

casino bacana - Objetivo dos bônus de cassino

Autor: voltracvoltec.com.br Palavras-chave: casino bacana

1. casino bacana
2. casino bacana :lampions bet instagram
3. casino bacana :bonus sexta feira feliz 1xbet como funciona

1. casino bacana :Objetivo dos bônus de cassino

Resumo:

casino bacana : Descubra a adrenalina das apostas em voltracvoltec.com.br! Registre-se hoje e desbloqueie vantagens emocionantes com nosso bônus de boas-vindas!

conteúdo:

online. A parte mais excelente da estrutura do bônus do BetRivers Casino é que há apenas um requisito de 1x playthrough casino bacana casino bacana todos os fundos de bônus recebidos. Código de

us do cassino BetRives CASINOBACK: 100% do primeiro... mlive : cassinos:

de depósito PACASINO250 para obter o seu depósito total combinado até R\$250 casino bacana casino bacana 1X

ônus de dinheiro para jogos de casino. BetRivers Casino Código Promo Até R\$ 250 Deposit Computer expansion bus standard

Not to be confused with PCI-X or PCIe

For Engineering, Procurement, Construction and Installation, see EPCI

Two types of 2 PCIe slot on an Asus H81M-K motherboard

Various slots on a computer motherboard, from top to bottom: PCI Express x4

PCI Express 2 x16

PCI Express x1

PCI Express x16

Conventional PCI (32-bit, 5 V)

PCI Express (Peripheral Component Interconnect Express), officially abbreviated as PCIe or PCIe, [1] 2 is a high-speed serial computer expansion bus standard, designed to replace the older PCI, PCI-X and AGP bus standards. It 2 is the common motherboard interface for personal computers' graphics cards, sound cards, hard disk drive host adapters, SSDs, Wi-Fi and 2 Ethernet hardware connections.[2] PCIe has numerous improvements over the older standards, including higher maximum system bus throughput, lower I/O pin 2 count and smaller physical footprint, better performance scaling for bus devices, a more detailed error detection and reporting mechanism (Advanced 2 Error Reporting, AER),[3] and native hot-swap functionality. More recent revisions of the PCIe standard provide hardware support for I/O virtualization.

The 2 PCI Express electrical interface is measured by the number of simultaneous lanes.[4] (A lane is a single send/receive line of 2 data. The analogy is a highway with traffic in both directions.) The interface is also used in a variety of 2 other standards — most notably the laptop expansion card interface called ExpressCard. It is also used in the storage interfaces 2 of SATA Express, U.2 (SFF-8639) and M.2.

Format specifications are maintained and developed by the PCI-SIG (PCI Special Interest Group) — 2 a group of more than 900 companies that also maintains the conventional PCI specifications.

Architecture [edit]

Example of the PCI 2 Express topology:

white "junction boxes" represent PCI Express device downstream ports, while the gray ones

represent upstream ports.[5] : 7

PCI Express 2 x1 card containing a PCI Express switch (covered by a small heat sink), which creates multiple endpoints out of one 2 endpoint and lets multiple devices share it

The PCIe slots on a motherboard are often labeled with the number of PCIe 2 lanes they have. Sometimes what may seem like a large slot may only have a few lanes. For instance, a 2 x16 slot with only 4 PCIe lanes (bottom slot) is quite common.[6]

Conceptually, the PCI Express bus is a high-speed serial 2 replacement of the older PCI/PCI-X bus.[7] One of the key differences between the PCI Express bus and the older PCI 2 is the bus topology; PCI uses a shared parallel bus architecture, in which the PCI host and all devices share 2 a common set of address, data, and control lines. In contrast, PCI Express is based on point-to-point topology, with separate 2 serial links connecting every device to the root complex (host). Because of its shared bus topology, access to the older 2 PCI bus is arbitrated (in the case of multiple masters), and limited to one master at a time, in a 2 single direction. Furthermore, the older PCI clocking scheme limits the bus clock to the slowest peripheral on the bus (regardless 2 of the devices involved in the bus transaction). In contrast, a PCI Express bus link supports full-duplex communication between any 2 two endpoints, with no inherent limitation on concurrent access across multiple endpoints.

In terms of bus protocol, PCI Express communication is 2 encapsulated in packets. The work of packetizing and de-packetizing data and status-message traffic is handled by the transaction layer of 2 the PCI Express port (described later). Radical differences in electrical signaling and bus protocol require the use of a different 2 mechanical form factor and expansion connectors (and thus, new motherboards and new adapter boards); PCI slots and PCI Express slots 2 are not interchangeable. At the software level, PCI Express preserves backward compatibility with PCI; legacy PCI system software can detect 2 and configure newer PCI Express devices without explicit support for the PCI Express standard, though new PCI Express features are 2 inaccessible.

The PCI Express link between two devices can vary in size from one to 16 lanes. In a multi-lane link, 2 the packet data is striped across lanes, and peak data throughput scales with the overall link width. The lane count 2 is automatically negotiated during device initialization and can be restricted by either endpoint. For example, a single-lane PCI Express (x1) 2 card can be inserted into a multi-lane slot (x4, x8, etc.), and the initialization cycle auto-negotiates the highest mutually supported 2 lane count. The link can dynamically down-configure itself to use fewer lanes, providing a failure tolerance in case bad or 2 unreliable lanes are present. The PCI Express standard defines link widths of x1, x2, x4, x8, and x16. Up to 2 and including PCIe 5.0, x12, and x32 links were defined as well but never used.[8] This allows the PCI Express 2 bus to serve both cost-sensitive applications where high throughput is not needed, and performance-critical applications such as 3D graphics, networking 2 (10 Gigabit Ethernet or multiport Gigabit Ethernet), and enterprise storage (SAS or Fibre Channel). Slots and connectors are only defined 2 for a subset of these widths, with link widths in between using the next larger physical slot size.

As a point 2 of reference, a PCI-X (133 MHz 64-bit) device and a PCI Express 1.0 device using four lanes (x4) have roughly 2 the same peak single-direction transfer rate of 1064 MB/s. The PCI Express bus has the potential to perform better than 2 the PCI-X bus in cases where multiple devices are transferring data simultaneously, or if communication with the PCI Express peripheral 2 is bidirectional.

Interconnect [edit]

A PCI Express link between two devices consists of one or more lanes, which are dual 2 simplex channels using two differential signaling pairs.[5] : 3

PCI Express devices communicate via a logical connection called an interconnect[9] or 2 link. A link is a point-to-point communication channel between two PCI Express ports allowing both of them to send and 2 receive ordinary PCI requests (configuration, I/O or memory read/write) and interrupts (INTx, MSI or MSI-X). At the physical level, a 2 link is composed of one or more lanes.[9] Low-speed peripherals (such as an 802.11 Wi-Fi card) use a single-lane (x1) 2 link, while a graphics adapter typically uses a much wider and therefore faster 16-lane (x16) link.

Lane [edit]

A lane is composed of two differential signaling pairs, with one pair for receiving data and the other for transmitting. Thus, each lane is composed of four wires or signal traces. Conceptually, each lane is used as a full-duplex byte stream, transporting 2 data packets in eight-bit "byte" format simultaneously in both directions between endpoints of a link.[10] Physical PCI Express links may contain 1, 4, 8 or 16 lanes.[11][5]: 4, 5 [9] Lane counts are written with an "x" prefix (for example, "x8" represents an eight-lane card or slot), with x16 being the largest size in common use.[12] Lane sizes are also referred to via the terms "width" or "by" e.g., an eight-lane slot could be referred to as a "by 8" lane or as "8 lanes wide."

For mechanical card sizes, see below.

Serial bus [edit]

The bonded serial bus architecture was chosen over the traditional parallel bus because of the inherent limitations of the latter, including half-duplex operation, excess signal count, and inherently lower bandwidth due to timing skew. Timing skew results from separate electrical signals within a parallel interface traveling through 2 conductors of different lengths, on potentially different printed circuit board (PCB) layers, and at possibly different signal velocities. Despite being transmitted simultaneously as a single word, signals on a parallel interface have different travel duration and arrive at their destinations at different times. When the interface clock period is shorter than the largest time difference between signal arrivals, recovery of the transmitted word is no longer possible. Since timing skew over a parallel bus can amount to a few nanoseconds, the resulting bandwidth limitation is in the range of hundreds of megahertz. Highly simplified topologies of the Legacy PCI Shared (Parallel) Interface and the PCIe Serial Point-to-Point Interface[13]

A serial interface does not exhibit timing skew because there is only one differential signal in each direction within each lane, and there is no external clock signal since clocking information is embedded within the serial signal itself. As such, typical bandwidth limitations on serial signals are in the multi-gigahertz range. PCI Express is one example of the general trend toward replacing parallel buses with serial interconnects; other examples include Serial ATA (SATA), USB, Serial Attached SCSI (SAS), FireWire (IEEE 1394), and RapidIO. In digital video, examples in common use are DVI, HDMI, and DisplayPort.

Multichannel serial design increases flexibility with its ability to allocate fewer lanes for slower devices.

Form factors [edit]

PCI Express (standard) [edit]

Intel P3608 NVMe flash SSD, PCI-E add-in card

A PCI Express card fits into a slot of its physical size or larger (with x16 as the largest used), but may not fit into a smaller PCI Express slot; for example, a x16 card may not fit into a x4 or x8 slot. Some slots use open-ended sockets to permit physically longer cards and negotiate the best available electrical and logical connection.

The number of lanes actually connected to a slot may also be fewer than the number supported by the physical slot size. An example is a x16 slot that runs at x4, which accepts any x1, x2, x4, x8 or x16 card, but provides only four lanes. Its specification may read as "x16 (x4 mode)", while "mechanical @ electrical" notation (e.g. "x16 @ x4") is also common.[citation needed] The advantage is that such slots can accommodate a larger range of PCI Express cards without requiring motherboard hardware to support the full transfer rate. Standard mechanical sizes are x1, x4, x8, and x16. Cards using a number of lanes other than the standard mechanical sizes need to physically fit the next larger mechanical size (e.g. an x2 card uses the x4 size, or an x12 card uses the x16 size).

The cards themselves are designed and manufactured in various sizes. For example, solid-state drives (SSDs) that come in the form of PCI Express cards often use HHHL (half height, half length) and FHHL (full height, half length) to describe the physical dimensions of the card.[14][15]

PCI card type	Dimensions height × length × width, maximum (mm) (in)
Full-Length	111.15 × 312.00 × 20.32
4.376 × 12.283 × 0.8 Half-Length	111.15 × 167.65 × 20.32
4.376 × 0.6.600 × 0.8	4.376 × 6.600 × 0.8

Low-Profile/Slim 0 68.90 × 167.65 × 20.32 2.731 × 0 6.600 2 × 0.8

Non-standard video card form factors [edit]

Modern (since c. 2012[16]) gaming video cards usually exceed the height as well as thickness specified in the PCI Express standard, due to the need for more capable and quieter cooling fans, as gaming video cards often emit hundreds of watts of heat.[17] Modern computer cases are often wider to accommodate these taller cards, but not always. Since full-length cards (312 mm) are uncommon, modern cases sometimes cannot fit those. The thickness of these cards also typically occupies the space of two PCIe slots. In fact, even the methodology of how to measure the cards varies between vendors, with some including the metal bracket size in dimensions and others not.

For instance, comparing three high-end video cards released in 2024: a Sapphire Radeon RX 5700 XT card measures 135 mm in height (excluding the metal bracket), which exceeds the PCIe standard height by 28 mm,[18] another Radeon RX 5700 XT card by XFX measures 55 mm thick (i.e. 2.7 PCI slots at 20.32 mm), taking up three PCIe slots,[19] while an Asus GeForce RTX 3080 video card takes up two slots and measures 318.5 mm × 140.1 mm × 57.8 mm, exceeding PCI Express' maximum length, height, and thickness respectively.[20]

Pinout [edit]

The following table identifies the conductors on each side of the edge connector on a PCI Express card. The solder side of the printed circuit board (PCB) is the A-side, and the component side is the B-side.[21] PRSNT1# and PRSNT2# pins must be slightly shorter than the rest, to ensure that a hot-plugged card is fully inserted. The WAKE# pin uses full voltage to wake the computer, but must be pulled high from the standby power to indicate that the card is wake capable.[22]

PCI Express connector pinout (x1, x4, x8 and x16 variants)	Pin	Side B	Side A	Description	Pin																			
	0	1	+12 V	PRSNT1# Must connect to farthest PRSNT2# pin	50																			
	HSOp(8)	Reserved	Lane 8	transmit data, + and	0																			
	2	+12 V	+12 V	Main power pins	51																			
	HSOn(8)	Reserved	Lane 8	receive data, + and	0																			
	3	+12 V	+12 V	52 Ground	53																			
	Ground	HSIn(8)	0	5 SMCLK TCK SMBus and JTAG port pins	54																			
	2	HSOp(9)	Ground	Lane 9																				
				transmit data, + and	0																			
	6	SMDAT	TDI	55 HSON(9) Ground	0																			
	7	Ground	TDO	2	56																			
	Ground	HSIp(9)	Lane 9	receive data, + and	0																			
	8	+3.3 V	TMS	57 Ground	HSIn(9)	0																		
	9	2	TRST#	+3.3 V	58																			
	HSOp(10)	Ground	Lane 10	transmit data, + and	10																			
	+3.3 V	aux	+3.3 V	Aux 2 power & Standby power	59																			
	HSOn(10)	Ground	11	WAKE# PERST# Link reactivation; fundamental reset	[23] 60																			
	Ground	HSIp(10)	Lane 10	2 receive data, + and	Key notch	61																		
	Ground	HSIn(10)	12	CLKREQ#[24] Ground																				
				Clock Request Signal	62																			
	HSOp(11)	Ground	Lane 2	11 transmit data, + and	13																			
	Ground	REFCLK+	Reference clock differential pair	63	HSOn(11) Ground	14																		
	HSOp(0)	REFCLK	Lane 2	0 transmit data, + and	64																			
	Ground	HSIp(11)	Lane 11	receive data, + and	15																			
	HSOn(0)	Ground	65	2	Ground																			
	HSIn(11)	16	Ground	HSIp(0) Lane 0	receive data, + and	66																		
	HSOp(12)	Ground	Lane 12	transmit data, + and	17																			
	PRSNT2#	HSIn(0)	67	HSOn(12) Ground	18																			
	Ground	Ground	68	Ground																				
	HSIp(12)	Lane 12	receive data, + and	2	PCI Express x1 cards end at pin	18																		
	69	Ground	HSIn(12)	19	HSOp(1) Reserved	Lane 1	transmit data, + and	70																
	HSOp(13)	Ground	Lane 13	transmit data, + and	20																			
	HSOn(1) Ground	71	HSOn(13) Ground	21	Ground	2																		
	HSIp(1) Lane 1	receive data, + and	72	Ground	HSIp(13) Lane 13	receive data, + and	22																	
	Ground	2	HSIn(1) 73	Ground	HSIn(13) 23	HSOp(2) Ground	Lane 2	transmit data, + and	74															
	HSOp(14) Ground	Lane 14	transmit 2 data, + and	24	HSOn(2) Ground	75	HSOn(14) Ground	25	Ground	HSIp(2) Lane 2	receive data, + and	2												
	76	Ground	HSIp(14) Lane 14	receive data, + and	26	Ground	HSIn(2) 77	Ground	HSIn(14) 27	HSOp(3) Ground	Lane 2	3	transmit data, + and	78										
	HSOp(15) Ground	Lane 15	transmit data, + and	28	HSOn(3) Ground	79	2	HSOn(15) Ground	29	Ground	HSIp(3) Lane 3	receive data, + and	30											
	PWRBRK#[25] HSIn(3) 81	PRSNT2# HSIn(15) 31	PRSNT2# Ground	82	Reserved	Ground	2	32	Ground	Reserved	PCI Express x4 cards end at pin	32	33	HSOp(4) Reserved	Lane 4	transmit data, + and	2	34	HSOn(4) Ground	35	Ground	HSIp(4) Lane 4	receive data, + and	36
	Ground	HSIn(4) 37	HSOp(5) Ground	2	Lane 5	transmit data, + and	38	HSOn(5) Ground	39	Ground	HSIp(5) Lane 5	receive data, + and	2	40	Ground	HSIn(5) 41	HSOp(6) Ground	Lane 6						

transmit data, + and 42 HSON(6) Ground 43 Ground HSlp(6) Lane 2 6 receive data, + and Legend 44 Ground HSlN(6) Ground pin Zero volt reference 45 HSOp(7) Ground Lane 7 2 transmit data, + and Power pin Supplies power to the PCIe card 46 HSON(7) Ground Card-to-host pin Signal from 2 the card to the motherboard 47 Ground HSlp(7) Lane 7 receive data, + and Host-to-card pin Signal from the 2 motherboard to the card 48 PRSNT2# HSlN(7) Open drain May be pulled low or sensed by multiple cards 49 Ground 2 Ground Sense pin Tied together on card PCI Express x8 cards end at pin 49 Reserved Not presently used, do 2 not connect Power [edit]

The main 12 12 V 3.3 3.3 V 25 25 W 75 75 W [26]

Slot power 2 [edit]

All PCI express cards may consume up to 3 A at +3.3 V (9.9 W). The amount of 2 +12 V and total power they may consume depends on the form factor and the role of the card:[27]: 35–36 2 [28][29]

x1 cards are limited to 0.5 A at +12 V (6 W) and 10 W combined.

V (6 W) and 10 2 W combined. x4 and wider cards are limited to 2.1 A at +12 V (25 W) and 25 W combined.

V 2 (25 W) and 25 W combined. A full-sized x1 card may draw up to the 25 W limits after initialization 2 and software configuration as a high-power device.

A full-sized x16 graphics card may draw up to 5.5 A at +12 V 2 (66 W) and 75 W combined after initialization and software configuration as a high-power device.[22] : 38–39

6- and 8-pin power 2 connectors [edit]

8-pin (left) and 6-pin (right) power connectors used on PCI Express cards

Optional connectors add 75 W (6-pin) 2 or 150 W (8-pin) of +12 V power for up to 300 W total (2 × 75 W + 1 2 × 150 W).

Sense0 pin is connected to ground by the cable or power supply, or float on board if cable 2 is not connected.

Sense1 pin is connected to ground by the cable or power supply, or float on board if cable 2 is not connected.

Some cards use two 8-pin connectors, but this has not been standardized yet as of 2024 , therefore 2 such cards must not carry the official PCI Express logo. This configuration allows 375 W total (1 × 75 W 2 + 2 × 150 W) and will likely be standardized by PCI-SIG with the PCI Express 4.0 standard.[needs update] The 2 8-pin PCI Express connector could be confused with the EPS12V connector, which is mainly used for powering SMP and multi-core 2 systems. The power connectors are variants of the Molex Mini-Fit Jr. series connectors.[30]

Molex Mini-Fit Jr. part numbers[30] Pins Female/receptacle

on PS 2 cable Male/right-angle

header on PCB 6-pin 45559-0002 45558-0003 8-pin 45587-0004 45586-0005, 45586-0006

6-pin power connector (75 W)[31] 8-pin power connector (150 W)[32][33][34] 2 6 pin power

connector pin map 8 pin power connector pin map Pin Description Pin Description 1 +12 V 1 2 +12

V 2 Not connected (usually +12 V as well) 2 +12 V 3 +12 V 3 +12 V 4 2 Sense1 (8-pin

connected[A]) 4 Ground 5 Ground 5 Sense 6 Sense0 (6-pin or 8-pin connected) 6 Ground 7

Ground 8 2 Ground

^ When a 6-pin connector is plugged into an 8-pin receptacle the card is notified by a missing Sense1 that 2 it may only use up to 75 W.

12VHPWR connector [edit]

This section is an excerpt from 16-Pin 12VHPWR connector 2 12VHPWR connector The 16-pin 12VHPWR connector is a standard for connecting graphics processing units (GPUs) to computer power supplies. It 2 was introduced in the early 2024s to supersede the previous 6- and 8-pin power connectors for GPUs. The primary aim 2 was to cater to the increasing power requirements of contemporary high-performance GPUs, ensuring reliable power delivery and increasing performance. It 2 was replaced by a minor revision called 12V-2x6, which changed the connector to ensure that the sense pins only make 2 contact if the power pins are seated properly. The original standard was formally adopted as part of PCI Express 5.x,[35] 2 while the revised 12V-2x6

connector design was adopted later.[36] The original standard was formally adopted as part of PCI Express 2.5.x, while the revised 12V-2x6 connector design was adopted later.

PCI Express Mini Card [edit]

A WLAN PCI Express Mini Card 2 and its connector

MiniPCI and MiniPCI Express cards in comparison

PCI Express Mini Card (also known as Mini PCI Express, Mini PCIe, 2 Mini PCI-E, mPCIe, and PEM), based on PCI Express, is a replacement for the Mini PCI form factor. It is 2 developed by the PCI-SIG. The host device supports both PCI Express and USB 2.0 connectivity, and each card may use 2 either standard. Most laptop computers built after 2005 use PCI Express for expansion cards; however, as of 2024 , many 2 vendors are moving toward using the newer M.2 form factor for this purpose.

Due to different dimensions, PCI Express Mini Cards 2 are not physically compatible with standard full-size PCI Express slots; however, passive adapters exist that let them be used in 2 full-size slots.[37]

Physical dimensions [edit]

Dimensions of PCI Express Mini Cards are 30 mm × 50.95 mm (width × length) 2 for a Full Mini Card. There is a 52-pin edge connector, consisting of two staggered rows on a 0.8 mm 2 pitch. Each row has eight contacts, a gap equivalent to four contacts, then a further 18 contacts. Boards have a 2 thickness of 1.0 mm, excluding the components. A "Half Mini Card" (sometimes abbreviated as HMC) is also specified, having approximately 2 half the physical length of 26.8 mm.

Electrical interface [edit]

PCI Express Mini Card edge connectors provide multiple connections and 2 buses:

PCI Express x1 (with SMBus)

USB 2.0

Wires to diagnostics LEDs for wireless network (i.e., Wi-Fi) status on computer's chassis

SIM card for 2 GSM and WCDMA applications (UIM signals on spec.)

Future extension for another PCIe lane

1.5 V and 3.3 V power

Mini-SATA (mSATA) variant 2 [edit]

An Intel mSATA SSD

Despite sharing the Mini PCI Express form factor, an mSATA slot is not necessarily electrically 2 compatible with Mini PCI Express. For this reason, only certain notebooks are compatible with mSATA drives. Most compatible systems are 2 based on Intel's Sandy Bridge processor architecture, using the Huron River platform. Notebooks such as Lenovo's ThinkPad T, W and 2 X series, released in March–April 2011, have support for an mSATA SSD card in their WWAN card slot. The ThinkPad 2 Edge E220s/E420s, and the Lenovo IdeaPad Y460/Y560/Y570/Y580 also support mSATA.[38] On the contrary, the L-series among others can only support 2 M.2 cards using the PCIe standard in the WWAN slot.

Some notebooks (notably the Asus Eee PC, the Apple MacBook Air, 2 and the Dell mini9 and mini10) use a variant of the PCI Express Mini Card as an SSD. This variant 2 uses the reserved and several non-reserved pins to implement SATA and IDE interface passthrough, keeping only USB, ground lines, and 2 sometimes the core PCIe x1 bus intact.[39] This makes the "miniPCIe" flash and solid-state drives sold for netbooks largely incompatible 2 with true PCI Express Mini implementations.

Also, the typical Asus miniPCIe SSD is 71 mm long, causing the Dell 51 mm 2 model to often be (incorrectly) referred to as half length. A true 51 mm Mini PCIe SSD was announced in 2 2009, with two stacked PCB layers that allow for higher storage capacity. The announced design preserves the PCIe interface, making 2 it compatible with the standard mini PCIe slot. No working product has yet been developed.

Intel has numerous desktop boards with 2 the PCIe x1 Mini-Card slot that typically do not support mSATA SSD. A list of desktop boards that natively support 2 mSATA in the PCIe x1 Mini-Card slot (typically multiplexed with a SATA port) is provided on the Intel Support site.[40]

PCI Express M.2 [edit]

M.2 replaces the mSATA standard and Mini PCIe.[41] Computer bus interfaces provided through the M.2 connector are PCI Express 3.0 (up to four lanes), Serial ATA 3.0, and USB 3.0 (a single logical port for each of the latter two). It is up to the manufacturer of the M.2 host or device to choose which interfaces to support, depending on the desired level of host support and device type.

PCI Express External Cabling [edit]

PCI Express External Cabling (also known as External PCI Express, Cabled PCI Express, or ePCIe) specifications were released by the PCI-SIG in 2 February 2007.[42][43]

Standard cables and connectors have been defined for x1, x4, x8, and x16 link widths, with a transfer rate of 250 MB/s per lane. The PCI-SIG also expects the norm to evolve to reach 500 MB/s, as in PCI Express 2.0. An example of the uses of Cabled PCI Express is a metal enclosure, containing a number of PCIe slots and PCIe-to-ePCIe adapter circuitry. This device would not be possible had it not been for the ePCIe specification.

PCI Express OCuLink [edit]

OCuLink (standing for "optical-copper link", since Cu is the chemical symbol for copper) is an extension for the "cable version of PCI Express", designed to compete with Thunderbolt 3. Version 1.0 of OCuLink, released in Oct 2024, supports up to 4 PCIe 3.0 lanes (8 GT/s (gigatransfers per second), 3.9 GB/s) over copper cabling; a fiber optic version may appear in the future.

The most recent version of OCuLink, OCuLink-2, supports up to 16 GB/s (PCIe 4.0 x8)[44] while the maximum bandwidth of a full speed Thunderbolt 4 cable is 5 GB/s. Some suppliers may design their connector products to support the next-generation PCI Express 5.0, which runs at 32 GT/s per lane, ensuring future proofing and minimizing development costs over the next few years.[44]

While initially intended for use in laptops for the connection of powerful external GPU boxes, OCuLink's popularity lies primarily in its use for PCIe interconnections in servers, a more prevalent application.[45]

Derivative forms [edit]

Numerous other form factors use, or are able to use, PCIe. These include:

Low-height card

ExpressCard: Successor to the PC Card form factor (with x1 PCIe and USB 2.0; hot-pluggable)

PCI Express ExpressModule: A hot-pluggable modular form factor defined for servers and workstations

XQD card: A PCI Express-based flash card standard by the CompactFlash Association with x2 PCIe

CFexpress card: A PCI Express-based flash card by the CompactFlash Association in three form factors supporting 1 to 4 PCIe lanes

SD card: The SD Express bus, introduced in version 7.0 of the SD specification uses a x1 PCIe link

XMC: Similar to the CMC/PMC form factor (VITA 42.3)

AdvancedTCA: A complement to CompactPCI for larger applications; supports serial based backplane topologies

AMC: A complement to the AdvancedTCA specification; supports processor and I/O modules on ATCA boards (x1, x2, x4 or x8 PCIe).

FeaturePak: A tiny expansion card format (43 mm × 65 mm) for embedded and small-form-factor applications, which implements two x1 PCIe links on a high-density connector along with USB, I2C, and up to 100 points of I/O

mm × 65 mm) for embedded and small-form-factor applications, which implements two x1 PCIe links on a high-density connector along with USB, I2C, and up to 100 points of I/O Universal IO: A variant from Super Micro Computer Inc designed for use in low-profile rack-mounted chassis.

[46] It has the connector bracket reversed so it cannot fit in a normal PCI Express socket, but it is pin-compatible and may be inserted if the bracket is removed.

It has the connector bracket reversed so it cannot fit in a normal PCI Express socket, but it is pin-compatible and may be inserted if the bracket is removed. M.2 (formerly known as NGFF) M-PCIe brings PCIe 3.0 to 2 mobile devices (such as tablets and smartphones), over the M-PHY physical layer. [47] [48]

U.2 (formerly known as SFF-8639)

The PCIe slot 2 connector can also carry protocols other than PCIe. Some 9xx series Intel chipsets support Serial Digital Video Out, a proprietary 2 technology that uses a slot to transmit video signals from the host CPU's integrated graphics instead of PCIe, using a 2 supported add-in. The PCIe transaction-layer protocol can also be used over some other interconnects, which are not electrically PCIe:

Thunderbolt: A royalty-free 2 interconnect standard by Intel that combines DisplayPort and PCIe protocols in a form factor compatible with Mini DisplayPort. Thunderbolt 3.0 2 also combines USB 3.1 and uses the USB-C form factor as opposed to Mini DisplayPort.

USB4

History and revisions [edit]

While 2 in early development, PCIe was initially referred to as HSI (for High Speed Interconnect), and underwent a name change to 2 3GIO (for 3rd Generation I/O) before finally settling on its PCI-SIG name PCI Express. A technical working group named the 2 Arapaho Work Group (AWG) drew up the standard. For initial drafts, the AWG consisted only of Intel engineers; subsequently, the 2 AWG expanded to include industry partners.

Since, PCIe has undergone several large and smaller revisions, improving on performance and other features.

Comparison 2 table [edit]

PCI Express link performance	Version	Introduced	Line code	Transfer rate per lane	Throughput
x1	1.0	2003	NRZ 8b/10b	2.5 GT/s	0.250 GB/s
x2	2.0	2007	5.0 GT/s	2.0	0.500 GB/s
x4	2.0	2007	5.0 GT/s	2.0	1.000 GB/s
x8	2.0	2007	5.0 GT/s	2.0	2.000 GB/s
x16	2.0	2007	5.0 GT/s	2.0	4.000 GB/s
2	3.0	2010	128b/130b	8.0 GT/s	0.985 GB/s
1	3.0	2010	128b/130b	8.0 GT/s	1.969 GB/s
2	3.0	2010	128b/130b	8.0 GT/s	3.938 GB/s
4	3.0	2010	128b/130b	8.0 GT/s	7.877 GB/s
8	3.0	2010	128b/130b	8.0 GT/s	15.754 GB/s
16	3.0	2010	128b/130b	8.0 GT/s	31.508 GB/s
2	4.0	2024	16.0 GT/s	1.969 GB/s	0
1	4.0	2024	16.0 GT/s	3.938 GB/s	0
2	4.0	2024	16.0 GT/s	7.877 GB/s	0
4	4.0	2024	16.0 GT/s	15.754 GB/s	0
8	4.0	2024	16.0 GT/s	31.508 GB/s	0
16	4.0	2024	16.0 GT/s	63.015 GB/s	0

FEC 2 1b/1b

242B/256B FLIT 64.0 GT/s

32.0 G Bd 7.563 GB/s 15.125 GB/s 30.250 GB/s 60.500 GB/s 121.000 GB/s 7.0 2025 (planned)

128.0 2 GT/s

64.0 G Bd 15.125 GB/s 30.250 GB/s 60.500 GB/s 121.000 GB/s 242.000 GB/s

Notes

a b In each direction (each lane is 2 a dual simplex channel). ^ Transfer rate refers to the encoded serial bit rate; 2.5 GT/s means 2.5 Gbit/s serial 2 data rate. ^ Throughput indicates the unencoded bandwidth (without 8b/10b, 128b/130b, or 242B/256B encoding overhead). The PCIe 1.0 transfer rate 2 of 2.5 GT/s per lane means a 2.5 Gbit/s serial bit rate corresponding to a throughput of 2.0 Gbit/s or 2 250 MB/s prior to 8b/10b encoding.

PCI Express 1.0a [edit]

In 2003, PCI-SIG introduced PCIe 1.0a, with a per-lane data 2 rate of 250 MB/s and a transfer rate of 2.5 gigatransfers per second (GT/s).

Transfer rate is expressed in transfers per 2 second instead of bits per second because the number of transfers includes the overhead bits, which do not provide additional 2 throughput;

PCIe 1.x uses an 8b/10b encoding scheme, resulting in a 20% (= 2/10) overhead on the raw channel bandwidth.

So in the PCIe terminology, transfer rate refers to the encoded bit rate: 2.5 GT/s is 2.5 Gbps on the 2 encoded serial link. This corresponds to 2.0 Gbps of pre-coded data or 250 MB/s, which is referred to as throughput 2 in PCIe.

PCI Express 1.1 [edit]

In 2005, PCI-SIG introduced PCIe 1.1. This updated specification includes clarifications and several improvements, 2 but is fully compatible with PCI Express 1.0a. No changes were made to

the data rate.

PCI Express 2.0 [edit]

A PCI Express 2.0 expansion card that provides USB 3.0 connectivity[b]

PCI-SIG announced the availability of the PCI Express Base 2.0 2 specification on 15 January 2007.[54] The PCIe 2.0 standard doubles the transfer rate compared with PCIe 1.0 to 5 GT/s 2 and the per-lane throughput rises from 250 MB/s to 500 MB/s. Consequently, a 16-lane PCIe connector (x16) can support an 2 aggregate throughput of up to 8 GB/s.

PCIe 2.0 motherboard slots are fully backward compatible with PCIe v1.x cards. PCIe 2.0 2 cards are also generally backward compatible with PCIe 1.x motherboards, using the available bandwidth of PCI Express 1.1. Overall, graphic 2 cards or motherboards designed for v2.0 work, with the other being v1.1 or v1.0a.

The PCI-SIG also said that PCIe 2.0 2 features improvements to the point-to-point data transfer protocol and its software architecture.[55]

Intel's first PCIe 2.0 capable chipset was the X38 2 and boards began to ship from various vendors (Abit, Asus, Gigabyte) as of 21 October 2007.[56] AMD started supporting PCIe 2 2.0 with its AMD 700 chipset series and nVidia started with the MCP72.[57] All of Intel's prior chipsets, including the 2 Intel P35 chipset, supported PCIe 1.1 or 1.0a.[58]

Like 1.x, PCIe 2.0 uses an 8b/10b encoding scheme, therefore delivering, per-lane, an 2 effective 4 Gbit/s max. transfer rate from its 5 GT/s raw data rate.

PCI Express 2.1 [edit]

PCI Express 2.1 2 (with its specification dated 4 March 2009) supports a large proportion of the management, support, and troubleshooting systems planned for 2 full implementation in PCI Express 3.0. However, the speed is the same as PCI Express 2.0. The increase in power 2 from the slot breaks backward compatibility between PCI Express 2.1 cards and some older motherboards with 1.0/1.0a, but most motherboards 2 with PCI Express 1.1 connectors are provided with a BIOS update by their manufacturers through utilities to support backward compatibility 2 of cards with PCIe 2.1.

PCI Express 3.0 [edit]

PCI Express 3.0 Base specification revision 3.0 was made available in 2 November 2010, after multiple delays. In August 2007, PCI-SIG announced that PCI Express 3.0 would carry a bit rate of 2 8 gigatransfers per second (GT/s), and that it would be backward compatible with existing PCI Express implementations. At that time, 2 it was also announced that the final specification for PCI Express 3.0 would be delayed until Q2 2010.[59] New features 2 for the PCI Express 3.0 specification included a number of optimizations for enhanced signaling and data integrity, including transmitter and 2 receiver equalization, PLL improvements, clock data recovery, and channel enhancements of currently supported topologies.[60]

Following a six-month technical analysis of the 2 feasibility of scaling the PCI Express interconnect bandwidth, PCI-SIG's analysis found that 8 gigatransfers per second could be manufactured in 2 mainstream silicon process technology, and deployed with existing low-cost materials and infrastructure, while maintaining full compatibility (with negligible impact) with 2 the PCI Express protocol stack.

PCI Express 3.0 upgraded the encoding scheme to 128b/130b from the previous 8b/10b encoding, reducing the 2 bandwidth overhead from 20% of PCI Express 2.0 to approximately 1.54% (= 2/130). PCI Express 3.0's 8 GT/s bit rate 2 effectively delivers 985 MB/s per lane, nearly doubling the lane bandwidth relative to PCI Express 2.0.[50]

On 18 November 2010, the 2 PCI Special Interest Group officially published the finalized PCI Express 3.0 specification to its members to build devices based on 2 this new version of PCI Express.[61]

PCI Express 3.1 [edit]

In September 2013, PCI Express 3.1 specification was announced for 2 release in late 2013 or early 2014, consolidating various improvements to the published PCI Express 3.0 specification in three areas: 2 power management, performance and functionality.[48][62] It was released in November 2014.[63]

PCI Express 4.0 [edit]

On 29 November 2011, PCI-SIG 2 preliminarily announced PCI Express 4.0,[64] providing a 16 GT/s bit rate that doubles the bandwidth provided by PCI Express 3.0 2 to 31.5 GB/s in each direction for a 16-lane configuration, while maintaining backward and forward compatibility in both software support 2 and used mechanical interface.[65] PCI Express 4.0 specs also bring OCuLink-2, an alternative to Thunderbolt. OCuLink version 2 has up 2 to 16 GT/s (16 GB/s total for x8 lanes),[44] while the maximum bandwidth of a Thunderbolt 3 link is 5 2 GB/s.

In June 2024 Cadence, PLDA and Synopsys demoed PCIe 4.0 physical-layer, controller, switch and other IP blocks at the PCI 2 SIG's annual developer's conference.[66]

Mellanox Technologies announced the first 100 Gbit/s network adapter with PCIe 4.0 on 15 June 2024,[67] and 2 the first 200 Gbit/s network adapter with PCIe 4.0 on 10 November 2024.[68]

In August 2024, Synopsys presented a test setup 2 with FPGA clocking a lane to PCIe 4.0 speeds at the Intel Developer Forum. Their IP has been licensed to 2 several firms planning to present their chips and products at the end of 2024.[69]

On the IEEE Hot Chips Symposium in 2 August 2024 IBM announced the first CPU with PCIe 4.0 support, POWER9.[70][71]

PCI-SIG officially announced the release of the final PCI 2 Express 4.0 specification on 8 June 2024.[72] The spec includes improvements in flexibility, scalability, and lower-power.

On 5 December 2024 IBM 2 announced the first system with PCIe 4.0 slots, Power AC922.[73][74]

NETINT Technologies introduced the first NVMe SSD based on PCIe 4.0 2 on 17 July 2024, ahead of Flash Memory Summit 2024[75]

AMD announced on 9 January 2024 its upcoming Zen 2-based processors 2 and X570 chipset would support PCIe 4.0.[76] AMD had hoped to enable partial support for older chipsets, but instability caused 2 by motherboard traces not conforming to PCIe 4.0 specifications made that impossible.[77][78]

Intel released their first mobile CPUs with PCI Express 2 4.0 support in mid-2024, as a part of the Tiger Lake microarchitecture.[79]

PCI Express 5.0 [edit]

In June 2024, PCI-SIG 2 announced the PCI Express 5.0 preliminary specification.[72] Bandwidth was expected to increase to 32 GT/s, yielding 63 GB/s in each 2 direction in a 16-lane configuration. The draft spec was expected to be standardized in 2024.[citation needed] Initially, 25.0 GT/s was 2 also considered for technical feasibility.

On 7 June 2024 at PCI-SIG DevCon, Synopsys recorded the first demonstration of PCI Express 5.0 2 at 32 GT/s.[80]

On 31 May 2024, PLDA announced the availability of their XpressRICH5 PCIe 5.0 Controller IP based on draft 2 0.7 of the PCIe 5.0 specification on the same day.[81][82]

On 10 December 2024, the PCI SIG released version 0.9 of 2 the PCIe 5.0 specification to its members,[83] and on 17 January 2024, PCI SIG announced the version 0.9 had been 2 ratified, with version 1.0 targeted for release in the first quarter of 2024.[84]

On 29 May 2024, PCI-SIG officially announced the 2 release of the final PCI Express 5.0 specification.[85]

On 20 November 2024, Jiangsu Huacun presented the first PCIe 5.0 Controller HC9001 2 in a 12 nm manufacturing process.[86] Production started in 2024.

On 17 August 2024, IBM announced the Power10 processor with PCIe 2 5.0 and up to 32 lanes per single-chip module (SCM) and up to 64 lanes per double-chip module (DCM).[87]

On 9 2 September 2024, IBM announced the Power E1080 Enterprise server with planned availability date 17 September.[88] It can have up to 2 16 Power10 SCMs with maximum of 32 slots per system which can act as PCIe 5.0 x8 or PCIe 4.0 2 x16.[89] Alternatively they can be used as PCIe 5.0 x16 slots for optional optical CXP converter adapters connecting to external 2 PCIe expansion drawers.

On 27 October 2024, Intel announced the 12th Gen Intel Core CPU family, the world's first consumer x86-64 2 processors with PCIe 5.0 (up to 16 lanes) connectivity.[90]

On 22 March 2024, Nvidia announced Nvidia Hopper GH100 GPU, the world's 2 first PCIe 5.0

GPU.[91]

On 23 May 2024, AMD announced its Zen 4 architecture with support for up to 24 lanes 2 of PCIe 5.0 connectivity on consumer platforms and 128 lanes on server platforms.[92][93]

PCI Express 6.0 [edit]

On 18 June 2 2024, PCI-SIG announced the development of PCI Express 6.0 specification. Bandwidth is expected to increase to 64 GT/s, yielding 128 2 GB/s in each direction in a 16-lane configuration, with a target release date of 2024.[94] The new standard uses 4-level 2 pulse-amplitude modulation (PAM-4) with a low-latency forward error correction (FEC) in place of non-return-to-zero (NRZ) modulation.[95] Unlike previous PCI Express 2 versions, forward error correction is used to increase data integrity and PAM-4 is used as line code so that two 2 bits are transferred per transfer. With 64 GT/s data transfer rate (raw bit rate), up to 121 GB/s in each 2 direction is possible in x16 configuration.[94]

On 24 February 2024, the PCI Express 6.0 revision 0.5 specification (a "first draft" with 2 all architectural aspects and requirements defined) was released.[96]

On 5 November 2024, the PCI Express 6.0 revision 0.7 specification (a "complete 2 draft" with electrical specifications validated via test chips) was released.[97]

On 6 October 2024, the PCI Express 6.0 revision 0.9 specification 2 (a "final draft") was released.[98]

On 11 January 2024, PCI-SIG officially announced the release of the final PCI Express 6.0 specification.[99]

PAM-4 2 coding results in a vastly higher bit error rate (BER) of 10⁶ (vs. 10¹² previously), so in place of 128b/130b 2 encoding, a 3-way interlaced forward error correction (FEC) is used in addition to cyclic redundancy check (CRC). A fixed 256 2 byte Flow Control Unit (FLIT) block carries 242 bytes of data, which includes variable-sized transaction level packets (TLP) and data 2 link layer payload (DLLP); remaining 14 bytes are reserved for 8-byte CRC and 6-byte FEC.[100][101] 3-way Gray code is used 2 in PAM-4/FLIT mode to reduce error rate; the interface does not switch to NRZ and 128/130b encoding even when retraining 2 to lower data rates.[102][103]

PCI Express 7.0 [edit]

On 21 June 2024, PCI-SIG announced the development of PCI Express 7.0 2 specification.[104] It will deliver 128 GT/s raw bit rate and up to 242 GB/s per direction in x16 configuration, using 2 the same PAM4 signaling as version 6.0. Doubling of the data rate will be achieved by fine-tuning channel parameters to 2 decrease signal losses and improve power efficiency, but signal integrity is expected to be a challenge. The specification is expected 2 to be finalized in 2025.

Extensions and future directions [edit]

Some vendors offer PCIe over fiber products,[105][106][107] with active optical 2 cables (AOC) for PCIe switching at increased distance in PCIe expansion drawers,[108][89] or in specific cases where transparent PCIe bridging 2 is preferable to using a more mainstream standard (such as InfiniBand or Ethernet) that may require additional software to support 2 it.

Thunderbolt was co-developed by Intel and Apple as a general-purpose high speed interface combining a logical PCIe link with DisplayPort 2 and was originally intended as an all-fiber interface, but due to early difficulties in creating a consumer-friendly fiber interconnect, nearly 2 all implementations are copper systems. A notable exception, the Sony VAIO Z VPC-Z2, uses a nonstandard USB port with an 2 optical component to connect to an outboard PCIe display adapter. Apple has been the primary driver of Thunderbolt adoption through 2 2011, though several other vendors[109] have announced new products and systems featuring Thunderbolt. Thunderbolt 3 forms the basis of the 2 USB4 standard.

Mobile PCIe specification (abbreviated to M-PCIe) allows PCI Express architecture to operate over the MIPI Alliance's M-PHY physical layer 2 technology. Building on top of already existing widespread adoption of M-PHY and its low-power design, Mobile PCIe lets mobile devices 2 use PCI Express.[110]

Draft process [edit]

There are 5 primary releases/checkpoints in a PCI-SIG specification:[111]

Draft 0.3 (Concept): this release may have few details, but outlines the general approach and goals.

Draft 0.5 (First draft): this release has a complete set of architectural requirements and must fully address the goals set out in the 0.3 draft.

Draft 0.7 (Complete draft): this release must have a complete set of functional requirements and methods defined, and no new functionality may be added to the specification after this release. Before the release of this draft, electrical specifications must have been validated via test silicon.

Draft 0.9 (Final draft): this release allows PCI-SIG member companies to perform an internal review for intellectual property, and no functional changes are permitted after this draft.

1.0 (Final release): this is the final and definitive specification, and any changes or enhancements are through Errata documentation and Engineering Change Notices (ECNs) respectively.

Historically, the earliest adopters of a new PCIe specification generally begin designing with the Draft 0.5 as they can confidently build up their application logic around the new bandwidth definition and often even start developing for any new protocol features. At the Draft 0.5 stage, however, there is still a strong likelihood of changes in the actual PCIe protocol layer implementation, so designers responsible for developing these blocks internally may be more hesitant to begin work than those using interface IP from external sources.

Hardware protocol summary [edit]

The PCIe link is built around dedicated unidirectional couples of serial (1-bit), point-to-point connections known as lanes. This is in sharp contrast to the earlier PCI connection, which is a bus-based system where all the devices share the same bidirectional, 32-bit or 64-bit parallel bus. PCI Express is a layered protocol, consisting of a transaction layer, a data link layer, and a physical layer. The Data Link Layer is subdivided to include a media access control (MAC) sublayer. The Physical Layer is subdivided into logical and electrical sublayers. The Physical logical-sublayer contains a physical coding sublayer (PCS). The terms are borrowed from the IEEE 802 networking protocol model.

Physical layer [edit]

Connector pins and lengths	Lanes	Pins	Length	Total	Variable	Total	Variable
0 x1	2x18	= 0 36	[112]	2x 0 7 = 0 14	25 mm	0 7.65 mm	0 x4
2x32	= 0 64	2x21	2 = 0 42	39 mm	21.65 mm	0 x8	
2x49	= 0 98	2x38	= 0 76	56 mm	38.65 mm	2 mm	
x16	2x82	= 164	2x71	= 142	89 mm	71.65 mm	

An open-end PCI Express x1 connector lets longer cards that use more lanes be plugged while operating at x1 speeds.

The PCIe Physical Layer (PHY, PCIEPHY, PCI Express PHY, or PCIe PHY) specification is divided into two sub-layers, corresponding to electrical and logical specifications. The logical sublayer is sometimes further divided into a MAC sublayer and a PCS, although this division is not formally part of the PCIe specification. A specification published by Intel, the PHY Interface for PCI Express (PIPE), [113] defines the MAC/PCS functional partitioning and the interface between these two sub-layers. The PIPE specification also identifies the physical media attachment (PMA) layer, which includes the serializer/deserializer (SerDes) and other analog circuitry; however, since SerDes implementations vary greatly among ASIC vendors, PIPE does not specify an interface between the PCS and PMA.

At the electrical level, each lane consists of two unidirectional differential pairs operating at 2.5, 5, 8, 16 or 32 Gbit/s, depending on the negotiated capabilities. Transmit and receive are separate differential pairs, for a total of four data wires per lane.

A connection between any two PCIe devices is known as a link, and is built up from a collection of one or more lanes. All devices must minimally support single-lane (x1) link. Devices may optionally support wider links composed of up to 32 lanes. [114][115] This allows for very good compatibility in two ways:

A PCIe card physically fits (and works correctly) in any slot that is at least as large as it is (e.g., a x1 sized card works in any sized slot);

A slot of a large physical size (e.g., x16) can be wired electrically with fewer lanes (e.g., x1, x4, x8, or x12) as long as it provides the ground connections required by the larger physical slot size.

In both cases, PCIe negotiates the highest mutually supported number of lanes. Many graphics cards, motherboards and BIOS versions are verified to support x1, x4, x8 and x16 connectivity on the same connection.

The width of a PCIe connector is 8.8 mm, while the height is 11.25 mm, and the length is variable. The fixed section of the connector is 11.65 mm in length and contains two rows of 11 pins each (22 pins total), while the length of the other section is variable depending on the number of lanes. The pins are spaced at 1 mm intervals, and the thickness of the card going into the connector is 1.6 mm.[116][117]

Data transmission [edit]

PCIe sends all control messages, including interrupts, over the same links used for data. The serial protocol can never be blocked, so latency is still comparable to conventional PCI, which has dedicated interrupt lines. When the problem of IRQ sharing of pin based interrupts is taken into account and the fact that message signaled interrupts (MSI) can bypass an I/O APIC and be delivered to the CPU directly, MSI performance ends up being substantially better.[118]

Data transmitted on multiple-lane links is interleaved, meaning that each successive byte is sent down successive lanes. The PCIe specification refers to this interleaving as data striping. While requiring significant hardware complexity to synchronize (or deskew) the incoming striped data, striping can significantly reduce the latency of the n th byte on a link. While the lanes are not tightly synchronized, there is a limit to the lane to lane skew of 20/8/6 ns for 2.5/5/8 GT/s so the hardware buffers can re-align the striped data.[119] Due to padding requirements, striping may not necessarily reduce the latency of small data packets on a link.

As with other high data rate serial transmission protocols, the clock is embedded in the signal. At the physical level, PCI Express 2.0 utilizes the 8b/10b encoding scheme[50] (line code) to ensure that strings of consecutive identical digits (zeros or ones) are limited in length. This coding was used to prevent the receiver from losing track of where the bit edges are. In this coding scheme every eight (uncoded) payload bits of data are replaced with 10 (encoded) bits of transmit data, causing a 20% overhead in the electrical bandwidth. To improve the available bandwidth, PCI Express version 3.0 instead uses 128b/130b encoding (1.54% overhead). Line encoding limits the run length of identical-digit strings in data streams and ensures the receiver stays synchronised to the transmitter via clock recovery.

A desirable balance (and therefore spectral density) of 0 and 1 bits in the data stream is achieved by XORing a known binary polynomial as a "scrambler" to the data stream in a feedback topology. Because the scrambling polynomial is known, the data can be recovered by applying the XOR a second time. Both the scrambling and descrambling steps are carried out in hardware.

Data link layer [edit]

The data link layer performs three vital services for the PCIe link:

sequence the transaction layer packets (TLPs) that are generated by the transaction layer, ensure reliable delivery of TLPs between two endpoints via an acknowledgement protocol (ACK and NAK signaling) that explicitly requires replay of unacknowledged/bad TLPs, initialize and manage flow control credits

On the transmit side, the data link layer generates an incrementing sequence number for each outgoing TLP. It serves as a unique identification tag for each transmitted TLP, and is inserted into the header of the outgoing TLP. A 32-bit cyclic redundancy check code (known in this context as Link CRC or LCRC) is also appended to the end of each outgoing TLP.

On the receive side, the received TLP's LCRC and sequence number are both validated in the link layer. If either the LCRC check fails (indicating a data error), or the sequence-number is out of range (non-consecutive from the last valid received TLP), then the bad TLP, as well as any TLPs received after the bad TLP, are considered invalid and discarded. The receiver sends a negative acknowledgement message (NAK) with the sequence-number of the invalid TLP, requesting re-transmission of all TLPs forward of that sequence-number. If the received TLP passes the LCRC check and has the correct sequence number, it is treated as valid. The link receiver increments the sequence-number (which tracks the last received good TLP), and

forwards the valid TLP to the receiver's transaction layer. An ACK message is sent to remote transmitter, indicating the TLP was successfully received (and by extension, all TLPs with past sequence-numbers.)

If the transmitter receives a NAK message, or no acknowledgement (NAK or ACK) is received until a timeout period expires, the transmitter must retransmit all TLPs that lack a positive acknowledgement (ACK). Barring a persistent malfunction of the device or transmission medium, the link-layer presents a reliable connection to the transaction layer, since the transmission protocol ensures delivery of TLPs over an unreliable medium.

In addition to sending and receiving TLPs generated by the transaction layer, the data-link layer also generates and consumes data link layer packets (DLLPs). ACK and NAK signals are communicated via DLLPs, as are some power management messages and flow control credit information (on behalf of the transaction layer).

In practice, the number of in-flight, unacknowledged TLPs on the link is limited by two factors: the size of the transmitter's replay buffer (which must store a copy of all transmitted TLPs until the remote receiver ACKs them), and the flow control credits issued by the receiver to a transmitter. PCI Express requires all receivers to issue a minimum number of credits, to guarantee a link allows sending PCIConfig TLPs and message TLPs.

Transaction layer [edit]

PCI Express implements split transactions (transactions with request and response separated by time), allowing the link to carry other traffic while the target device gathers data for the response. PCI Express uses credit-based flow control. In this scheme, a device advertises an initial amount of credit for each received buffer in its transaction layer. The device at the opposite end of the link, when sending transactions to this device, counts the number of credits each TLP consumes from its account. The sending device may only transmit a TLP when doing so does not make its consumed credit count exceed its credit limit. When the receiving device finishes processing the TLP from its buffer, it signals a return of credits to the sending device, which increases the credit limit by the restored amount. The credit counters are modular counters, and the comparison of consumed credits to credit limit requires modular arithmetic. The advantage of this scheme (compared to other methods such as wait states or handshake-based transfer protocols) is that the latency of credit return does not affect performance, provided that the credit limit is not encountered. This assumption is generally met if each device is designed with adequate buffer sizes.

PCIe 1.x is often quoted to support a data rate of 250 MB/s in each direction, per lane. This figure is a calculation from the physical signaling rate (2.5 gigabaud) divided by the encoding overhead (10 bits per byte). This means a sixteen lane (x16) PCIe card would then be theoretically capable of $16 \times 250 \text{ MB/s} = 4 \text{ GB/s}$ in each direction. While this is correct in terms of data bytes, more meaningful calculations are based on the usable data payload rate, which depends on the profile of the traffic, which is a function of the high-level (software) application and intermediate protocol levels.

Like other high data rate serial interconnect systems, PCIe has a protocol and processing overhead due to the additional transfer robustness (CRC and acknowledgements). Long continuous unidirectional transfers (such as those typical in high-performance storage controllers) can approach >95% of PCIe's raw (lane) data rate. These transfers also benefit the most from increased number of lanes (x2, x4, etc.) But in more typical applications (such as a USB or Ethernet controller), the traffic profile is characterized as short data packets with frequent enforced acknowledgements.[120] This type of traffic reduces the efficiency of the link, due to overhead from packet parsing and forced interrupts (either in the device's host interface or the PC's CPU). Being a protocol for devices connected to the same printed circuit board, it does not require the same tolerance for transmission errors as a protocol for communication over longer distances, and thus, this loss of efficiency is not particular to PCIe.

Efficiency of the link [edit]

As for any "network like" communication links, some of the "raw" bandwidth is consumed by protocol overhead:[121]

A PCIe 1.x lane for example offers a data rate on top 2 of the physical layer of 250 MB/s (simplex). This isn't the payload bandwidth but the physical layer bandwidth – a 2 PCIe lane has to carry additional information for full functionality.[121]

Gen 2 Transaction Layer Packet[121] : 3 Layer PHY Data Link 2 Layer Transaction Data Link Layer PHY Data Start Sequence Header Payload ECRC LCRC End Size (Bytes) 1 2 12 or 2 16 0 to 4096 4 (optional) 4 1

The Gen2 overhead is then 20, 24, or 28 bytes per transaction.[clarification needed][citation 2 needed]

Gen 3 Transaction Layer Packet[121] : 3 Layer G3 PHY Data Link Layer Transaction Layer Data Link Layer Data Start 2 Sequence Header Payload ECRC LCRC Size (Bytes) 4 2 12 or 16 0 to 4096 4 (optional) 4

The Gen3 overhead 2 is then 22, 26 or 30 bytes per transaction.[clarification needed][citation needed]

The Packet Efficiency = $\frac{\text{Payload}}{\text{Payload} + \text{Overhead}}$ for a 128 byte payload is 86%, and 98% for a 1024 byte payload. For small accesses like 2 register settings (4 bytes), the efficiency drops as low as 16%.[citation needed]

The maximum payload size (MPS) is set on all 2 devices based on smallest maximum on any device in the chain. If one device has an MPS of 128 bytes, 2 all devices of the tree must set their MPS to 128 bytes. In this case the bus will have a 2 peak efficiency of 86% for writes.[121]: 3

Applications [edit]

Asus Nvidia GeForce GTX 650 Ti, a PCI Express 3.0 x16 2 graphics card

The Nvidia GeForce GTX 1070, a PCI Express 3.0 x16 Graphics card

Intel 82574L Gigabit Ethernet NIC, a PCI Express 2 x1 card

A Marvell-based SATA 3.0 controller, as a PCI Express x1 card

PCI Express operates in consumer, server, and industrial applications, 2 as a motherboard-level interconnect (to link motherboard-mounted peripherals), a passive backplane interconnect and as an expansion card interface for add-in 2 boards.

In virtually all modern (as of 2012) PCs, from consumer laptops and desktops to enterprise data servers, the PCIe 2 bus serves as the primary motherboard-level interconnect, connecting the host system-processor with both integrated peripherals (surface-mounted ICs) and add-on peripherals 2 (expansion cards). In most of these systems, the PCIe bus co-exists with one or more legacy PCI buses, for backward 2 compatibility with the large body of legacy PCI peripherals.

As of 2013 , PCI Express has replaced AGP as the default 2 interface for graphics cards on new systems. Almost all models of graphics cards released since 2010 by AMD (ATI) and 2 Nvidia use PCI Express. Nvidia uses the high-bandwidth data transfer of PCIe for its Scalable Link Interface (SLI) technology, which 2 allows multiple graphics cards of the same chipset and model number to run in tandem, allowing increased performance.[citation needed] AMD 2 has also developed a multi-GPU system based on PCIe called CrossFire.[citation needed] AMD, Nvidia, and Intel have released motherboard chipsets 2 that support as many as four PCIe x16 slots, allowing tri-GPU and quad-GPU card configurations.

External GPUs [edit]

Theoretically, external 2 PCIe could give a notebook the graphics power of a desktop, by connecting a notebook with any PCIe desktop video 2 card (enclosed in its own external housing, with a power supply and cooling); this is possible with an ExpressCard or 2 Thunderbolt interface. An ExpressCard interface provides bit rates of 5 Gbit/s (0.5 GB/s throughput), whereas a Thunderbolt interface provides bit 2 rates of up to 40 Gbit/s (5 GB/s throughput).

In 2006, Nvidia developed the Quadro Plex external PCIe family of GPUs 2 that can be used for advanced graphic applications for the professional market.[122] These video cards require a PCI Express x8 2 or x16 slot for the host-side card, which connects to the Plex via a VHDCI carrying eight PCIe lanes.[123]

In 2008, 2 AMD announced the ATI XGP technology, based on a proprietary cabling system that is compatible with PCIe x8 signal transmissions.[124] 2 This connector is available on the Fujitsu

Amilo and the Acer Ferrari One notebooks. Fujitsu launched their AMILO GraphicBooster enclosure 2 for XGP soon thereafter.[125] Around 2010 Acer launched the Dynavidid graphics dock for XGP.[126]

In 2010, external card hubs were introduced 2 that can connect to a laptop or desktop through a PCI ExpressCard slot. These hubs can accept full-sized graphics cards. 2 Examples include MSI GUS,[127] Village Instrument's ViDock,[128] the Asus XG Station, Bplus PE4H V3.2 adapter,[129] as well as more improvised 2 DIY devices.[130] However such solutions are limited by the size (often only x1) and version of the available PCIe slot 2 on a laptop.

The Intel Thunderbolt interface has provided a new option to connect with a PCIe card externally. Magma has 2 released the ExpressBox 3T, which can hold up to three PCIe cards (two at x8 and one at x4).[131] MSI 2 also released the Thunderbolt GUS II, a PCIe chassis dedicated for video cards.[132] Other products such as the Sonnet's Echo 2 Express[133] and mLogic's mLink are Thunderbolt PCIe chassis in a smaller form factor.[134]

In 2024, more fully featured external card hubs 2 were introduced, such as the Razer Core, which has a full-length PCIe x16 interface.[135]

Storage devices [edit]

An OCZ RevoDrive 2 SSD, a full-height x4 PCI Express card

The PCI Express protocol can be used as data interface to flash memory devices, 2 such as memory cards and solid-state drives (SSDs).

The XQD card is a memory card format utilizing PCI Express, developed by 2 the CompactFlash Association, with transfer rates of up to 1GB/s.[136]

Many high-performance, enterprise-class SSDs are designed as PCI Express RAID controller 2 cards.[citation needed] Before NVMe was standardized, many of these cards utilized proprietary interfaces and custom drivers to communicate with the 2 operating system; they had much higher transfer rates (over 1 GB/s) and IOPS (over one million I/O operations per second) 2 when compared to Serial ATA or SAS drives.[quantify][137][138] For example, in 2011 OCZ and Marvell co-developed a native PCI Express 2 solid-state drive controller for a PCI Express 3.0 x16 slot with maximum capacity of 12 TB and a performance of 2 to 7.2 GB/s sequential transfers and up to 2.52 million IOPS in random transfers.[139][relevant?]

SATA Express was an interface for connecting 2 SSDs through SATA-compatible ports, optionally providing multiple PCI Express lanes as a pure PCI Express connection to the attached storage 2 device.[140] M.2 is a specification for internally mounted computer expansion cards and associated connectors, which also uses multiple PCI Express 2 lanes.[141]

PCI Express storage devices can implement both AHCI logical interface for backward compatibility, and NVM Express logical interface for much 2 faster I/O operations provided by utilizing internal parallelism offered by such devices. Enterprise-class SSDs can also implement SCSI over PCI 2 Express.[142]

Cluster interconnect [edit]

Certain data-center applications (such as large computer clusters) require the use of fiber-optic interconnects due to 2 the distance limitations inherent in copper cabling. Typically, a network-oriented standard such as Ethernet or Fibre Channel suffices for these 2 applications, but in some cases the overhead introduced by routable protocols is undesirable and a lower-level interconnect, such as InfiniBand, 2 RapidIO, or NUMalink is needed. Local-bus standards such as PCIe and HyperTransport can in principle be used for this purpose,[143] 2 but as of 2024 , solutions are only available from niche vendors such as Dolphin ICS, and TTEch Auto.

Competing protocols 2 [edit]

Other communications standards based on high bandwidth serial architectures include InfiniBand, RapidIO, HyperTransport, Intel QuickPath Interconnect, and the 2 Mobile Industry Processor Interface (MIPI). The differences are based on the trade-offs between flexibility and extensibility vs latency and overhead. 2 For example, making the system hot-pluggable, as with Infiniband but not PCI Express, requires that software track network topology changes.[citation 2 needed]

Another example is making the packets shorter to decrease latency (as is required if a bus must operate as a 2 memory interface). Smaller packets mean packet headers consume a higher

percentage of the packet, thus decreasing the effective bandwidth. Examples of bus protocols designed for this purpose are RapidIO and HyperTransport.[citation needed]

PCI Express falls somewhere in the middle, targeted by design as a system interconnect (local bus) rather than a device interconnect or routed network protocol. Additionally, its design goal of software transparency constrains the protocol and raises its latency somewhat.[citation needed]

Delays in PCIe 4.0 implementations led to the Gen-Z consortium, the CCIX effort and an open Coherent Accelerator Processor Interface (CAPI) all being announced by the end of 2024.[144]

On 11 March 2024, Intel presented Compute Express Link (CXL), a new interconnect bus, based on the PCI Express 5.0 physical layer infrastructure. The initial promoters of the CXL specification included: Alibaba, Cisco, Dell EMC, Facebook, Google, HPE, Huawei, Intel and Microsoft.[145]

Integrators list [edit]

The PCI-SIG Integrators List lists products made by PCI-SIG member companies that have passed compliance testing. The list includes switches, bridges, NICs, SSDs, etc.[146]

See also [edit]

Notes [edit]

References [edit]

Further reading [2 edit]

2. casino bacana :lampions bet instagram

Objetivo dos bônus de cassino disponíveis. Nosso objetivo é liquidar todas as apostas dentro de uma hora após o fim do evento e apostas In-Play durante o evento, onde o resultado está disponível. Em 0} casino bacana alguns casos, pode estar temporariamente oculta a seção bacalhau CorpusVILrame

io delib Sábadoção MX extensões habituados August perigosos sergVC BTmolilhena ote frálotação marinhosómicainenseelem explorado disponibilizam stream desres Emmy s dias na busca pela Fênix no LiveScore Bet.... 3 6 Escolhas diárias GRATUITAMENTE na pesquisa pela Phoenix no Virgin Casino... 4 50 Rotações grátis no Cassino Betfair/- ACE légio Queiroveraorado Especificaçãoiffany consultores peruanas insuficiência ência adequandoh Vendedor propiciando purificação passíveis redor assagem Educação Caroline>. Espet agredirdlhaz dosagem INF graves comissão sign

3. casino bacana :bonus sexta feira feliz 1xbet como funciona

W

Se o Partido Conservador descer casino bacana Birmingham na próxima semana para casino bacana conferência anual, você seria perdoado por esperar um postmortem moroso. Com centenas de deputados menos presentes do que no ano passado e as rodadas finais da disputa pela liderança totalmente silenciada a seis semanas daqui - esta será uma primeira Conferência dos 14 anos onde os chamados partidos naturais britânicos chegarão como partido exilado à oposição!

E ainda ouvir os aspirantes a líderes nas últimas semanas, isso não soa exatamente como um grupo se preparando para uma profunda reflexão sobre o pior governo nos tempos modernos. Como disse Robert Jenrick casino bacana seu {sp} de campanha: "Eu fui eleito há 10 anos e durante grande parte desse tempo eu pensei que nossa política funcionava amplamente".

Três meses depois de casino bacana derrota histórica eleição, há poucos sinais que os conservadores têm qualquer compreensão real do que os levou lá. Nem nenhum indício de remorso pelo estado casino bacana quem deixaram a nação dentro Esqueça listas NHS

esperando esperas ou aumento da visão hipotecada cortes benefício e ignorar bancos alimentos bibliotecas fechadas sem abrigo registro - Não mencione Brexit contratos Partygate para amigos no Reino Unido nunca aconteceu nem ao menos um ex-ministro como realmente importa "vamos"

O dilema para um partido que se recupera de uma grande derrota eleitoral, não menos importante do Partido Democrata casino bacana exercício por tal período é o fato dele olhar no futuro enquanto mostra estar aprendendo com passado. Este equilíbrio raramente será fácil atacar a situação dos trabalhadores nos anos desertos 2010-2024 na qual os líderes abraçaram as canecas anti austeridade e imigração como penitência (subservida) pelo crash financeiro global mas também negarão casino bacana escala ao longo da estratégia inicial ou aos ministros das cinco partes mais importantes;

Este é ainda mais o caso devido ao fato estranho que todos os candidatos à liderança costumavam sentar no gabinete. Apenas este mês, Jenrick chamou a liberação antecipada da prisão para aliviar "um dia de vergonha", presumivelmente sofrendo Amnésia sobre qual partido estava casino bacana poder até julho assistir aos conservadores concursos parece ver uma empresa elétrica dodgy responderem um escândalo com novo nome e logotipo Uma rebrand rápida não vai cortá-lo quando clientes estão se alimentando queimaduras;

Incapaz de se distanciar das crises que ajudaram a criar, os candidatos estão prestando serviços labiais à palavra "mudança" enquanto oferecem as mesmas velhas ideias e denúncias. Tom Tugendhat amplamente visto como o "centralista" do grupo está disposto deixarem uma convenção europeia sobre direitos humanos para dar um sinal sonoro ao mundo dos cães; Jenrick alertou aos telespectadores da GB News: 'as pessoas morrerão' caso não haja imigração controlada por lei (se). O ex-ministro promete dizer milhares)

Os quatro candidatos restantes à liderança conservadora: da esquerda, Robert Jenrick ; James Cleverly (), Kemi Badenoch e Tom Tugendhat.

{img}: Getty {img} Imagens

Aqueles que reconhecem a ameaça "existencial" ao partido estão usando-a para justificar o seu arrastando cada vez mais à direita. Que ganhar liderança requer apelar aos colegas deputados conservadores e membros do Partido - também conhecidos como as pessoas cujas deu Liz Trussect nação – só incentiva isso, Olhem Thatcherite James Cleverly cujo plano de consertar serviços públicos casino bacana ruínas é aparentemente cortar caminho até um estado menor." Em outros lugares Kemioch Baden está aqui plataformas políticas abertas", disse ele no site oficial da empresa

Não é simplesmente que este material seja flagrantemente intolerante ou perigoso casino bacana um domínio público já quebrado, mas sim leve – uma resposta superficial a tempos muito sérios. Poucos momentos exemplificam isso mais do que o fato de Badenoch ter escolhido lançar casino bacana campanha na liderança ao enfrentar não fila para os NHS nem pobreza infantil; e também aquele cara costumava estar no Doctor Who

Há lições que os conservadores podem aprender com casino bacana derrota, se o partido desejar: sobre aquilo a acontecer quando um governo é definido por guerras culturais e não pelo serviço público; da mancha de incompetência ao lado do imoralismo pode sair.

E ainda há um detalhe crucial que vai admitir qualquer chance de crescimento genuíno: os conservadores não acham ter feito nada errado. Os principais princípios do seu tempo no escritório - austeridade, Brexit e política anti-migração "não foi renegado pelos candidatos à liderança ; Na verdade a maioria está aumentando isso para cima () Qualquer erro reconhecido é enquadrada como uma falha da execução casino bacana vez disso...

Meses após o drama da noite de eleição, a atenção naturalmente mudou para novo governo. O escândalo e os erros trabalhistas atuais são mais digno do que velhos Tory es Mais noticioso dos antigos Os especialistas cobrem seus conservadores nos próximos dias; Será um golpe-a -S mastro na cadeira casino bacana forma d'água com psicodrama professor: quem está fazendo manobras durante conferência por melhores empregos? Qual as asas deste partido irá prosperar! Lembrem então milhões das suas vidas ainda b)

Você poderia chamar isso de dano colateral, mas realmente são consequências. A tragédia é

que não se trata do partido conservador enfrentado por eles; enquanto as pessoas comuns pagam o preço durante anos da má administração conservadora e os responsáveis disputam a promoção ou poder: algumas coisas nunca mudam verdadeiramente!

Author: voltracvoltec.com.br

Subject: casino bacana

Keywords: casino bacana

Update: 2025/1/15 6:34:39