



PCI Express® 3.0 Frequently Asked Questions PCI-SIG

1) What is PCI Express® (PCIe) 3.0? What are the requirements for this evolution of the PCIe® architecture?

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 will be 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, the PCI-SIG believes a 65nm process or better will be required to optimize on silicon area and power.

2) What is the bit rate for PCIe 3.0 and how does it compare to prior generations of PCIe?

The bit rate for PCIe 3.0 is 8GT/s. This bit rate represents the most optimum tradeoff between manufacturability, cost, power and compatibility.

The 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 and silicon die size.

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

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 4Gb/s. 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 8Gb/s 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	2Gb/s	~250MB/s	~8GB/s
PCIe 2.x	5.0GT/s	4Gb/s	~500MB/s	~16GB/s
PCIe 3.0	8.0GT/s	8Gb/s	~1GB/s	~32GB/s

Total bandwidth represents the aggregate interconnect bandwidth in both directions.

4) Does this mean that PCIe is finished at 8GT/s? What comes next?

The PCI-SIG will study the requirements of its members and of the industry for the next generation of the PCIe architecture following the successful release of the PCIe 3.0 specifications. Higher signaling rates depend on a number of factors. The PCI-SIG is committed to delivering the most robust and high-performance I/O interconnect specifications, while at the same time maintaining an uncompromised focus on low cost, low power, high volume manufacturability and compatibility, by taking advantage of breakthroughs in signaling technologies and silicon process capabilities.

5) Will PCIe 3.0 specifications only deliver a signaling rate increase?

The PCIe 3.0 specifications will comprise the Base and the Card Electromechanical (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 will be made to the electrical section to comprehend 8GT/s signaling. As the technology definition progresses through the 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 will be rolled into the PCIe 3.0 specification revisions as they become available. The final PCIe 3.0 specification will consolidate all ECN and errata published since the release of the PCIe 2.1 specification, as well as interim errata.

6) Will PCIe 3.0 products be compatible with existing PCIe 1.x and PCIe 2.x products?

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 will be 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.

7) What are the PCIe protocol extensions, and how do they improve PCIe interconnect performance?

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.

8) When will PCIe 3.0 specifications be available?

The PCI-SIG released the PCIe 3.0 specification 0.9 revision for member review in August 2010. It is anticipated that the final specification will become available in November 2010.

9) What timeframe are PCIe 3.0 products expected in the marketplace?

The PCI-SIG does not comment on member products. However, based on architecture updates and PCI-SIG enabling activities, it is estimated that PCIe 3.0 products may emerge two to three quarters following the completion of the specifications.

10) What is 8b/10b encoding?

8b/10b encoding is a byte-oriented coding scheme that maps each byte of data into one or two 10-bit characters. 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.

11) What is scrambling? How does scrambling impact the PCIe 3.0 architecture?

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.

12) What is equalization? How is Tx equalization different from Rx equalization? What is trainable equalization?

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 the PCI-SIG to deliver the most optimum combination of channel and silicon enhancements at the lowest cost for the most common topologies.

13) Why is a new generation of PCIe needed?

The 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, the PCI-SIG begins to study the requirements for technology evolution to keep abreast of these performance and feature requirements.

14) What are the initial target applications for PCIe 3.0?

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.

15) Will PCIe 3.0 enable greater power delivery to cards?

The PCIe Card Electromechanical (CEM) 3.0 specification will consolidate all previous form factor power delivery specifications, including the 150W and the 300W specifications.

16) Will PCIe 3.0 be more expensive to implement than PCIe 2.x?

The PCI-SIG attempts to define and evolve the PCIe architecture in a manner consistent with low-cost and high-volume manufacturability considerations. While the 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.

17) Will there be a new compliance specification developed for PCIe 3.0?

For each revision of its specification, the 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. The 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 compliance Integrators List.