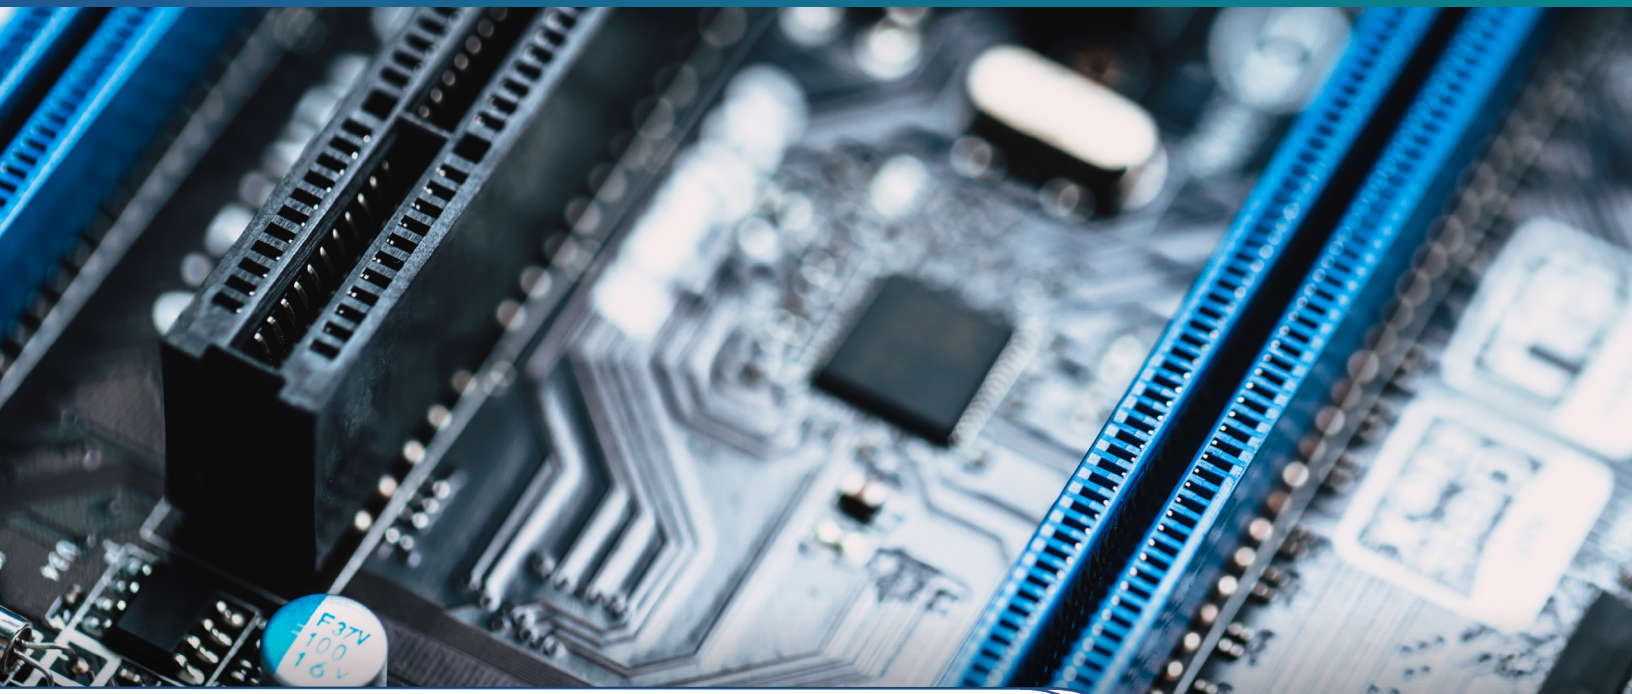

WHITE PAPER

Evolution of PCIe® Standards and Test Requirements

Designers of data center systems and devices face ever-increasing demands to provide faster speeds. Ethernet network interfaces in data centers are moving to speeds of 400 gigabits per second (Gb/s) and will likely double in the coming years. Peripheral Component Interconnect Express (PCIe) 5.0, the latest iteration of the standard, will enable the mass adoption of 400 gigabit Ethernet (GE) technologies in the data center.

PCIe is a core technology used in many types of computer servers and endpoint devices. The PCI Special Interest Group (PCI-SIG®) defines specifications and compliance tests that guarantee the interoperability of PCIe systems. PCIe 5.0 provides full-duplex bandwidth of approximately 128 gigabytes per second (GB/s) for a 16-lane system.

The PCIe standard has evolved from PCIe 1.0, released in 2003 supporting 2.5 gigatransfers per second (GT/s), to PCIe 5.0, released in 2019 supporting 32 GT/s.

PCIe is scalable, and slots come in different configurations of bidirectional lanes: x1, x4, x8, x16, x32. The number represents the number of lanes in the PCIe slot. For example, a PCIe x1 slot provides one lane and transmits data at 1 bit per cycle. A PCIe x2 slot provides two lanes and transmits data at 2 bits per cycle, and so on. PCIe cards fit interchangeably into slots, but the bandwidth available depends upon the version of the PCIe standard of the card.

KEYSIGHT TECHNOLOGIES

Figure 1 shows the evolution of the standard and the bandwidth doubling with each generation.

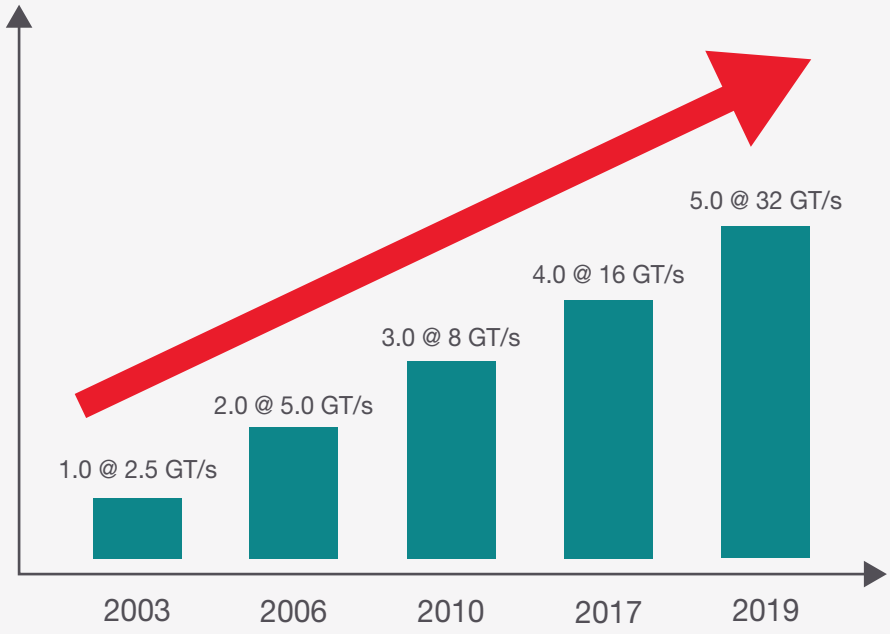

Figure 1: Evolution of the PCIe standard

As a serial point-to-point, multilane interconnect between two devices, PCIe communicates directly with devices via a switch that directs data flow. As a result, it is possible to perform "hot swapping" or "hot plugging" of cards in PCIe slots without shutting down the system. PCI-SIG requires vendors to conduct compliance testing for interoperability that include a list of preset or transmitter equalization settings which must be electrically verified. These presets ensure optimal compensation for channel loss characteristics and optimize signal integrity on the link. There are three categories of testing required to reach official PCI-SIG approved compliance: the physical layer, the data link layer, and general interoperability.

Different types of test equipment must perform different categories of tests:

1. Oscilloscopes validate physical-layer transmitter (Tx) test.
2. Bit error ratio testers (BERTs) validate physical-layer receiver (Rx) test.
3. Protocol analyzers validate data link layers.

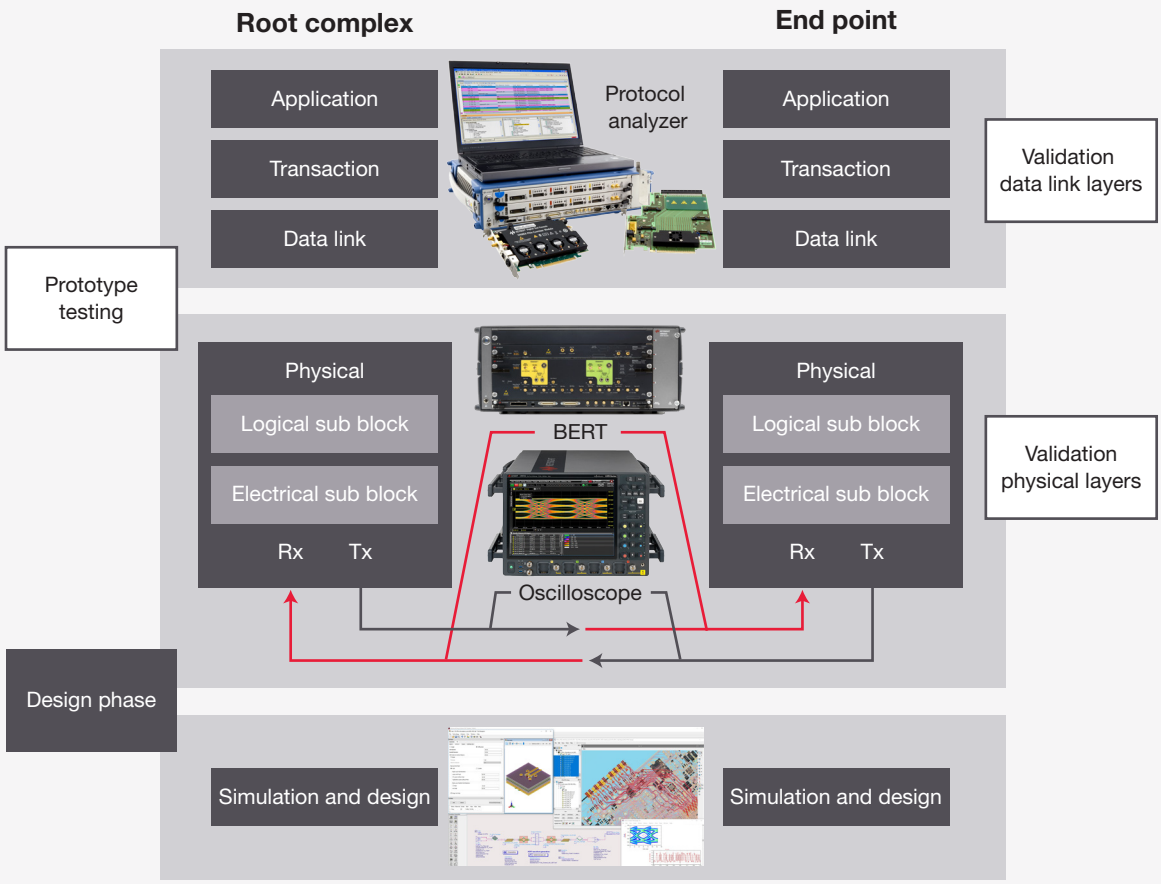

Figure 2: PCIe is a layered protocol, split between physical properties tested by oscilloscopes, BERTs, and protocol analyzers

This paper looks at the differences between the various generations of the PCIe standard. It will help guide you through compliance testing and debug to ensure that your designs meet the specifications set forth by the standards.

Summary of Standards

PCIe 3.0

In PCIe 2.0, the bit rate is 5 GT/s, but with the 20% performance overhead of the 8b/10b encoding scheme, the delivered bandwidth is 4 Gb/s. PCIe 3.0 and later versions use more efficient 128b/130b encoding, whittling the overhead down to a modest 1.5%.

By removing this overhead, the interconnect bandwidth doubled to 8 Gb/s with the implementation of the PCIe 3.0 specification while preserving compatibility with version 2.0 software and mechanical interfaces. With full backward compatibility, PCIe 3.0 provides the same topologies and channel reach for client and server configurations as in PCIe 2.0.

PCIe 1.x and 2.x cards seamlessly plug into PCIe 3.0-capable slots, and vice versa, operating at the highest performance levels supported by those configurations. The PCIe 3.0 specifications comprise the Base and Card Electromechanical (CEM) specifications. The electrical section of the PCIe 3.0 Base Specification defines electrical performance at the integrated circuit (IC) level and supports 8 GT/s signaling.

An eye diagram, appropriately named since it has the appearance of a human eye, indicates the signal quality in the digital domain from the receiver's perspective. It provides valuable insight during design, debugging, and maintenance. As the speed of PCIe increases with each iteration of the standard, signal quality suffers, as exhibited by eye closure. Longer channel lengths also reduce signal quality. Validation testing of the physical layer is more challenging as speeds and channel distances increase. Speeds of 8 GT/s in PCIe 3.0 severely degrade signals at the receiver, viewed (un-equalized) as a closed-eye diagram on an oscilloscope. To enable correct communication, the Tx and Rx need to agree on what level constitutes a one and a zero, and employ several techniques, such as equalization and de-emphasis, to generate a clean eye at the receiver.

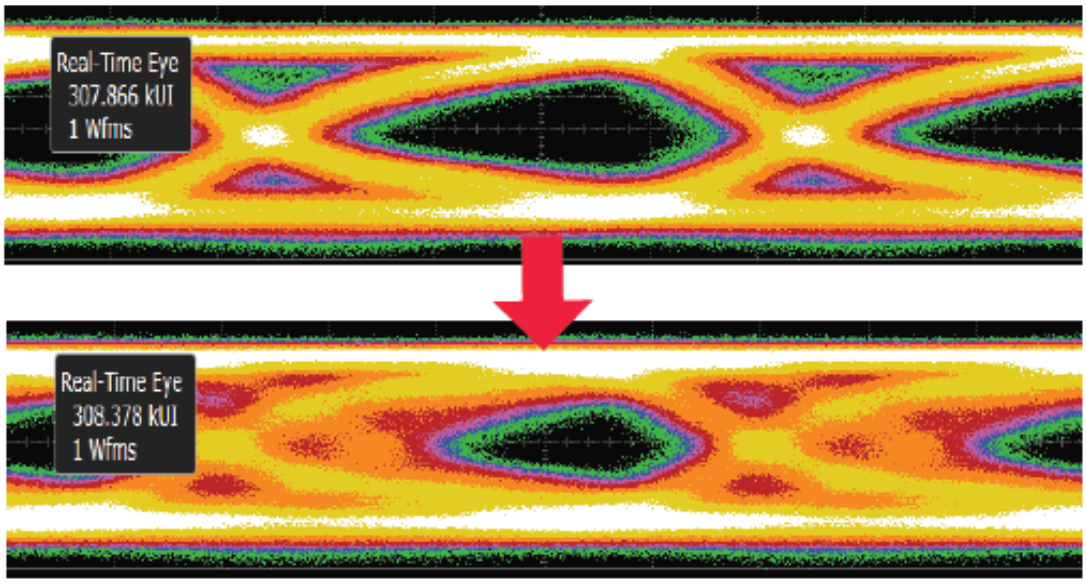

Figure 3: Example of horizontal eye closure (in this case caused by power supply noise)

The PCIe 3.0 standard added receiver equalization and transmitter de-emphasis tests, which are critical for success at 8 GT/s and above. Equalization can occur at the transmitter, the receiver, or both. PCIe 1.x and PCIe 2.x specify a simple form of equalization called transmitter de-emphasis.

De-emphasis reduces the low-frequency energy seen by the receiver. Equalization reduces the effects of greater channel loss at high frequencies. Receiver equalization implementation requires various types of algorithms; the two most common are linear and decision feedback (DFE). Transmitter de-emphasis equalization occurs at the transmitter, while DFE equalization occurs at the receiver. Receiver equalization at the receiver can also include continuous time linear equalization (CTLE) in combination with the DFE.

PCIe 4.0

The PCIe 4.0 standard (Gen 4) debuted in 2017, seven years after completion of PCIe 3.0. Compared with its predecessor, Gen 4 doubled the data rate from 8 to 16 Gb/s. Gen 4 architecture is compatible with prior generations of the technology, from software to clocking architecture to mechanical interfaces.

From a protocol and encoding standpoint, Gen 4 looks a lot like PCIe Gen 3, sharing many elements in common, including 128/130-bit encoding. At first glance, Gen 4 has more in common with Gen 3 than Gen 3 does with Gen 2. But when you increase the speed of the device, you are automatically sending higher frequencies through the same channel. Insertion loss, or attenuation caused by resistance in the link during electrical signal transmission, increases with higher frequency rates.

At 16 GT/s, the Gen 4 signal significantly attenuates in a typical FR4 (the most common printed circuit board, or PCB, material) channel. As a result, ensuring the signal integrity of PCIe 4.0 designs requires additional testing as signal loss at 16 GT/s (PCIe 4.0) is a much bigger concern than at 8 GT/s (PCIe 3.0). Indeed, it is not possible to send a signal on the same 50-cm channel and two connectors without deploying some type of repeater technology, or retimer. PCIe 4.0 added a retimer section to the specification to extend the reach of the channel, increasing the test complexity of devices and especially systems.

The Gen 4 spec consists of three components:

1. The Base Specification is what you need to develop a new ASIC.
2. The Card Electromechanical (CEM) specification defines system behavior at the point of the CEM connector.
3. The Test Specification defines what you need to test to comply with the PCIe 4.0 CEM specification standard.

Some might refer to PCIe Gen 4 as a closed-eye specification. That means that even if you have a perfect transmitter — for example, a transmitter that has effectively zero jitter — by the time you connect the transmitter to a channel, inter-symbol interference will force the eye to close. Your ability to successfully transmit PCIe 4.0 signals depends on the ability of your receiver's equalization strategy to open the eye back up.

Link equalization also becomes more important. As a Gen 4 device links up to another Gen 4 (16 GT/s) device, it undertakes a two-step process. The Gen 4 device will first link up at Gen 3 speeds (8 GT/s), and then if successful, it will repeat the link equalization process to reach 16 GT/s.

With the introduction of Gen 4, lane margining at the receiver allows designers to assess the performance variation tolerance of their system at the receiver. Understanding performance variation is essential because signal performance varies from one card to another. These variances cause increased channel loss, cross talk, and channel discontinuities, resulting in more system noise, deteriorated jitter performance, and signal eye closure.

Lane margining helps system designers assess their design's performance variation tolerance by obtaining margin information with a PCIe 4.0 PHY receiver and controller. The receiver eye width and (optionally, depending on the silicon capabilities) eye height can be monitored in-band from data reported by the PCIe 4.0 PHY. Lane margining is an in-band method by which the sample point of the device's receiver can be adjusted in time (eye width) or in voltage (eye height).

PCIe 5.0

PCI-SIG released the PCIe 5.0 specification in May 2019. The fast release of the Gen 5 interface, completed in less than two years, was a welcome change after the seven-year wait for PCIe 4.0. PCIe 5.0 doubles the transfer rate once again, reaching 32 GT/s while maintaining low power and backward compatibility with previous generations. Gen 5 promises up to 128 GB/s of throughput via an x16 configuration, enabling 400GE speeds in the data center.

Together, 400GE speeds and PCIe 5.0 enable applications such as artificial intelligence (AI), machine learning, gaming, visual computing, storage, and networking. These advances allow users to drive innovation in 5G, cloud computing, and hyperscale data centers.

The PCIe 5.0 standard is a relatively straightforward extension of 4.0. It takes the same approach for Tx and Rx testing used for Gen 4 and similar methods for calibrating eye width and height for receiver stressed jitter testing. The new standard further reduces latency and tolerates higher signal loss for long-reach applications. PCIe 5.0 uses the 128b/130b encoding scheme that debuted with PCIe 3.0 and compatible CEM connectors.

In general, there are minimal spec changes — only those needed to enable another bump in speed or to implement electrical changes to improve signal integrity and the mechanical performance of connectors.

PCI-SIG expects PCIe 4.0 and PCIe 5.0 to co-exist for a while, with PCIe 5.0 used for high-performance needs that crave the most throughput. They include graphical processing units for AI workloads and networking applications.

Keysight provides physical layer Tx test, Rx test, and interconnect design solutions for all generations of the PCIe standard. Figure 4.0 below illustrates the test equipment required for PCIe 5.0.

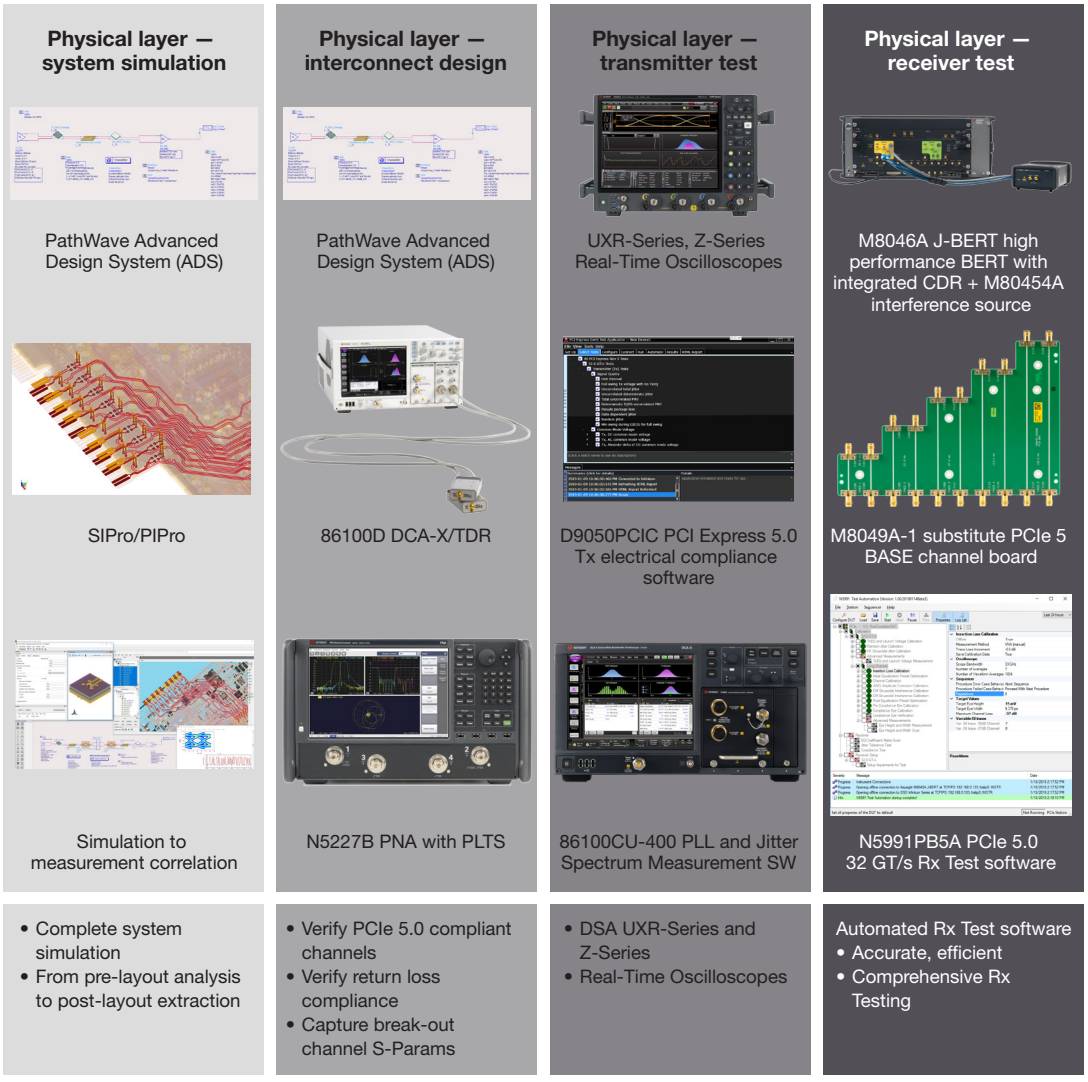

Figure 4: Keysight's PCIe 5.0 test solutions

Conclusion

In June 2019, PCI-SIG announced it would release the standards for PCIe 6.0 in 2021. However, PCIe 6.0 products will not reach the market until at least the end of 2022, if not 2023. The PCIe 6.0 interface plans to double the transfer rate once again to 64 GT/s, providing throughput of 256 GB/s over the same maximum of 16 lanes. PCIe 6.0 will also be backward compatible with previous PCIe generations.

Each generation of the PCIe standard doubled the data transfer rate and increased the complexity of test. Regardless of which generation of the standard you are working on, you need a test solution approved by PCI-SIG to ensure that your products comply with the standard and get to market faster. Keysight provides a total-solution approach to test all generations of the PCIe standard, so you can focus on your next design, rather than spending time learning the details of the test procedures and requirements.

For more information: www.keysight.com/find/pcie

Learn more at: www.keysight.com

For more information on Keysight Technologies' products, applications or services, please contact your local Keysight office. The complete list is available at: www.keysight.com/find/contactus

KEYSIGHT TECHNOLOGIES