



## PCI-X® 2.0 Frequently Asked Questions

**Q1: How are PCI-X versions 1.0 and 2.0 related to PCI? Are they the same?**

A1: PCI-X 2.0 is the next generation of PCI. It builds upon previous generations of PCI including the electricals, protocols, signals names and functions, connectors, etc. It also maintains backward compatibility with conventional PCI. It is the next logical advance in the world's most popular PC bus. There have been many generations of PCI, which all build upon each other. The PCI bus began with a 32-bit / 33MHz specification. Over time, to increase performance, 64-bit and 66MHz versions were introduced. To increase the bus speed and reduce latency PCI-X 1.0 was developed, with a maximum clock speed of 133 MHz. PCI-X 1.0 also introduced improved protocols, such as the split-transaction protocol which allows more efficient use of bus bandwidth, resulting in throughput gains beyond the simple increases in clock speed and bus width. Because of the demand for even higher throughput and to improve error correction, the PCI-X 2.0 specification was developed. It extends the bus frequency to 266 MHz and 533MHz and adds advanced features like ECC, while still maintaining backward compatibility to the first generation.

**Q2: What speed grades are in the PCI-X 2.0 specification?**

A2: There are 4 speed grades in the PCI-X 2.0 specification: PCI-X 66, PCI-X 133, PCI-X 266, and PCI-X 533. The PCI-X 66 and PCI-X 133 speed grades were included in the PCI-X 1.0 specification; they support 66MHz, and 133MHz PCI-X respectively. 100MHz PCI-X has been implemented in the market by using PCI-X 133 adapter cards. Both PCI-X 266 and PCI-X 533 are new to PCI-X 2.0; they are the 266MHz and 533MHz versions of the specification. All four speed grades are included in the PCI-X 2.0 specification.

**Q3: Is PCI-X 2.0 backward compatible with all generations of PCI?**

A3: Yes. PCI-X 2.0 is logically backward compatible to all previous generations and speed grades of PCI and PCI-X. PCI-X 266 and PCI-X 533 devices are electrically compatible with 3.3V and 1.5V IO buffers only; they are not compatible with 5V PCI. (Note: The latest version of the PCI local-bus specification (version 3.0) obsoletes 5V-only add-in cards and 5V-only system slots.)

**Q4: What new features are introduced in the PCI-X 2.0 specification?**

A4: The PCI-X 2.0 specification includes new features such as Error Correction Code (ECC), 1.5V signaling, source-synchronous strobes, device ID messages, and a 16-bit version. The ECC improves the robustness of the interface. Likewise, the 1.5V signaling and strobes improve the performance so that the bus can run at 533MHz. Device ID messages are designed to enable a whole new class of peer-to-peer transfer applications. The 16-bit version of the bus is designed for embedded applications where bandwidth can be traded-off to reduce device pin counts.

**Q5: What are the initial target applications for PCI-X 266 and PCI-X 533?**

A5: The initial target applications for PCI-X 2.0 technology are the workstation and server segments of the computer industry. PCI-X 266 and PCI-X 533 will provide the greatest value in those applications that require the highest possible bandwidth and where customers need and want low cost and backward compatibility with existing PCI slots: 10 gigabit Ethernet, 10 gigabit Fibre Channel, InfiniBand and multi-function cards requiring large aggregate bandwidth. These high bandwidth applications, typical of servers and workstations, need the performance offered by PCI-X 2.0. Backward compatibility is typically a must-have for investment protection in these markets.

**Q6: How can the 16-bit versions of PCI-X be used?**

A6: The 16-bit application of PCI-X is uniquely suited for applications that require a reduced cost interface. The 16-bit functionality can be implemented either by creating a stand-alone 16-bit bus, or by breaking a 64-bit bus into 4 different segments of 16-bits each. For example, a peer-to-peer bridge with a 64-bit PCI-X 533 interface on its secondary side could alternatively implement 4 independent buses, each 16 bits wide, to support 4 independent loads each running at 533MHz. This capability may be particularly interesting for embedded applications.

**Q7: Will PCI-X 2.0 be used in desktops?**

A7: Adoption of PCI-X 2.0 in the desktop marketplace will be determined by the needs of the developers. It is not the purpose of the PCI-SIG® to determine when and where new technologies are adopted. For high-end desktops and workstations PCI-X has an interesting value proposition. A single 16-bit or 32-bit, PCI-X 66 bus can support 4 loads, and operates at 3.3V signaling levels (removing the added cost of 5V tolerance). Because PCI-X 2.0 maintains the same low-cost interface, while providing options for reducing pin count, lower cost versions can be created.

**Q8: Does PCI-X 2.0 support only point-to-point loads? Or, does it support multi-drop too?**

A8: PCI-X 2.0 supports both point-to-point loads and multi-drop loads. A PCI-X bus running at 66MHz can support four card slots. A PCI-X bus running at 100MHz can support two slots (when using PCI-X 133 adapter cards). At higher speeds (PCI-X 133, PCI-X 266, and PCI-X 533) one slot is supported.

In embedded environments where components are soldered down and connectors aren't used there exists the possibility for multiple loads at 133MHz. However, these electrical environments are not defined in the specification, since interoperability isn't an issue in embedded applications. It is left to the user to validate multiple embedded loads at the higher clock frequencies.

**Q9: PCI-X 533 electrical parameters appear as “design goals” in the specification. What does this mean, and when will the final numbers be released?**

A9: The PCI-X Electrical and Mechanical Specification includes fully validated and released parameters for PCI-X 266 and “design goals” for PCI-X 533. “Design goals” have been fully validated to guarantee interoperability, but are not fully released. This allows for the possibility of further refinement of the PCI-X 533 design parameters to improve manufacturing yields. The PCI-SIG believes that PCI-X 266 adequately addresses the bandwidth requirements for today's devices and that the full release of PCI-X 533 parameters should occur after suppliers have gained adequate experience building PCI-X 266 devices.

**Q10: Why is the time-to-market for PCI-X 266 and PCI-X 533 shorter than that of other technology transitions?**

A10: The transition to PCI-X 266 and PCI-X 533 is shorter than that of other technologies because it leverages and reuses much of the technology in the PCI-X 1.0 specification. The architecture, state machine, bus functional model, device drivers, signals and signal functionality, pin-outs, connector, test suites, form factors, layouts, and design tools are all either identical to PCI-X 1.0, or are highly leveraged.

**Q11: What test tools and other infrastructure will be available to support development of PCI-X 266 and PCI-X 533 products?**

A11: Analysis, stimulus, and parametric test tools for PCI-X 2.0 are readily available. Since PCI-X 2.0 is a natural extension of PCI-X 1.0, test tools and infrastructure supporting PCI-X 1.0 development are extensible to support PCI-X 2.0.

**Q12: What enhancements are anticipated to the PCI-X specification in the future?**

A12: Enhancements to the PCI-X specification will come directly from the feedback of PCI-SIG members. This is an important factor in determining the direction of new PCI-SIG standards. It is believed that the PCI-X 2.0 specification has all the functionality and features that are required to support today's applications. However, future versions of the PCI-X specification may include new features to support new functionality as required by its members. Feedback from the PCI-SIG membership will be important in determining which features to include in the specification.

**Q13: Why does PCI-X 2.0 provide such cost-effective bandwidth?**

A13: PCI-X 2.0 is highly cost-effective throughout the entire cost structure. PCI-X 2.0 uses very little silicon for its protocol engine and uses very little silicon for its physical interface. It requires minimal redesign from earlier versions of PCI, maintains the same low cost connector, uses similar design tools and testing methodology and equipment. It uses the same board form factor that has already been designed into previous systems and requires no new BIOS, device drivers, or operating systems. For high-bandwidth systems it maintains similar pin counts to serial technologies.

**Q14: How can I get a copy of the PCI-X 2.0 specification?**

A14: A benefit of membership in the PCI-SIG is access to both published specifications and draft specifications (see membership benefits at: [http://www.pcisig.com/membership/about\\_us/](http://www.pcisig.com/membership/about_us/)). The specification is also available for purchase by non-members. Current specifications can be obtained either through the PCI-SIG Web site: [www.pcisig.com](http://www.pcisig.com), or by calling +1 (503) 291-2569.