



PCI-SIG FAQ

Specifications

- [PCI Express® 4.0 Specification](#)
- [PCI Express® 3.1 Specification](#)
- [PCI Express® 3.0 Specification](#)
- [PCI Express® General](#)
- [Mobile PCI Express® \(M-PCIe\) Specification](#)
- [PCI Express® M.2 Specification](#)
- [PCI Express® OCuLink Specification](#)
- [PCI Express® Form Factor SFF-8639 Specification](#)

Initiatives

- [Storage applications](#)
- [Low-power applications](#)

PCI Express® 4.0 Specification

Q: What is PCI Express® (PCIe®) 4.0? What are the requirements for this evolution of the PCIe architecture?

A: PCIe 4.0 is the next evolution of the ubiquitous and general-purpose PCI Express I/O specification. At 16GT/s bit rate, the interconnect performance bandwidth will double PCIe 3.0 specification bandwidth capabilities, while preserving backward compatibility with software and mechanical interfaces. The key requirement for evolving the PCIe architecture is to continue to provide performance scaling consistent with bandwidth demand from a variety of applications with low cost, low power and minimal perturbations at the platform level. One of the main factors in the wide adoption of the PCIe architecture is its sensitivity to high-volume manufacturing capabilities and materials such as FR4 boards and low-cost connectors. Another important requirement is the manufacturability of products using the most widely available silicon process technology.

Q: What is the release status of the PCIe 4.0 specification?

A: The final PCIe 4.0 specification, including form factor specification updates, is expected to be available in early 2017. The timing of the specification maturity is a function of the participation and contributions of PCI-SIG members as the technical work groups consider and debate technology choices and capabilities.

Q: What is the bit rate for the PCIe 4.0 specification and how does it compare to prior generations of PCIe?

A: Based on PCI-SIG's feasibility analysis, the bit rate for the PCIe 4.0 specification will be 16GT/s. This bit rate represents the optimum tradeoff between performance, manufacturability, cost, power and compatibility. PCI-SIG's analysis covered multiple topologies and confirmed the potential feasibility of 16GT/s signaling with low-cost enablers. PCI-SIG has determined that 16GT/s on copper doubles the bandwidth over the PCIe 3.0 specification and is technically feasible at approximately PCIe 3.0 power levels.

Q: How will PCIe 4.0 meet the needs of Big Data applications?

A: The PCIe 4.0 specification delivers 16GT/s bit rates, meeting demand for higher bandwidth in Big Data applications such as HPC, data centers, workstations/client platforms, embedded systems, among others. It facilitates narrower link widths and cost savings through pin reduction, while preserving backwards compatibility with all previous PCIe specifications.

Q: What are the results of the feasibility testing for the PCIe 4.0 specification?

A: The preliminary data confirmed a 16GT/s interconnect can be manufactured in mainstream silicon process technology and can be deployed with existing low-cost materials and infrastructure, while maintaining compatibility with previous generations of PCIe architecture.

Q: What were the requirements outlined for the feasibility analysis?

A: In assessing potential improvements to the connector, materials, silicon and channel improvements, PCI-SIG required compatibility, low-cost and high-volume manufacturing to be maintained.

Q: Will PCIe 4.0 products be compatible with existing PCIe 1.x, PCIe 2.x and PCIe 3.x products?

A: PCI-SIG is proud of its long heritage of developing compatible architectures and its members have consistently producing compatible and interoperable products. In keeping with this tradition, the PCIe 4.0 architecture will be compatible with prior generations of this technology, from software to clocking architecture to mechanical interfaces. That is to say PCIe 1.x, 2.x and 3.x cards will seamlessly plug into PCIe 4.0-enabled slots and operate at the highest possible performance levels. Similarly, all PCIe 4.0 cards will plug into PCIe 1.x-, PCIe 2.x- and PCIe 3.x-capable slots and operate at the highest performance levels supported by those configurations.

Q: Why is a new generation of PCIe architecture needed?

A: PCI-SIG responds to the needs of its members and the industry at large. Since inception, the PCI Express architecture has been designed to support adoption in multiple applications, from SoCs to high-performance servers, even as performance and feature requirements have evolved in newer product generations. In today's technology landscape that requires low power, high performance I/O, PCIe 4.0 will offer 16GT/s

Q: What are the initial target applications for the PCIe 4.0 architecture?

A: The PCIe 4.0 specification will address applications requiring increased bandwidth at a lower cost including server, workstation, desktop PC, notebook PC, tablets, embedded systems, peripheral devices, high-performance computing markets and more. The target implementations are entirely at the discretion of the designer.

Q: Is PCIe 4.0 architecture more expensive to implement than PCIe 3.x?

A: PCI-SIG attempts to define and evolve the PCIe architecture in a manner consistent with low-cost and high-volume manufacturability considerations. While PCI-SIG cannot comment on design choices and implementation costs, optimized silicon die size and power consumption continue to be important considerations that inform PCIe specification development and architecture evolution.

Q: Will there be a new compliance specification developed for the PCIe 4.0 specification?

A: PCI-SIG offers an industry-leading and robust compliance program to help ensure seamless and consistent interoperability for plug-and-play across multiple platforms. For each revision of its specification, PCI-SIG develops compliance tests and related collateral consistent with the requirements of the new architecture. These compliance requirements are incremental in nature and build on the prior generation of the architecture. PCI-SIG anticipates releasing compliance specifications as they mature along with corresponding tests and measurement criteria. Each revision of the PCIe architecture maintains its own criteria for product interoperability and admission into the PCI-SIG Integrators List.

PCI Express® 3.1 Specification

Q: What is the PCIe® 3.1 specification?

A: The PCIe 3.1 specification adds seven PCIe 3.0 ECNs to the PCIe 3.0 base specification. This specification includes power, performance and functionality ECNs developed since the release of the original PCIe 3.0 base specification. ECNs included are:

- M-Pcie
- L1 Power management Substates with CLKREQ#
- Enhanced Downstream Port Containment
- Lightweight Notification
- Precision Time Measurement
- Separate Refclk Independent SSC
- Process Address Space ID

PCI Express® 3.0 Specification

Q: What categories of products can be tested for PCIe 3.0 compliance?

A: PCIe 3.0 products including components (endpoints, switches and bridges, root complex), add-in cards and PC-AT compatible motherboards/systems can be certified with the PCIe 3.0 compliance program.

Q: What is PCI Express® (PCIe®) 3.0?

A: The PCIe 3.0 specification offers up to an 8GT/s bit rate and interconnect performance bandwidth that is doubled over PCIe 2.0. This bit rate represents the most optimum tradeoff between manufacturability, cost, power and compatibility. In providing full compatibility, the same topologies and channel reach as in PCIe 2.0 are supported for both client and server configurations. For the PCIe 3.0 architecture, PCI-SIG® believes a 65nm process or better will be required to optimize on silicon area and power.

Q: How does the PCIe 3.0 8GT/s “double” the PCIe 2.0 5GT/s bit rate?

A: The PCIe 2.0 bit rate is specified at 5GT/s, but with the 20 percent performance overhead of the 8b/10b encoding scheme, the delivered bandwidth is actually 4Gbps. PCIe 3.0 removes the requirement for 8b/10b encoding and uses a more efficient 128b/130b encoding scheme instead. By removing this overhead, the interconnect bandwidth can be doubled to 8Gbps with the PCIe 3.0 specification. This bandwidth is the same as an interconnect running at 10GT/s with the 8b/10b encoding overhead. That said, the PCIe 3.0 specification delivers the same effective bandwidth without the prohibitive penalties associated with 10GT/s signaling, such as PHY design complexity and increased silicon die size and power.

The following table summarizes the bit rate and approximate bandwidths for the various generations of the PCIe architecture:

PCIe architecture	Raw bit rate	Interconnect/link bandwidth	Bandwidth per lane per direction	Total bandwidth for x16 link
PCIe 1.x	2.5GT/s	2Gbps	~250MB/s	~8GB/s
PCIe 2.x	5.0GT/s	4Gbps	~500MB/s	~16GB/s
PCIe 3.0	8.0GT/s	8Gbps	~1GB/s	~32GB/s
PCIe 4.0	16.0GT/s	16GB/s	~2GB/s	~64GB/s

Total bandwidth represents the aggregate interconnect bandwidth in both directions.

Q: Does the PCIe 3.0 specification only deliver a signaling rate increase?

A: The PCIe 3.0 specification is comprised of the Base and the Card Electro-mechanical (CEM) specifications. Within the PCIe 3.0 Base specification, which defines a chip-to-chip interface, updates have been made to the CEM specification section to comprehend 8GT/s signaling. As the technology definition progresses through PCI-SIG specification development process, additional ECNs and errata will be incorporated with each review cycle. For example, the PCIe protocol extensions that address interconnect latency and other platform resource usage considerations have been rolled into the PCIe 3.0 specification revisions. The final PCIe 3.0 specification consolidated all ECN and errata published since the release of the PCIe 2.1 specification, as well as interim errata.

Q: Are PCIe 3.0 products compatible with existing PCIe 1.x and PCIe 2.x products?

A: In keeping with its tradition of developing compatible architectures, PCI-SIG designed the PCIe 3.0 architecture to be fully compatible with prior generations of this technology. That is to say PCIe 1.x and 2.x cards will seamlessly plug into PCIe 3.0-capable slots and operate at their highest performance levels. Similarly, all PCIe 3.0 cards will plug into PCIe 1.x- and PCIe 2.x-capable slots and operate at the highest performance levels supported by those configurations.

Q: When was the PCIe 3.0 specification made available?

A: PCI-SIG released the PCIe 3.0 specification on November 17, 2010.

Q: What is 8b/10b encoding?

A: 8b/10b encoding is a byte-oriented coding scheme that maps each byte of data into a 10-bit symbol. It guarantees a deterministic DC wander and a minimum edge density over a per-bit time continuum. These two characteristics permit AC coupling and a relaxed clock data recovery implementation. Since each byte of data is encoded as a 10-bit quantity, this encoding scheme guarantees that in a multi-lane system, there are no bubbles introduced in the lane striping process.

Q: What is scrambling? How does scrambling impact the PCIe 3.0 architecture?

A: Scrambling is a technique where a known binary polynomial is applied to a data stream in a feedback topology. Because the scrambling polynomial is known, the data can be recovered by running it through a feedback topology using the inverse polynomial. Scrambling affects the PCIe architecture at two levels: the PHY layer and the protocol layer immediately above the PHY. At the PHY layer, scrambling introduces more DC wander than an encoding scheme such as 8b/10b; therefore, the Rx circuit must either tolerate the DC wander as margin degradation or implement a DC wander correction capability. Scrambling does not guarantee a transition density over a small number of unit intervals, only over a large number. The Rx clock data recovery circuitry must be designed to remain locked to the relative position of the last data edge in the absence of subsequent edges. At the

protocol layer, an encoding scheme such as 8b/10b provides out-of-band control characters that are used to identify the start and end of packets. Without an encoding scheme (i.e. scrambling only) no such characters exist, so an alternative means of delineating the start and end of packets is required. Usually this takes the form of packet length counters in the Tx and Rx and the use of escape sequences. The choice for the scrambling polynomial is currently under study.

Q: What are the target applications for PCIe 3.0?

A: PCIe 3.0 technology has been broadly adopted in graphics cards, SSDs, network adapters, storage devices and switches. In the future, other applications may put additional bandwidth and performance demands on the PCIe architecture.

Q: Does PCIe 3.0 enable greater power delivery to cards?

A: The PCIe Card Electro-mechanical (CEM) 3.0 specification consolidates all previous form factor power delivery specifications, including the 150W and the 300W specifications.

Q: Is PCIe 3.0 more expensive to implement than PCIe 2.x?

A: PCI-SIG attempts to define and evolve the PCIe architecture in a manner consistent with low-cost and high-volume manufacturability considerations. While PCI-SIG cannot comment on design choices and implementation costs, optimized silicon die size and power consumption continue to be overarching imperatives that inform PCIe specification development and architecture evolution.

PCI Express® (General)

Q: What are the PCIe protocol extensions, and how do they improve PCIe interconnect performance?

A: The PCIe protocol extensions are primarily intended to improve interconnect latency, power and platform efficiency. These protocol extensions pave the way for better access to platform resources by various compute and I/O intensive applications as they interact with and through the PCIe interconnect hierarchy. There are multiple protocol extensions and enhancements being developed and they range in scope from data reuse hints, atomic operations, dynamic power adjustment mechanisms, loose transaction ordering, I/O page faults, BAR resizing and so on. Together, these protocol extensions will increase PCIe deployment leadership in emerging and future platform I/O usage models by enabling significant platform efficiencies and performance advantages.

Q: What is equalization? How is Tx equalization different from Rx equalization? What is trainable equalization?

A: Equalization is a method of distorting the data signal with a transform representing an approximate inverse of the channel response. It may be applied either at the Tx, the Rx, or both. A simple form of equalization is Tx de-emphasis as specified in PCIe 1.x and PCIe 2.x, where data is sent at full swing after each polarity transition and is sent at reduced swing for all bits of the same polarity thereafter. De-emphasis reduces the low frequency energy seen by the Rx. Since channels exhibit greater loss at high frequencies, the effect of equalization is to reduce these effects. Equalization may also be used to compensate for ripples in the channel that occur due to reflections from impedance discontinuities such as vias or connectors. Equalization may be implemented using various types of algorithms; the two most common are linear (LE) and decision feedback (DFE). Linear equalization may be implemented at the Tx or the Rx, while DFE is implemented at the Rx. Trainable equalization refers to the ability to adjust the tap coefficients. Each combination of Tx, channel, and Rx will have a unique set of coefficients yielding an optimum signal-to-noise

ratio. The training sequence consists of adjustments to the tap coefficients while applying a quality metric to minimize the error. The choice for the type of equalization to require in the next revision of the PCIe specifications depends largely on the interconnect channel optimizations that can be derived at the lowest cost point. It is the intent of PCI-SIG to deliver the most optimum combination of channel and silicon enhancements at the lowest cost for the most common topologies

Mobile PCIe (M-PCIe™) Specification

Q: What is the M-PCIe specification?

A: The M-PCIe specification enables the PCIe architecture to operate over the MIPI Alliance M-PHY physical layer, thereby extending the benefits of the PCIe I/O standard to mobile devices including laptops, tablets and smartphones. This layered architecture of PCIe I/O technology facilitates the integration of power-efficient M-PHY with its extensible protocol stack.

Q: Why did PCI-SIG adapt PCIe protocols to operate over the MIPI® Alliance M-PHY® specification?

A: As PCs become lighter and thinner, and tablets and smartphones become more functional, consumers want seamless, always on/always connected functionality from their computing devices. To respond to these market expectations, device manufacturers need efficient, intelligent I/O technologies. The PCIe architecture satisfies all of these requirements, and with the adaptation to operate over the MIPI Alliance's M-PHY, a mobile physical layer specification, it can deliver consistent high performance in power-constrained platforms such as laptops, tablets and smartphones. By delivering this technology, the PCI-SIG is meeting the emerging needs of its members and the industry.

Q: Is M-PCIe significantly better than PCIe for mobile applications?

A: In addition to the link active power improvements with the MIPI M-PHY, the Mobile PCIe (M-PCIe) specification provides a number of other benefits for mobile devices. M-PCIe defines a simplified link management architecture (e.g. Link Training Status State Machine [LTSSM] and related logic) which can greatly lower development and validation costs. M-PCIe allows for asymmetric link width design for optimization of power-sensitive devices. In a multi-lane topology where more data flows in one direction than the other, asymmetric links yield better performance with lower power consumption. Also, the M-PHY provides a flexible range of data rates that can be tuned to the bandwidth needs of the application.

Q: What is the primary benefit of using PCIe architecture with the MIPI M-PHY?

A: By reducing I/O technology proliferation and by maximizing the reuse of existing IP, component and device manufacturers can recoup their investments faster; drastically reduce the time for product development and validation; flexibly comprehend high-performance applications for optimum user experience and hasten the delivery of innovative solutions to the market. The mobile and handset industry can realize these benefits by adopting PCIe architecture adapted to run over M-PHY.

Q: What are the main applications of the PCIe adaptation on the MIPI M-PHY?

A: The initial application is expected to be high-performance wireless LAN (WiGig). In addition, it is expected that this technology will be adopted in future storage applications in various topologies due to the anticipated migration of storage attach points from SATA to PCIe technology. As a power-efficient, general-purpose load-store I/O architecture, component and device designers can implement this technology in other I/O expansion

usage models of their choosing as well. It is anticipated that mobile platform developers will incorporate this specification into thin laptops, tablets and cellular phones.

Q: Is there a need for new software to support the PCIe adaptation on the MIPI M-PHY?

A: The adaptation of the PCIe architecture on the MIPI M-PHY requires no new software. It reuses the existing, ubiquitous support in all major Operating Systems (e.g. pci.sys bus driver on Windows platforms). This includes existing support for device discovery, configuration and control.

PCI Express® M.2 Form Factor Specification

Q: What is the PCI Express® M.2 specification?

A: The PCI Express® M.2 specification is a form factor designed for ultra-light and thin mobile platforms. The specification is a natural transition from the Mini Card and Half-Mini Card to a smaller form factor in both size and volume. This specification enables a variety of devices with multiple sockets, keys, and form factors (modules and BGAs/LGAs). In the future, the M.2 specification will incorporate a new definition for BGA SSDs. The new definition will give system manufacturers the flexibility in form factor choice for their solid state storage solutions. Ultimately, this will expand the potential product offerings for SSD manufacturers and give them the option to populate M.2 modules with the BGA device.

Q: Why was the M.2 specification developed and how did it evolve from the NGFF (Next Generation Form Factor)?

A: NGFF began as a mobile module and socket definition effort outside of PCI-SIG. Later it was brought under the auspices of PCI-SIG and renamed as M.2. PCI-SIG owns the preceding standard, the PCI Express Mini Card, sometimes referred to as Mini CEM, which primarily focuses on communication devices. The Serial ATA International Organization (SATA-IO) leveraged this form factor and definition to create its mSATA specification. With NGFF, later renamed as M.2, there was a golden opportunity to define a single specification which would serve the needs of both communication devices and solid state storage devices in a modular form factor. An MOU was established between two standards bodies: PCI-SIG and SATA-IO enabling members from both organizations to collaborate on the specification development. The new specification name, M.2, was included in the list of proposed names presented by the PCI Express Mini Working Group (the group developing the specification) to the PCI-SIG Board of Directors. It was selected as the preferred name of the specification, yet it carries no specific meaning.

Q. What are the different lengths of M.2 storage devices and how do storage developers take advantage of the extra space?

A: M.2 defines 42mm, 80mm and 110mm form factors primarily for design flexibility and I/O expansion through a mobile-friendly connector. Client platforms, along with server and micro-server platforms can take advantage of the flexible M.2 form factor for I/O expansion. Storage device developers can also utilize the space for higher capacity modules or for design flexibility in a single-sided longer module vs. a double-sided shorter module.

Q: What types of usage models is the PCI Express® M.2 specification designed for?

A: The specification has the flexibility to support high-end and scalable performance to power and space-constrained platforms, such as laptops, tablets and smartphones.

Q: What formats does the M.2 form factor come in?

A: The M.2 form factor comes in a connector or a soldered-down format. The connector format has single-sided modules for low profile solutions or dual-sided modules for increased integration, while the soldered-down format has single-sided for use in low profile applications.

Q: Does the PCIe® M.2 specification support other small form factor technologies?

A: Yes, the specification is designed to include support for multiple technologies, including but not limited to Wi-Fi®, Bluetooth®, SSD and WWAN.

Q: When did the PCIe® M.2 specification become available?

A: Revision 1.0 of the PCIe M.2 specification was released to PCI-SIG members in December 2013.

PCI Express® OCuLink Specification

Q: What is the PCI Express OCuLink specification?

A: The PCI Express OCuLink specification is focused on small, low-cost cable solutions targeting form factor client and mobile external enclosures and internal enclosure usage models. The internal usage targets PCIe-attached storage such as storage systems and PCIe add-in cards, while external usage extends PCIe I/O outside of the box for external attached storage.

Q: How many lanes and at what speed can an OCuLink cable support?

A: OCuLink has a basic transfer rate of 8 GT/s with room to scale higher. One external and one internal OCuLink connector can support four PCIe lanes with up to 32 Gbps in each direction.

Q: What is the OCuLink cable made from?

A: OCuLink, meaning “optical copper link,” has cables made from copper; however, there are plans to offer optical cables in the future.

Q: What is the release status of the OCuLink specification?

A: Revision 0.9 of the OCuLink specification is currently under review. Revision 1.0 is expected to be published in Q4 2015.

Q: What is the relationship between the PCI Express OCuLink specification and the PCI Express External Cabling specification?

A: The PCI Express OCuLink specification is intended to augment the PCI Express specification suite, targeting specific usage models that the PCI Express External Cabling specification does not already address. The two specifications will continue to evolve independently as members and the industry require.

PCI Express® SFF-8639 Form Factor Specification

Q: What is SFF-8639?

A: SFF-8639 is a new connector designed for PCIe attached storage devices by linking SSDs to a backplane in a single server or rack-mounted server cluster. SFF-8639 functions as a server module for high-density SSD storage attach, supporting up to 8.0 GT/s data rate per direction and allowing SFF-8639 devices and existing storage products to co-exist.

Q: What are the benefits for using a SFF-8639 form factor?

A: It can enable hot plug and hot swap (optional) and maximize module interoperability for end-users. It also sports an extensible design to meet future bandwidth needs.

Q: When was the SFF-8639 form factor specification released?

A: A 0.7 version of the specification was released in March 2014.

Q: What voltage rails does the SFF-8639 connector support?

A: The SFF-8639 connector supports +3.3 Vaux (not required for all platforms), +5V and +12V.

Q: What are the optional and required auxiliary signals supported by the SFF-8639 module connector?

A: Required:

- REFCLK-/REFCLK+, WAKE#, PERST#, PRSNT#, IFDET#, ACTIVITY#, reference clock (REFCLKB), and PERSTB#

Optional:

- SMBCLK, SMBDAT, DUALLINKEN# and CLKREQ#

Q: What are the power delivery requirements for SFF-8639 modules bay and system?

A: All SFF-8639 bays require one power rail: +12V, with a second, optional 3.3 Vaux rail. Systems that provide SFF-8639 bays are required to provide the +12V rail to every SFF-8639 bay in the system.

Storage

Q: What capabilities does PCIe offer for storage?

A: PCIe is a prominent attach point for storage and networking devices. PCIe storage is embodied in NVMe specification for PCIe-SSD interconnects and creates a hot-pluggable backbone for high-density SSD storage attachments.

Q: Will PCIe continue to be a preferred storage I/O interconnect?

A: PCI-SIG firmly believes PCIe technology will continue to outpace other interconnect technologies in both units and bandwidth/capacity. With the data explosion driving SSD adoption, there are strong projections that the PCIe architecture will continue to provide unbeatable performance and low power features for storage devices.

Q: How can PCIe technology be applied in storage applications?

A: PCI-SIG offers several form factors that deliver connectivity and expansion for storage applications, such as:

- SFF-8639 server module for high-density SSD storage attach
- M.2 form factor specification
- OCuLink cable form factor

Low-Power

Q: Is the PCIe architecture able to support low power and battery-based applications?

A: Yes, PCIe specifications can help reduce power consumption in smaller embedded, handheld and mobile device platforms. The PCIe architecture offers flexible lane width

configurations and speed selection supporting low power solutions for desktops, servers and storage solutions.

Q: How does the PCIe architecture support low-power states?

A: The PCIe architecture provides L1 sub-states to improve energy distribution on platforms, reduce microwatts and to lower power in idle mode. Additionally, half-swing state which operates at 400 mW was introduced in the PCIe 1.0 base specification. The upcoming PCIe 4.0 specification will cut power usage even further with quarter-swing state at 200 mW.