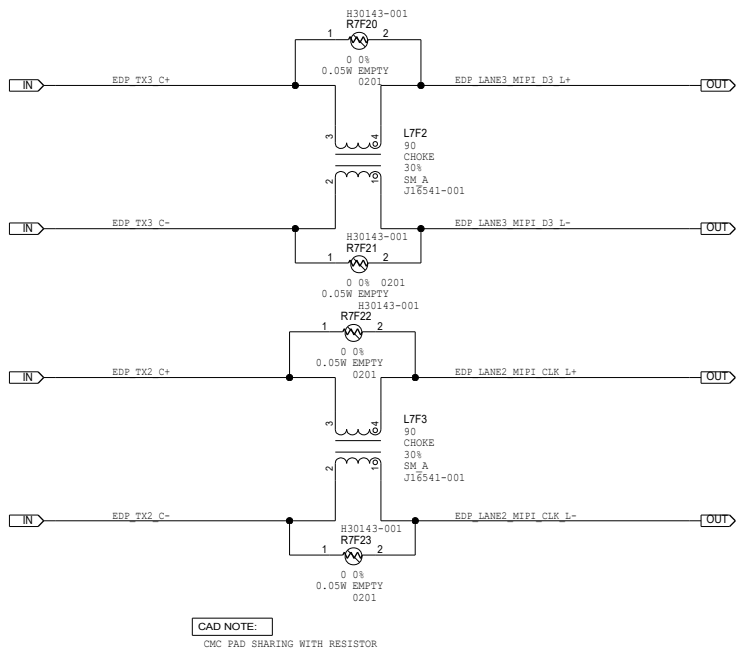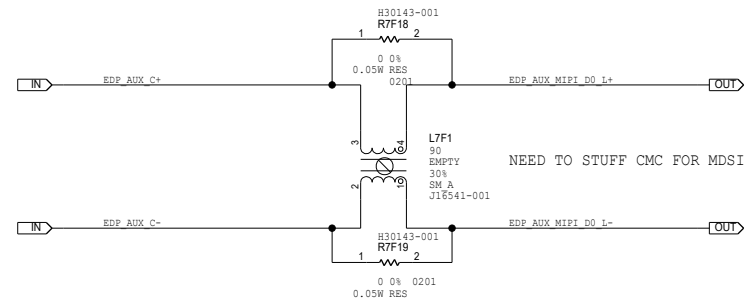
Figure 75:   Alternative eDP Example (Sheet 2 of 6): Backlight Control Options
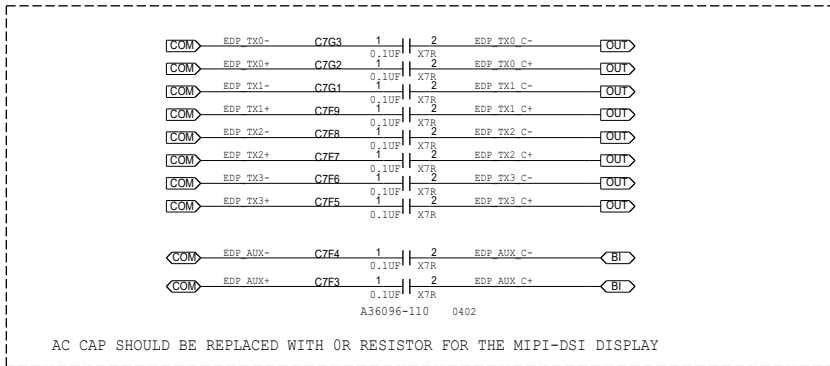
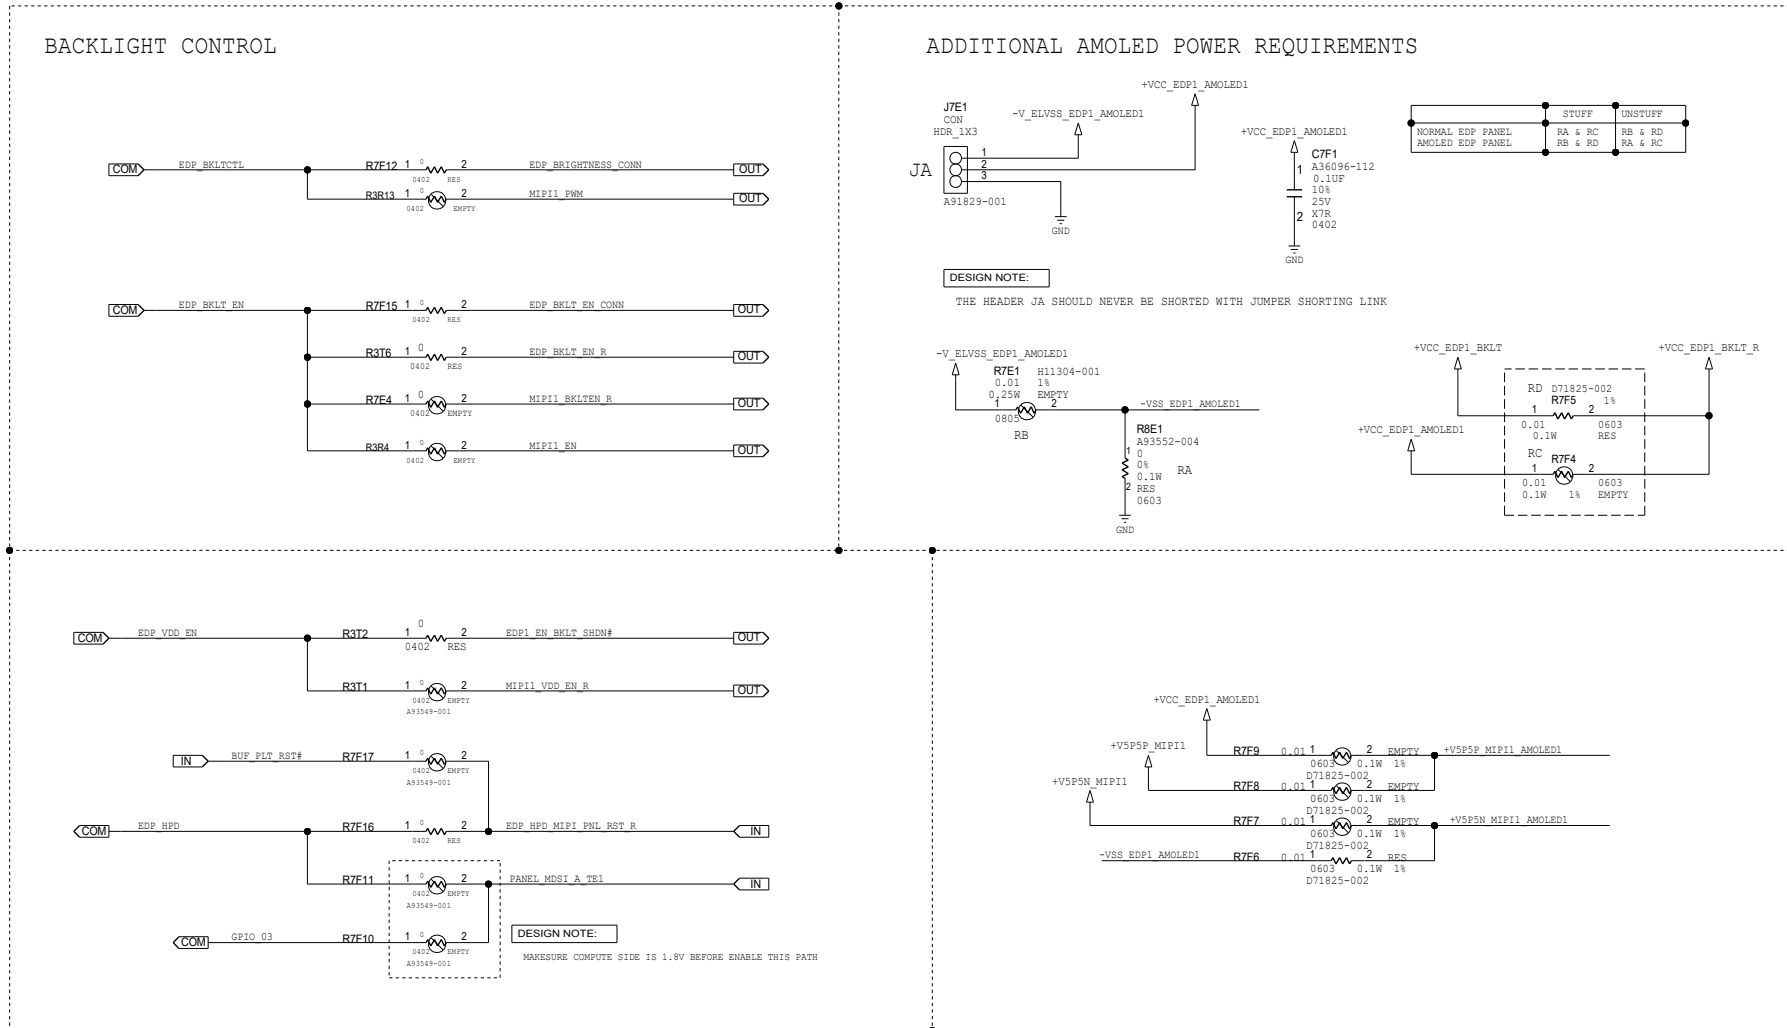*Figure 76:   Alternative eDP Example (Sheet 3 of 6): Connector to Display Panel Assembly*
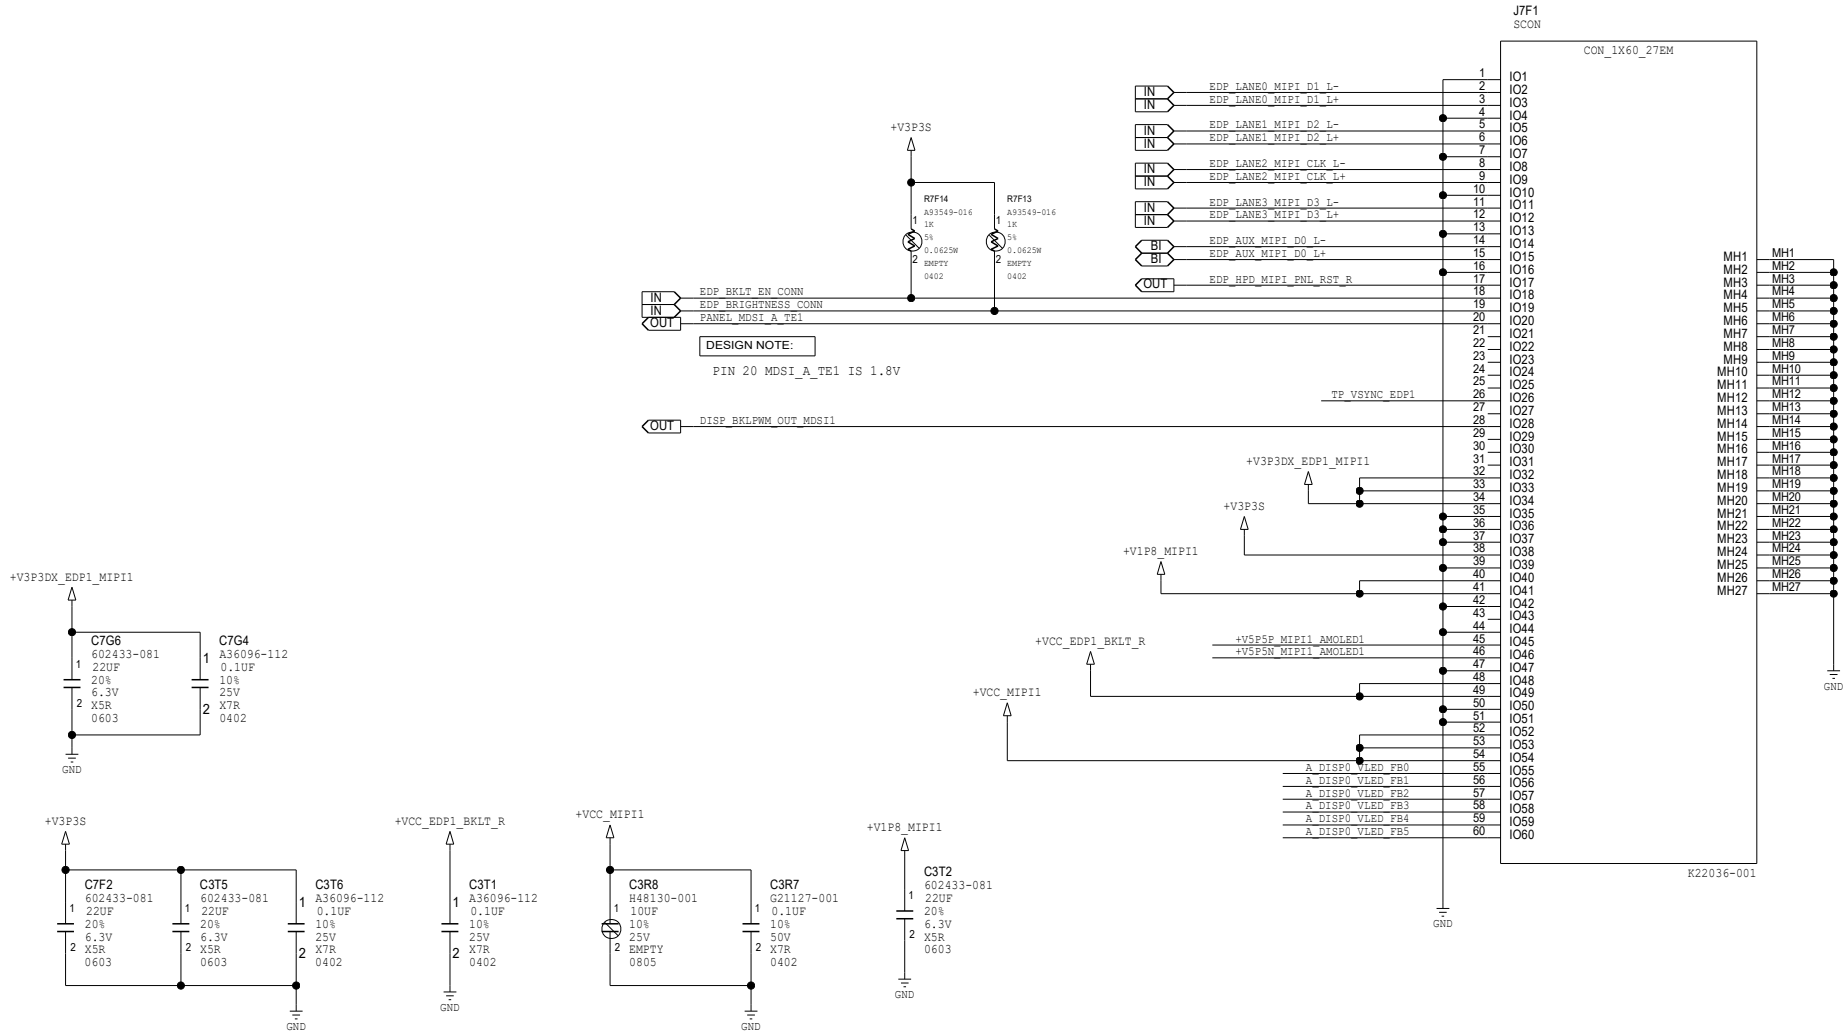
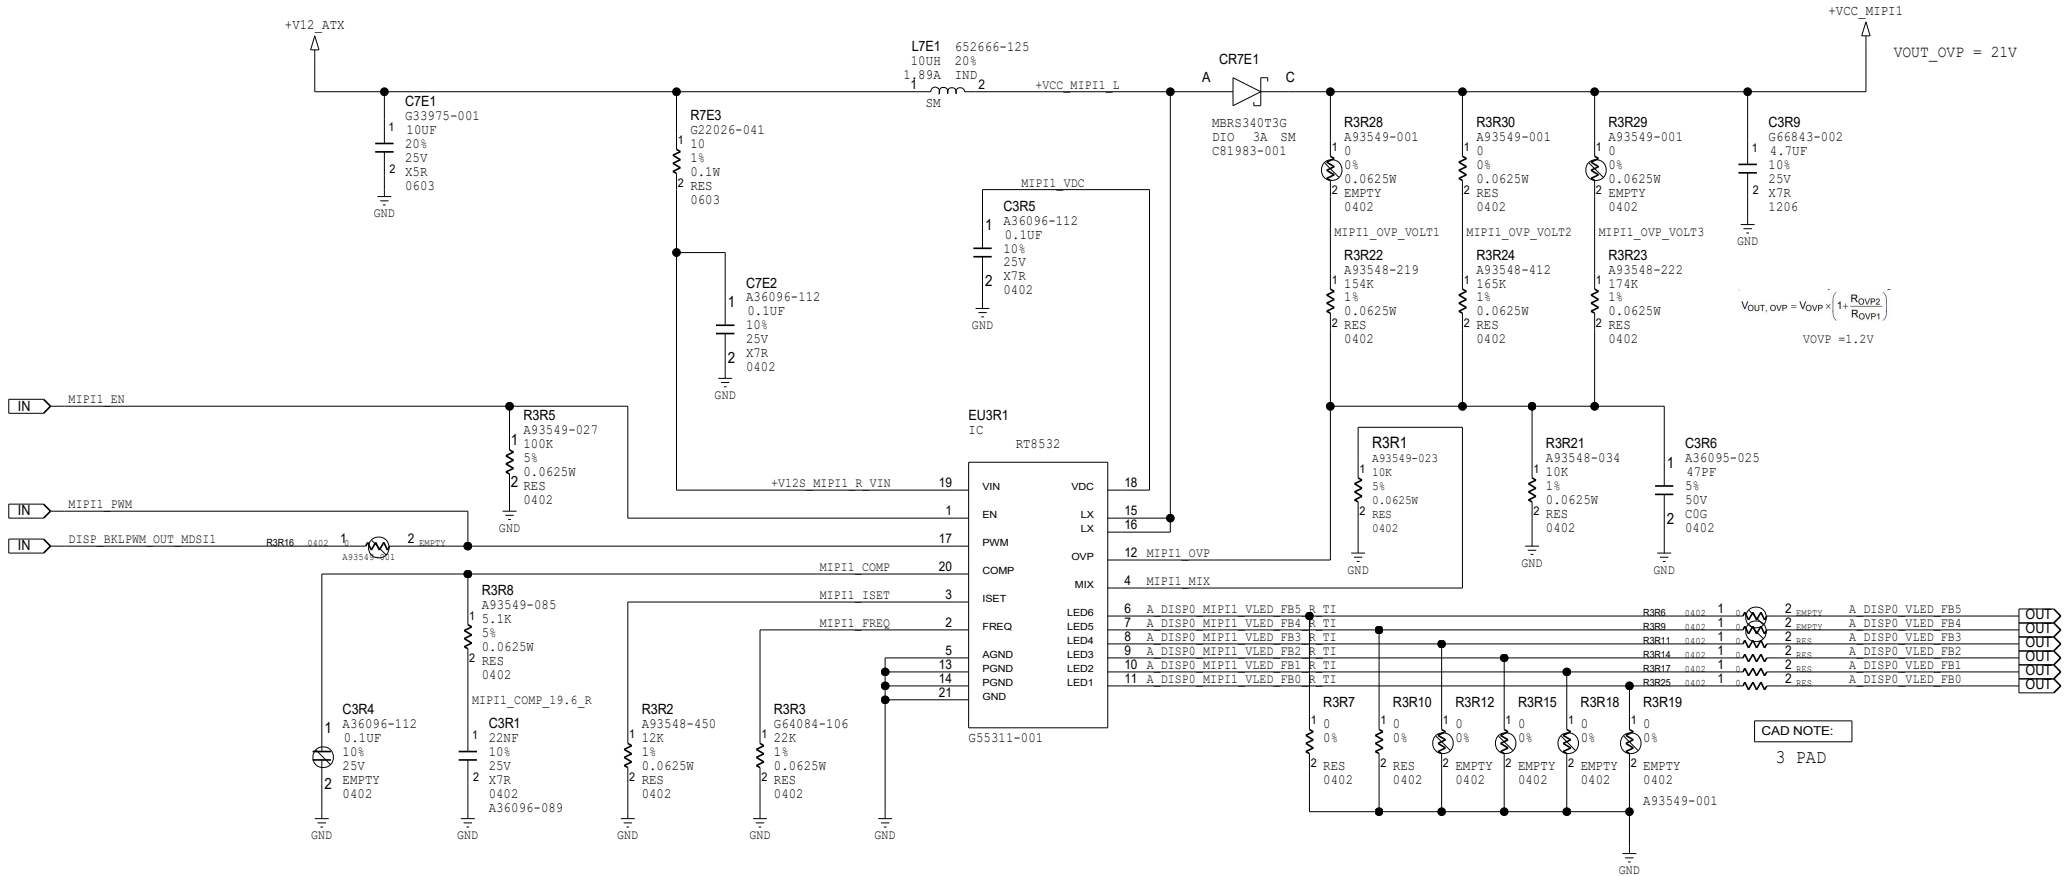*Figure 77: Alternative eDP Example (Sheet 4 of 6): Backlight LED Driver*

Figure 78:   Alternative eDP Example (Sheet 5 of 6): Split Rail (Pos / Neg) PS for AMOLED
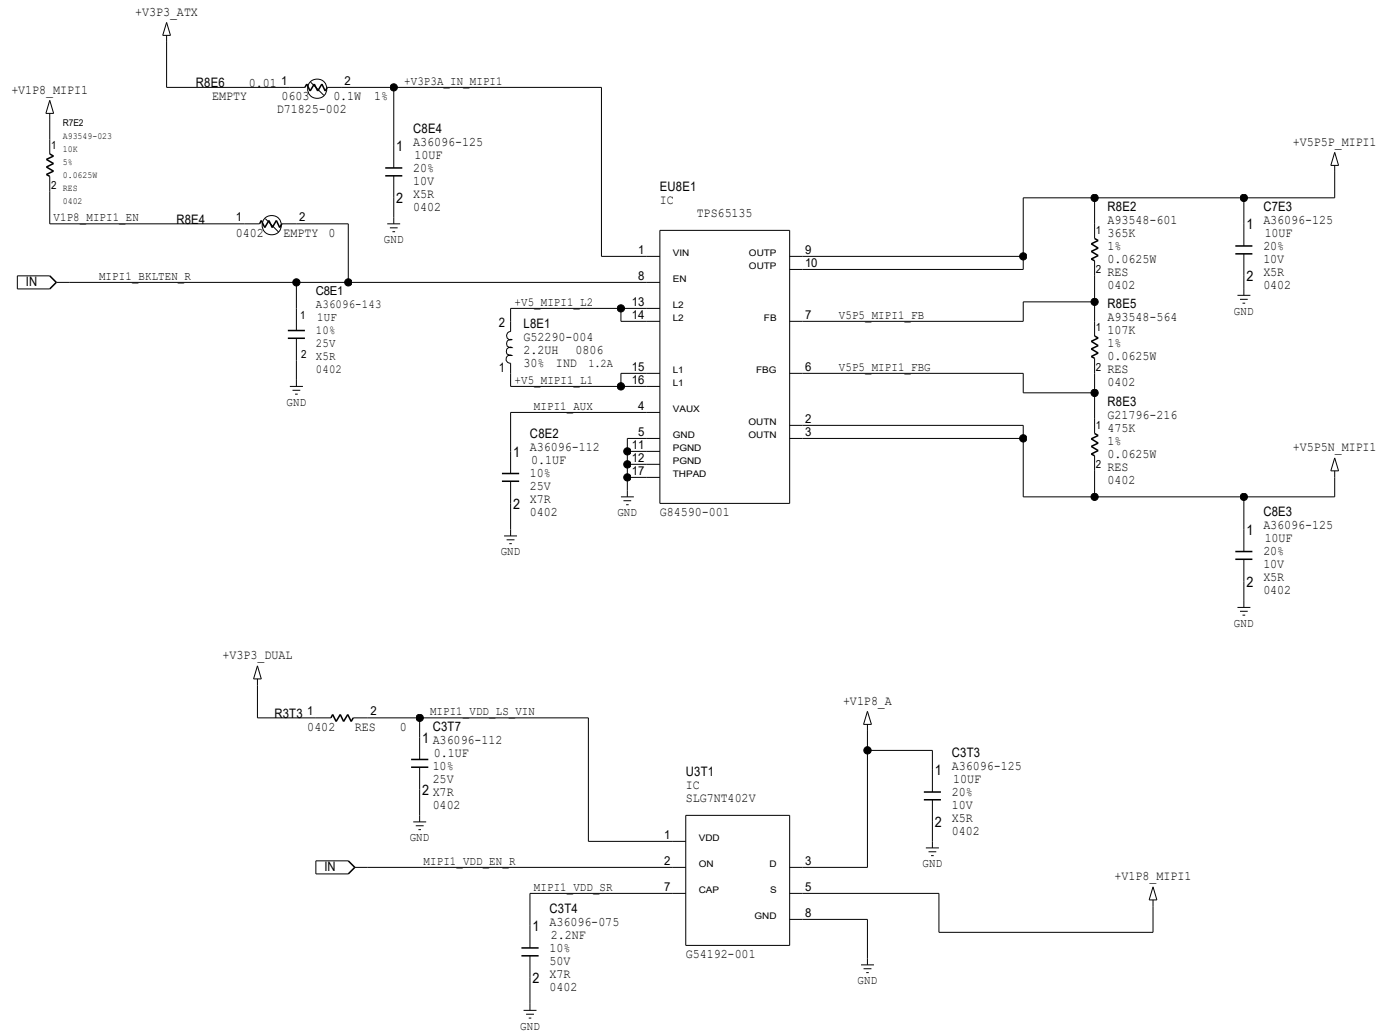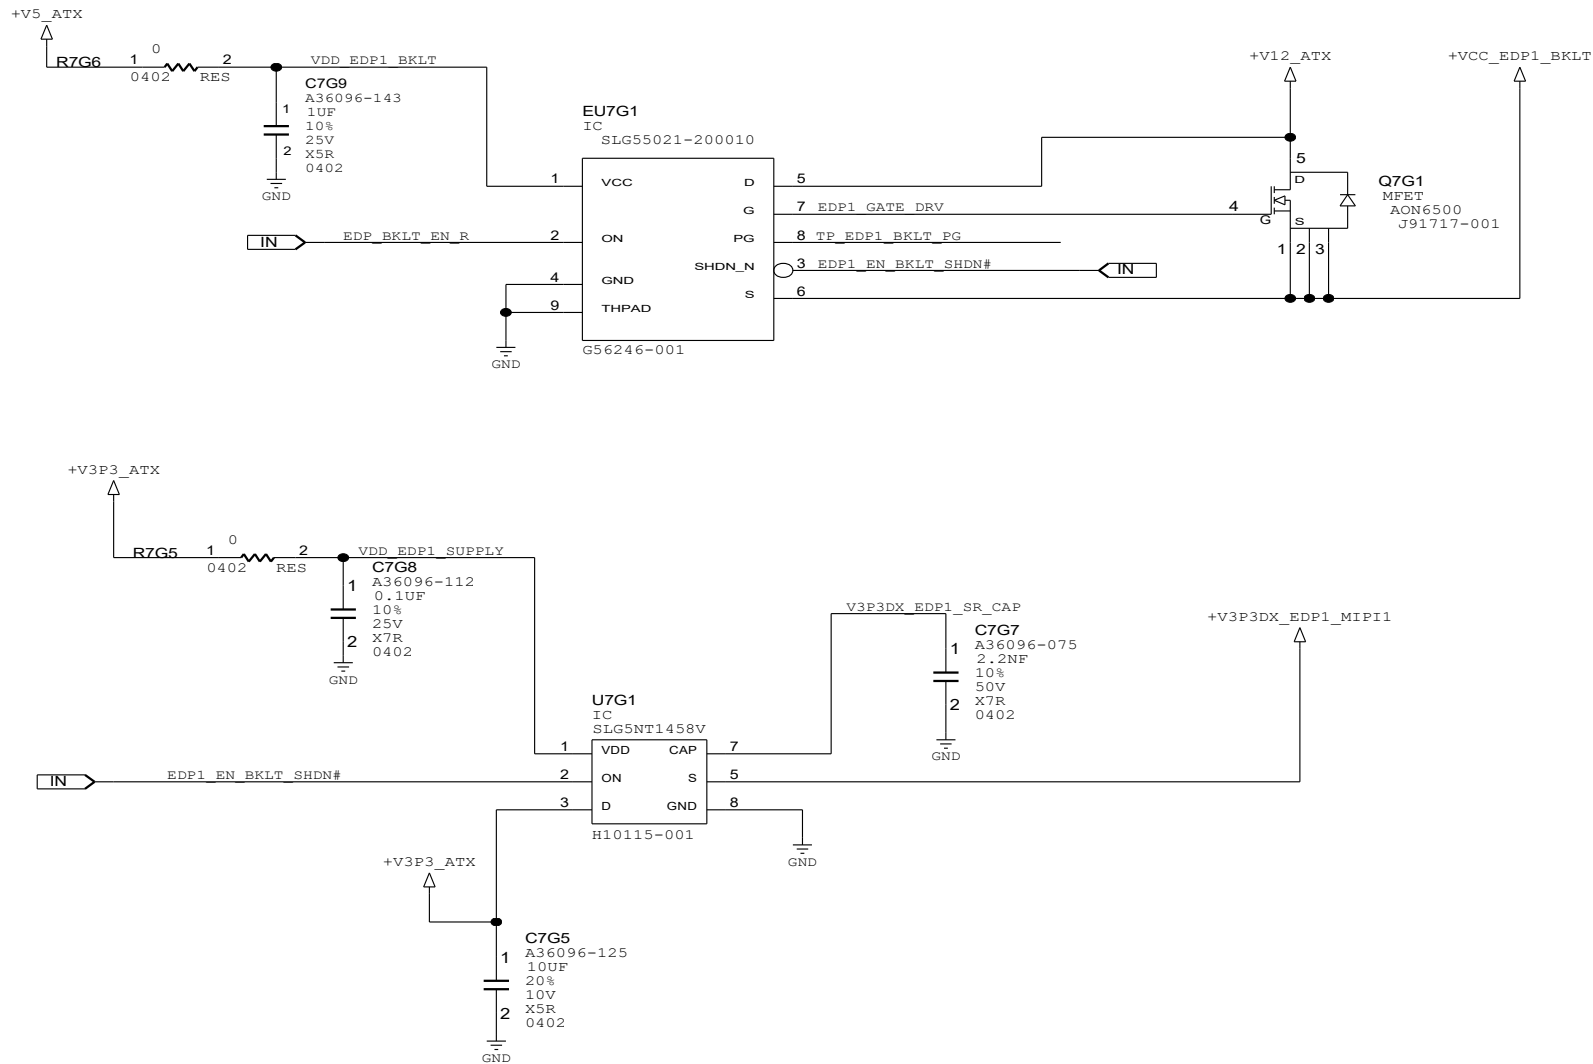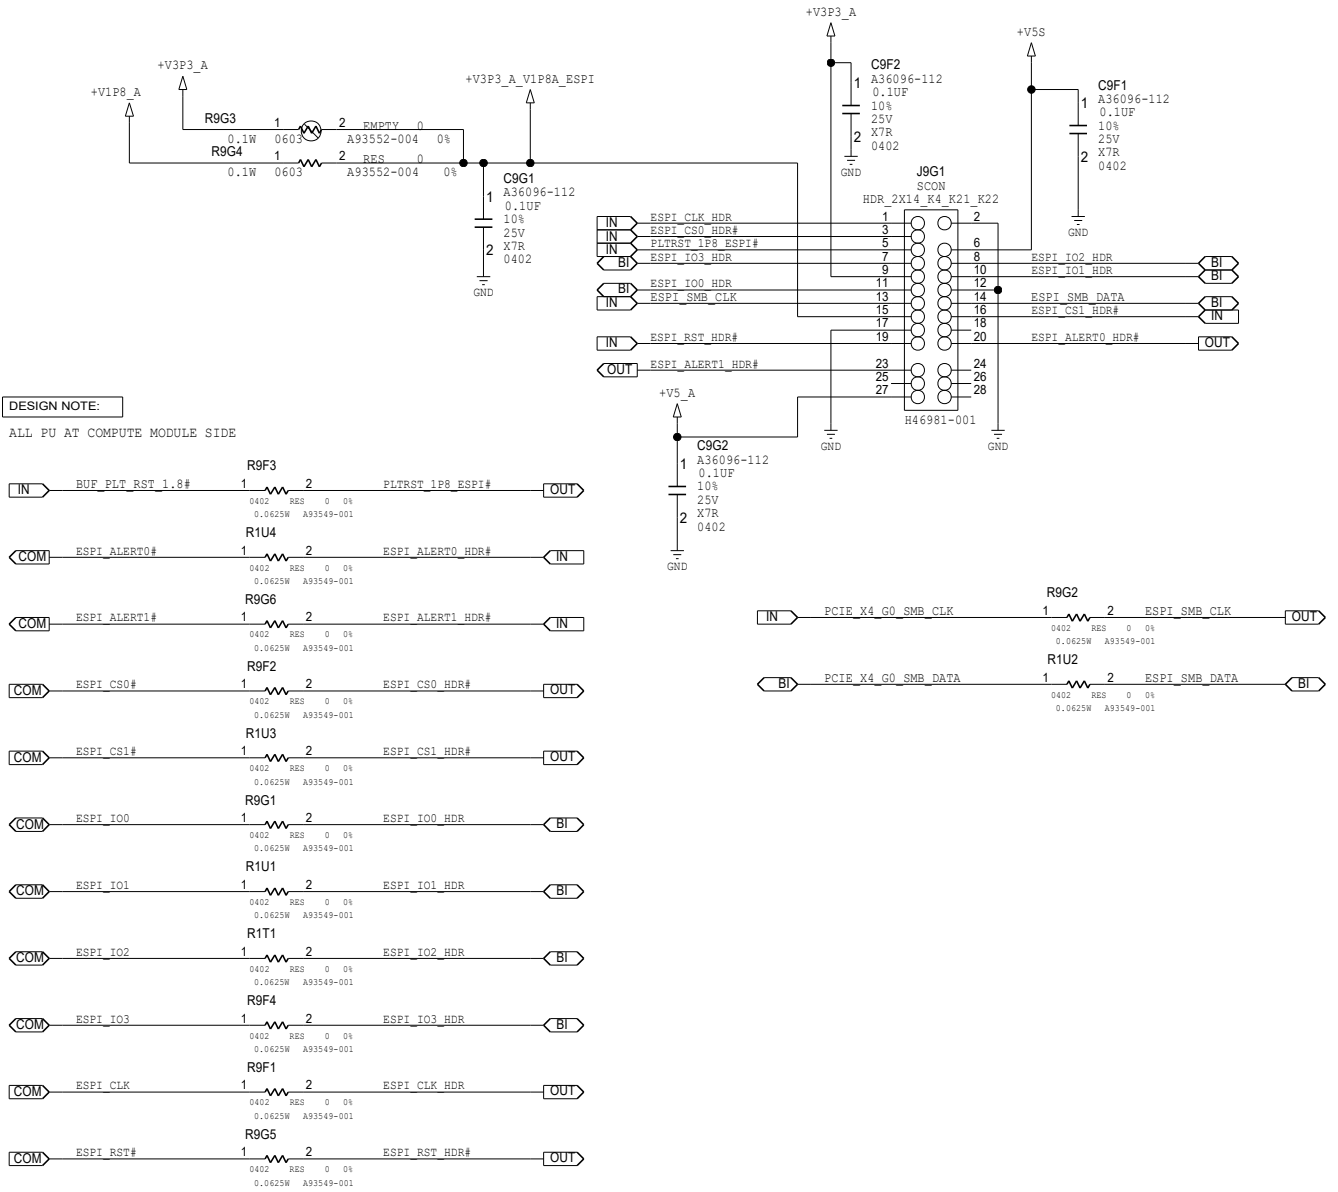
Figure 79: Alternative eDP Example (Sheet 6 of 6): High Side Gate Driver for eDP Backlight

6.3. Appendix C: eSPI Header Example

Figure 80: eSPI Header Example

6.4. Appendix D: Useful Books – General x86 Computer Topics

Table 35: General Books on x86 Computer Topics

Title	Author	Note
PCI Express System Architecture	Ravi Budruk, Don Anderson, Tom Shanley	www.mindshare.com
PCI System Architecture (4th Edition)	Tom Shanley, Don Anderson	www.mindshare.com
Universal Serial Bus System Architecture	Don Anderson	www.mindshare.com
SATA Storage Technology	Don Anderson	www.mindshare.com
Protected Mode Software Architecture (The PC System Architecture Series)	Tom Shanley	www.mindshare.com
The Unabridged Pentium 4	Tom Shanley	www.mindshare.com
Building the Power-Efficient PC: A Developer's Guide to ACPI Power Management, First Edition	Jerzy Kolinski, Ram Chary, Andrew Henroid, and Barry Press	Intel Press, 2002, ISBN 0-9702846-8-3
Hardware Bible	Winn L. Rosch	SAMS, 1997, 0-672-30954-8
The Indispensable PC Hardware Book	Hans-Peter Messmer	Addison-Wesley, 1994, ISBN 0-201-62424-9
The PC Handbook: For Engineers, Programmers, and Other Serious PC Users, Sixth Edition	John P. Choisser and John O. Foster	Annabooks, 1997, ISBN 0-929392-36-1
PC Hardware in a Nutshell, 3rd Edition	Robert Bruce Thompson and Barbara Fritchman Thompson	O'Reilly, 2003, ISBN 0-596-00513-X
PCI & PCI-X Hardware and Software Architecture & Design, Fifth Edition	Edward Solari and George Willse	Annabooks, Intel Press, 2001, ISBN 0-929392-63-9
PCI System Architecture	Tom Shanley and Don Anderson	Addison-Wesley, 2000, ISBN 0-201-30974-2
PCI Express Electrical Interconnect Design: Practical Solutions for Board-level Integration and Validation, First Edition	Dave Coleman, Scott Gardiner, Mohamad Kolberhdari, and Stephen Peters	Intel Press, 2005, ISBN 0-9743649-9-1
Introduction to PCI Express: A Hardware and Software Developer's Guide, First Edition	Adam Wilen, Justin Schade, and Ron Thornburg	Intel Press, 2003, ISBN 0-9702846-9-1
Serial ATA Storage Architecture and Applications, First Edition	Knut Grimsrud and Hubbert Smith	Intel Press, 2003, ISBN 0-9717861-8-6
USB Design by Example, A Practical Guide to Building I/O Devices, Second Edition	John Hyde	Intel Press, ISBN 0-9702846-5-9
Universal Serial Bus System Architecture, Second Edition	Don Anderson and Dave Dzatko	Mindshare, Inc., ISBN 0-201-30975-0
Printed Circuits Handbook, Fourth Edition	Clyde F. Coombs Jr.	McGraw-Hill, 1996, ISBN 0—07-012754-9
High Speed Signal Propagation, First Edition	Howard Johnson and Martin Graham	Prentice Hall, 2003, ISBN 0-13-084408-X
High Speed Digital Design: A Handbook of Black Magic, First Edition	Howard Johnson	Prentice Hall, ISBN: 0133957241
C Programmer's Guide to Serial Communications, Second Edition	Joe Campbell	SAMS, 1987, ISBN 0-672-22584-0
The Programmer's PC Sourcebook, Second Edition	Thom Hogan	Microsoft Press, 1991, ISBN 1-55615-321-X
The Undocumented PC, A Programmer's Guide to I/O, CPUs, and Fixed Memory Areas	Frank van Gilluwe	Addison-Wesley, 1997, ISBN 0-201-47950-8
VHDL Modeling for Digital Design Synthesis	Yu-Chin Hsu, Kevin F. Tsai, Jessie T. Liu and Eric S. Lin	Kluwer Academic Publishers, 1995, ISBN: 0-7923-9597-2

6.5.　　　Appendix E: Revision History

Table 36: Revision History

Release Rev	Date	Author	Notes / Changes
1.0	Mar 17, 2021	C. Eder	CDG preliminary version with Ethernet KR and KR4 CEI diagrams.
2.0	Jan 14, 2022	S. Milnor C. Eder	Release Candidate for first version of complete COM-HPC CDG.
			Revise Figure 35 to show a 50V capacitor for C4V20 Revise Figure 41 to show a 50V capacitor for C5W7 Section 3.14.1. Page 96 – insert short statement about adding HD Audio support to pending COM-HPC Base Spec Rev 1.1 due to lack of SoundWire support Remove references to code names for unreleased Intel products ADL and ICL Revise code name references to show only Intel document numbers
			Incorporate nVent Change Requests
			Change Rev to RC2.0c, change date, re-issue. No other changes.
			Add note to Figure 39 (USB4 ESD diodes) explaining diode positioning Replace Figures 67 and 73 (JSOM diagrams) with revised Figures that call out metric M2.5 hardware rather than 4-40 Imperial hardware
			Formal PICMG release of CDG Revision 2.0
2.1	Aug 10, 2023	S. Milnor C. Eder	Section 4.3.3 SATA Design Rule Summary Delete erroneous requirements in Ref lines 7 and 8 Renumber Rev 2.0 Ref lines 9, 10, 11, 12 to Rev 2.1 Ref lines 7, 8, 9, 10 Remove TX pair to RX pair length matching requirement (Ref line 9 in Rev 2.0, Ref line 7 in Rev 2.1) Section 4.3.4 PCIe Design Rule Summary Remove TX pair to RX pair length matching requirement (Ref line 8) Add missing information on serpentine trace distance to self (Ref line 11)
			Updated Figure 8 (CEI Marvel 88E1543 "Alaska" typo) Updated copyright claims to include year 2022 Update Samtec patent claims in Sections 1.71 and 1.72 This is per request from Samtec patent lawyer on 4/25/2022
			Section 1.7.2. Page 11 Update Samtec "Unnecessary" patent claims per 3/3/2023 and 3/27/2023 input from Samtec Section 1.7.4. Page 12 Add year 2023 to Copyright claim Copyright updated to include 2023 in various other parts of the document Section 1.8. Pages 14 and 13 Correct PCI and PCIe abbreviation explanation ("Interconnect" not "Interface") Fully expand the SATA abbreviation explanation ("Advanced Technology" instead of "AT") Section 3.6.9. Page 57 Add two PCIe Gen 5 capable redriver parts to Table 9 Section 3.10. Page 86 Correct HPD level translator IC reference from U54 to U49
			Formal PICMG release of CDG Revision 2.1 USB4 schematics (Figures 36 through 41) deliberately blurred to satisfy Intel NDA concerns
2.2	Feb 16, 2024	S. Milnor C. Eder	Add numerous notes in a ***distinctive bold blue italic font*** explaining design notes for COM-HPC Mini Remove references to the Intel JHL8040R "Burnside Bridge" USB4 retimer Add in references to the Intel JHL9040R "Hayden Bridge" USB4 retimer Remove detailed USB4 schematics (Figures 36 through 41) due to Intel NDA concerns Add in an Intel provided Hayden Bridge USB4 block diagram for the COM-HPC Client (Figure 36) Add in an Intel provided Hayden Bridge USB4 block diagram for the COM-HPC Mini (Figure 37) Figure numbering and caption details are different between from Figure 36 on in this revision relative to Rev. 2.1 Update Module Type Detection Table in Section 3.24. to include the Mini (Table 21) Add a new Figure for COM-HPC Type detection for the Mini (Figure 57 in revision 2.2)