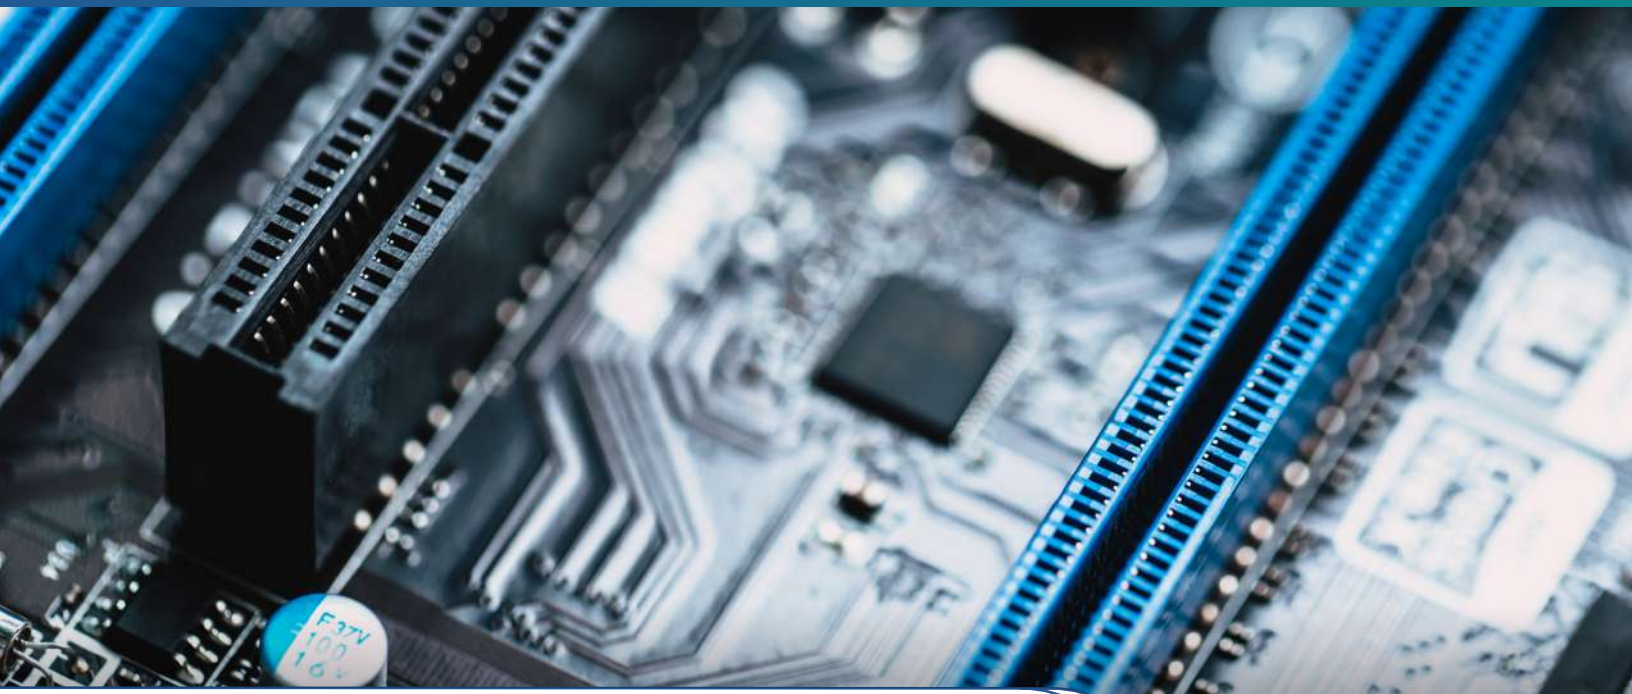

WHITE PAPER

The 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 800 gigabits per second (Gbps) and will likely double in the coming years. The Peripheral Component Interconnect Express (PCIe®) 6.0 standard, the standard's latest iteration, will enable the mass adoption of 800 gigabit Ethernet (GE) technologies in the data center.

PCIe is a core technology that many types of computer servers and endpoint devices use. The PCI Special Interest Group (PCI-SIG®) defines specifications and compliance tests that guarantee the interoperability of PCIe systems. PCIe 6.0 provides full-duplex bandwidth of approximately 256 Gbps for a 16-lane system.

Evolving from PCIe 1.0 in 2003 and supporting 2.5 gigatransfers per second (GT/s), the current PCIe standard, PCIe 6.0, came out in 2022 and supports 64 GT/s.

PCIe is scalable and slots come in different configurations of bidirectional lanes: x1, x4, x8, x16, x32. The numbers represent 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 forth. 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.

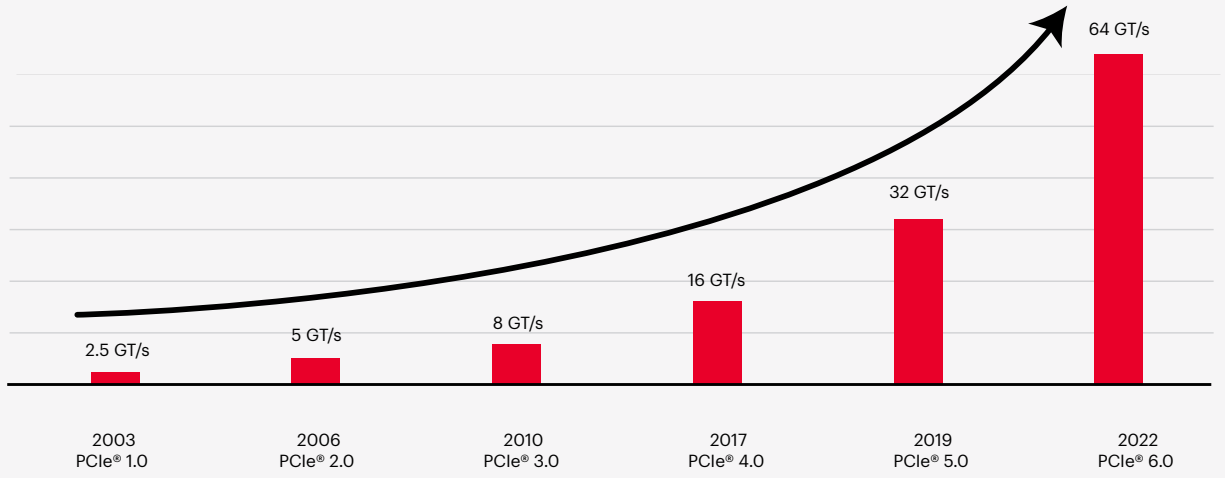

Figure 1: Evolution of the PCIe standard

As a serial point-to-point multilane interconnect between two devices, PCIe communicates directly with devices via a switch that manages data flow. As a result, the PCIe specification has provisions for "hot swapping" or "hot plugging" of PCIe devices without shutting down the system.

The Peripheral Component Interconnect Special Interest Group (PCI-SIG) requires vendors to conduct compliance testing for interoperability that includes a list of preset or transmitter equalization settings which must be electrically verified. We define these presets as a ratio of pre-shoots to de-emphasis, and it is important to note that the presets for PCIe 6.0 differ from those in previous versions. You must conduct three categories of testing to reach official PCI-SIG-approved compliance: the physical layer, the data link layer, and general interoperability.

Different categories of tests require different equipment setups:

1. Oscilloscopes validate physical-layer transmitter (Tx) testing.
2. Oscilloscopes and bit error ratio testers (BERTs) validate physical-layer link equalization testing (Tx & Rx).
3. Protocol analyzers validate data link layers.

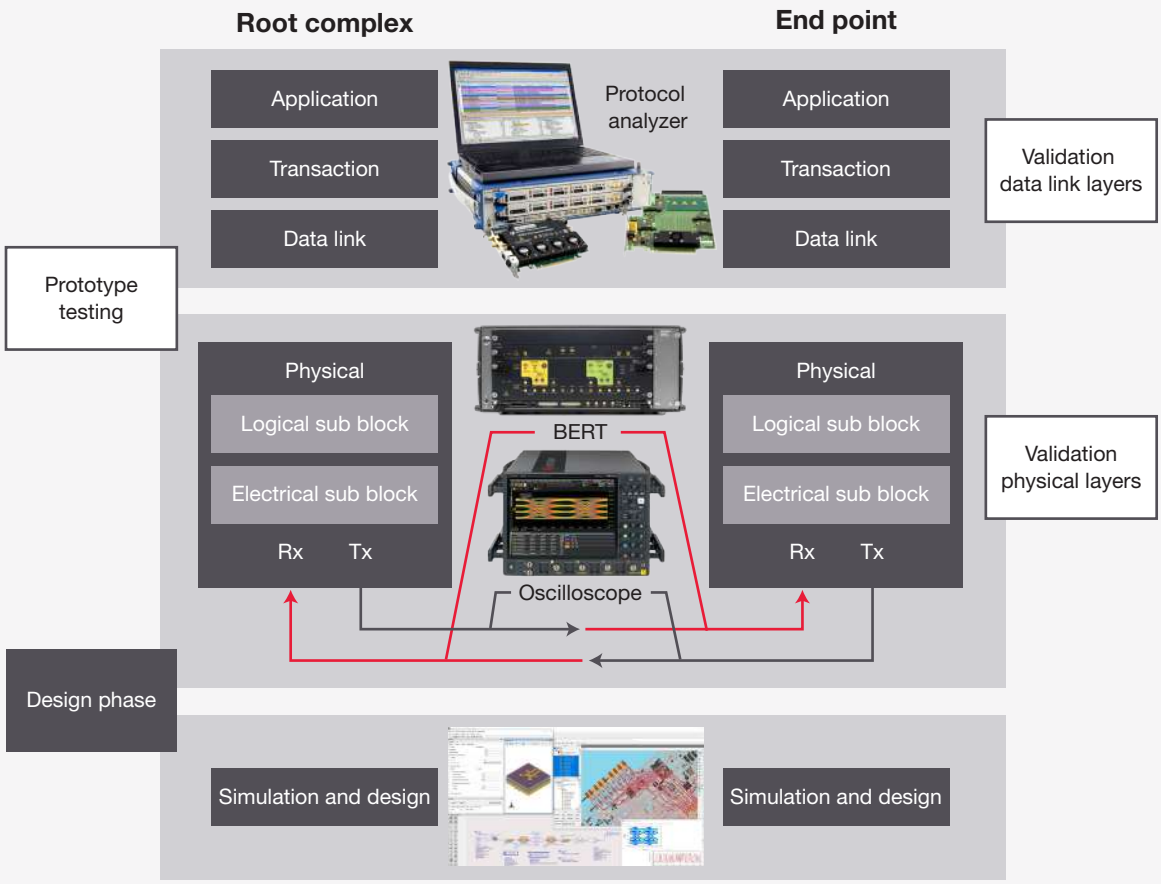

Figure 2. PCIe is a layered protocol; the testing of physical properties is split between oscilloscopes, BERTs, and protocol analyzers

This white paper examines 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. When reviewing the standard, you will notice that the PCIe specification framework consists of multiple specifications:

1. The Base specification will help you develop a new application-specific integrated circuit (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 CEM specification standard.

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 that support 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, an apt name 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 you can see by the eye closure in Figure 3. 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; this will appear (un-equalized) as a closed-eye diagram on an oscilloscope. To enable accurate communication, the Tx and Rx need to agree on what level constitutes a one and a zero, and employ techniques such as equalization and de-emphasis, to generate a clean eye at the receiver.

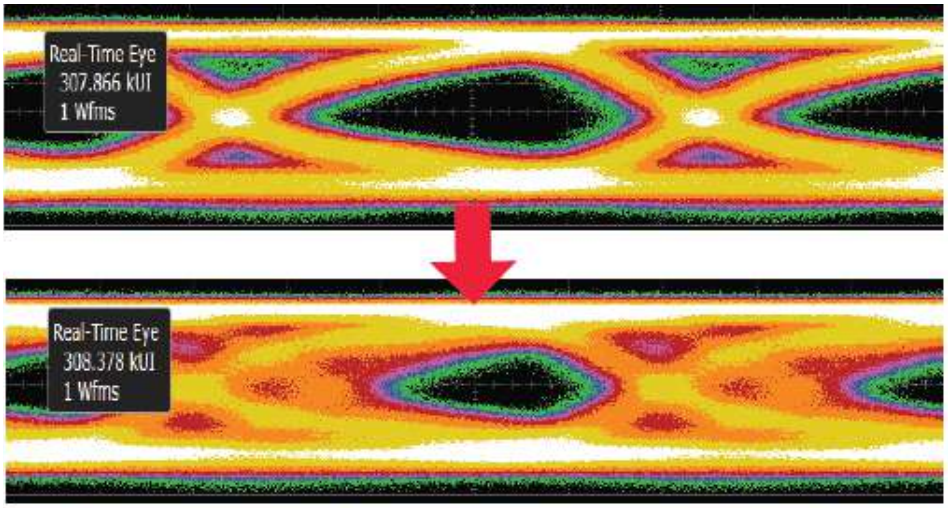

Figure 3. Example of horizontal eye closure

The PCIe 3.0 standard adds 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 that the receiver picks up. 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.

To enable maximum distance between transmitters and receivers, PCIe 3.0 introduced an active equalization adaptation process, where the receiver tunes the transmitters pre-shoot and de-emphasis for the best fit to its own equalization capabilities for the specific transmission line. The new capability required a completely new physical layer test, the link equalization tests for receiver and transmitter. The link equalization receiver test checks if a receiver can tune its link partner's transmitter equalization under worst-case stress conditions. The link equalization transmitter tests check if the transmitter performs the changes physically and logically as requested by the link partner's receiver.

PCIe 4.0

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

From a protocol and encoding standpoint, PCIe 4.0's maximum transfer rate of 16 GT/s looks a lot like the 8 GT/s of PCIe 3.0. They share many common elements, including 128/130-bit encoding. At first glance, PCIe 4.0 has more in common with PCIe 3.0 than PCIe 3.0 does with PCIe 2.0. However, when you increase the speed of the device, you automatically send higher frequencies through the same channel. Resistance in the link during electrical signal transmission causes insertion loss, or attenuation and increases with higher frequency rates.

At 16 GT/s, the PCIe 4.0 signal significantly attenuates in a typical FR4 channel, the most common printed circuit board material. 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.

Even with an increase in test complexity, the number of tests for PCIe 3.0 for 8 GT/s is higher than those for PCIe 4.0 for 16 GT/s. This is because PCIe 3.0 requires testing three different channel scenarios: short, medium, and long. PCIe 4.0 tests the long channel scenario only.

As with PCIe 3.0, PCIe 4.0 is sometimes referred to as a closed-eye specification. That means that even if you have a perfect transmitter, a transmitter with essentially 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.

There is a two-step process for link equalization as a PCIe 4.0 device supporting 16 GT/s, links to another PCIe 4.0 device which supports 16 GT/s. First, you establish a link at 8 GT/s. Then, if successful, the link equalization process repeats to reach 16 GT/s.

With PCIe 4.0, designers should assess the performance variation tolerance of their system. Understanding performance variation is essential because signal performance varies from one card to another. These variances cause increases in channel loss, cross talk, and channel discontinuities that result in more system noise, a deterioration in jitter performance, and signal eye closure.

PCIe 5.0

PCI-SIG released the PCIe 5.0 specification in May 2019. Completed in less than two years, the fast release of PCIe 5.0, 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. PCIe 5.0 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 that PCIe 4.0 uses, 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.

New in PCIe 5.0 is an optional speed bypass mode for the link training which enables training from 2.5 GT/s, directly to 32 GT/s and skips 8 GT/s and 16 GT/s. This helps to reduce the link bring-up time in systems where you know the transmitter, channel, and receiver conditions, like in embedded systems. There is a new training path for link equalization testing for 32 GT/s. This means that for the first time, you can perform link equalization RX and TX testing on lanes other than lanes that behave like a lane 0, using a single lane test setup.

In general, there are minimal spec changes other than the ones you need to enable another bump in speed, or to implement electrical changes to improve signal integrity and the mechanical performance of connectors.

PCIe 6.0

The PCI-SIG released the PCIe 6.0 specification in January 2022. PCIe 6.0 technology is the first PCI Express standard to use pulse amplitude modulation 4-level (PAM4) signal encoding (see Figure 4), enabling PCIe 6.0 devices to achieve twice the throughput of PCIe 5.0 devices while maintaining the same channel bandwidth. PCIe 6.0 technology reaches up to 64 GT/s while maintaining low power and backward compatibility with previous generations. PCIe 6.0 promises up to 256GB/s of throughput via an x16 configuration, enabling 800GE speeds in the data center.

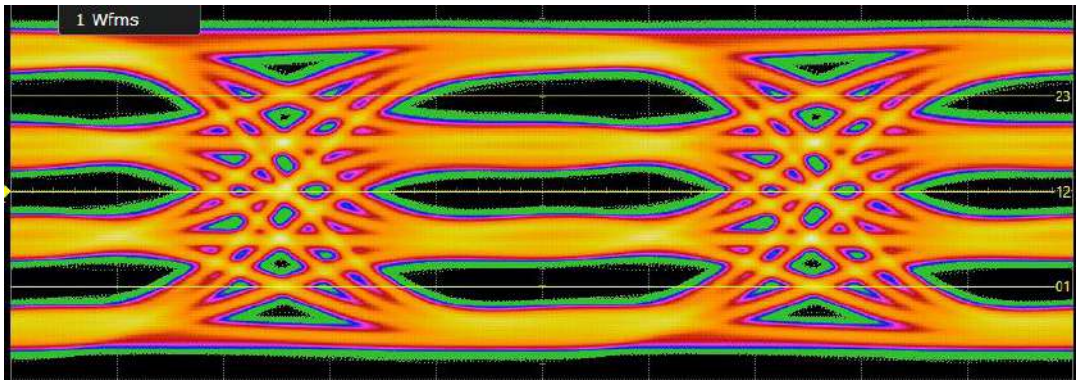

Figure 4. PCI Express 6.0 signal at 32GBaud/64GT/s

Together, 800GE speeds and PCIe 6.0 technology enable applications such as AI, machine learning, gaming, visual computing, storage, and networking. These advances allow users to drive innovation in 5G, cloud computing, hyperscale data centers, and beyond.

The PCIe 6.0 standard is a significant upgrade over PCIe 5.0 technology with its use of PAM4 signaling. However, it uses the same high-level approach for Tx and Rx testing while adding some new transmitter measurements that are specific to PAM4

encoding. Similar to previous generations, PCIe 6.0 devices employ transmitter and receiver equalization for 64 GT/s operation along with requiring forward error correction (FEC).

In addition to these electrical changes, PCIe 6.0 introduces flow control unit (FLIT) encoding. PCIe FLITs have a static payload size which eliminates the need for extra packet overhead, thus reducing overall latency. FEC causes a latency increase, but the increase is offset by FLIT encoding.

Because of the close timing of the specification releases, and the time to market for new devices, users can expect PCIe 5.0 and PCIe 6.0 to co-exist for a while. However, applications which require high performance and throughput will utilize PCIe 6.0. Examples of these high-performance applications include graphical processing units for AI workloads, high throughput networking applications, and Compute Express Link (CXL) interconnect.

Keysight provides physical layer Tx test, Rx test, and interconnect design solutions for all generations of the PCIe standard. Figure 5 illustrates the test equipment that you need for PCIe 5.0.

Physical Layer – System Simulation	**Physical Layer - Interconnect Design**	**Physical Layer – Transmitter Test**	**Physical Layer – Receiver Test**	**Protocol Layer Test**
Complete system simulation from pre-layout analysis to post-layout extraction	Verify PCIe compliant channels and return loss compliance	Automated Tx test software for characterization and compliance	Automated Rx test software for PCIe Base and CEM spec versions	Deep protocol analysis capability with root complex and endpoint emulation
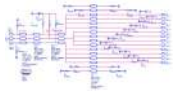 PathWave Advanced Design System (ADS)	 PathWave Advanced Design System (ADS)	 SW02PCIE PCI Express full Tx test suite	 N5991PB6A PCIe 6.0 64 GT/s Rx test software	 P5551A PCIe 5.0 Exerciser
 SIPro/PIPro	 N1000A DCA-X sampling oscilloscope w/ TDR	 UXR-Series oscilloscope	 M8040A BERT w/ integrated CDR & interference source	 P5552A PCIe 5.0 Analyzer
 Simulation to measurement correlation	 N5227B PNA w/ PLTS	 N1000A DCA-X w/ FlexDCA jitter & spectrum measurement software	 M8049A ISI channel boards	 P5563A Test Backplane

Figure 5. Keysight's PCIe test solutions

Conclusion

The PCI-SIG released the PCIe 6.0 specification in January 2022 and estimates that it will take 1 to 1.5 years for PCIe 6.0 products to reach the market. This means the soonest PCIe 6.0 devices will be available is early to mid-2023. While maintaining backwards compatibility to previous generations, the PCIe 6.0 interface doubles the transfer rate to 64 GT/s and provides throughput of 256 GB/s over the same maximum of 16 lanes. You can achieve this by adopting PAM4 signaling for PCIe which doubles the data encoded in each transition. The introduction of this new encoding scheme adds a layer of complexity that will require robust test solutions.

Regardless of which generation of the PCIe specification you are working on, you need a test solution the PCI-SIG aligns with 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 specification 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