PCI Express® M.2 Specification

Q: What is the PCI Express® M.2 Specification?
A: The PCI Express® M.2 Specification is a next generation form factor for ultra-light and thin 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.

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 performance and provide scalable performance to power-constrained platforms, such as tablets and smartphones. The specification is designed as a tunable I/O technology for developers to create the perfect balance of power and performance.

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 will the PCIe® M.2 Specification be available?
A: Revision 1.0 of the PCIe M.2 specification was released to PCI-SIG members in December 2013.

M-PCIe™ Specification

Q: Why is PCI-SIG adapting 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 M-PHY specification it can deliver consistent high performance in power-constrained platforms such as ULT 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 M-PCIe technology 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, can 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 today 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: This adaptation of the PCIe architecture 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® 3.0 Specification

Q: When was the PCIe 3.0 Integrators List published?
A: The first PCIe 3.0 member company products to pass PCIe 3.0 compliance testing were published on May 17, 2013. The compliance testing was done at the PCI-SIG Workshop #85.

Q: What categories of products can be tested for PCI-SIG 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 PCI-SIG PCIe 3.0 compliance program.

Q: What is PCI Express® (PCIe®) 3.0? What are the requirements for this evolution of the PCIe architecture?
A: PCIe 3.0 is the next evolution of the ubiquitous and general-purpose PCI Express I/O standard. At 8GT/s bit rate, the interconnect performance bandwidth is doubled over PCIe 2.0, while preserving 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 leading 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 materials and tolerances such as FR4 boards, low-cost clock sources, connectors and so on. In providing full compatibility, the same topologies and channel reach as in PCIe 2.0 are supported for both client and server configurations. Another important requirement is the manufacturability of products using the most widely available silicon process technology. 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: What is the bit rate for PCIe 3.0 and how does it compare to prior generations of PCIe?
A: The bit rate for PCIe 3.0 is 8GT/s. This bit rate represents the most optimum tradeoff between manufacturability, cost, power and compatibility. PCI-SIG analysis covered multiple topologies and configurations, including servers. All of these studies confirmed the feasibility of 8GT/s signaling with low-cost enablers and with minimal increases in power, silicon die size and complexity.

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 implementation of the PCIe 3.0 specification. This bandwidth is the same as an interconnect running at 10GT/s with the 8b/10b encoding overhead. In this way, the PCIe 3.0 specifications deliver the same effective bandwidth, but 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 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

Total bandwidth represents the aggregate interconnect bandwidth in both directions.

Q: Do PCIe 3.0 specifications only deliver a signaling rate increase?
A: The PCIe 3.0 specifications comprise the Base and the Card Electro-mechanical (CEM) specifications. There may be updates to other form factor specifications as the need arises. Within the Base specification, which defines a chip-to-chip interface, updates have been made to the electrical section to comprehend 8GT/s signaling. As the technology definition progresses through PCI-SIG specification development process, additional ECN and errata will be incorporated with each review cycle. For example, the current 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 consolidates all ECN and errata published since the release of the PCIe 2.1 specification, as well as interim errata.

Q: Will PCIe 3.0 products be compatible with existing PCIe 1.x and PCIe 2.x products?
A: PCI-SIG is proud of its long heritage of developing compatible architectures and its members have consistently produced compatible and interoperable products. In keeping with this tradition, the PCIe 3.0 architecture is fully compatible with prior generations of this technology, from software to clocking architecture to mechanical interfaces. 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. The following chart summarizes the interoperability between various generations of PCIe and the resultant interconnect performance level:

Transmitter Device	Receiver Device	Channel	Interconnect Data Rate
8GHz	8GHz	8GHz	8.0GT/s
5GHz	5GHz	5GHz	5.0GT/s
2.5GHz	2.5GHz	2.5GHz	2.5GT/s
8GHz	5GHz	8GHz	5.0GT/s
8GHz	2.5GHz	8GHz	2.5GT/s
5GHz	2.5GHz	5GHz	2.5GT/s
…	…	…	…

In short, the notion of the compatible highest performance level is modeled after the mathematical least common denominator (LCD) concept. Also, PCIe 3.0 products will need to support 8b/10b encoding when operating in a pre-PCIe 3.0 environment.

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: 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 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

Q: What are the target applications for PCIe 3.0?
A: It is expected that graphics, Ethernet, InfiniBand, storage and PCIe switches will continue to drive the bandwidth evolution for the PCIe architecture and these applications are the current targets of the PCIe 3.0 technology. 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® 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. The new specification is intended to update PCIe 3.0 with 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® 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 be doubled over the PCIe 3.0 specification, 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, low-cost connectors and so on.

Q: What is the release status of the PCIe 4.0 specification?
A: The final PCIe 4.0 specifications, including form factor specification updates, are expected to be available in late 2016/early 2017. The timing of the specification maturity is a function of the participation and contributions of PCI-SIG members as the technical workgroups consider and debate technology choices, capabilities and analyses.

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 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 analysis covered multiple topologies. All of these studies confirmed the potential feasibility of 16GT/s signaling with low-cost enablers.

Q: What are the results of the feasibility testing for the PCIe 4.0 specification?
A: After technical analysis, the PCI-SIG has determined that 16 GT/s on copper, which will double the bandwidth over the PCIe 3.0 specification, is technically feasible at approximately PCIe 3.0 power levels. The preliminary data also confirms that 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. In addition, the PCI-SIG will investigate advancements in active and idle power optimizations as they become available.

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 that compatibility, low-cost and high-volume manufacturing 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 produced compatible and interoperable products. In keeping with this tradition, the PCIe 4.0 architecture is 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-capable slots and operate at the highest performance levels possible. 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. As applications evolve to consume the I/O bandwidth provided by the current generation of the PCIe architecture, PCI-SIG begins to study the requirements for technology evolution to keep abreast of performance and feature requirements.

Q: What are the initial target applications for the PCIe 4.0 architecture?
A: The PCIe 4.0 specification will address the many applications pushing for increased bandwidth at a low 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 been a new compliance specification developed for the PCIe 4.0 specification?
A: For each revision of its specification, PCI-SIG develops compliance tests and related collateral consistent with the requirements of the new architecture. All of 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 technology maintains its own criteria for product interoperability and admission into the PCI-SIG Integrators List.