

# The Making of a Server

Mickey Applebaum

To ensure that your company's servers run smoothly, you should purchase network components that meet the system requirements for the network operating system you choose. One of the most important—and often overlooked—system requirements for NetWare 5 and NetWare 4.2 is that you use a server-class computer as the server. (For a complete list of system requirements, read the documentation included with the version of NetWare you are using. You can also access in-depth product information, including system requirements, for Novell products at <http://www.novell.com/catalog>.) Although the system requirements for NetWare 3 and NetWare 2 do not include a server-class computer, Novell recommends that you use a server-class computer for all versions of NetWare.

How do you identify a server-class computer when virtually every computer that is available today appears to have the same internal design? This article outlines the capabilities a server-class computer should have. In particular, this article examines the key component of a server-class computer: the bus.

## IDENTIFYING A SERVER

Because servers typically store critical data that users must be able to access at all times, you cannot risk losing this data because of a hardware failure. To prevent data loss, you should always demand fault-tolerant capabilities when you are choosing a computer to act as a server.




Figure 1. The PCI shared bus architecture

For example, you should ensure that the computer includes Error Correction Code (ECC) memory, which provides the most basic fault-tolerant capabilities. You should also look for other components that provide fault-tolerant capabilities, including hot-swap hard drive bays, hot-swap Peripheral Component Interconnect (PCI) slots, redundant SCSI host interface or Redundant Array of Independent Drives (RAID) host interface, redundant power supplies, and redundant cooling fans. You may also want to look for fused or heat-protected circuitry, which can prevent data loss and physical damage in the event that the computer is exposed to excess heat and voltage.

Identifying a computer's fault-tolerant capabilities is relatively simple. You should carefully read the manufacturer's specification

sheet, which should include a complete list of the computer's components. You can also contact the manufacturer to inquire about a particular component.

You should also purchase a computer with a high bus throughput. With older computers that include a MicroChannel bus or an Extended Industry Standard Architecture (EISA) bus, you can easily calculate bus throughput. Both MicroChannel and EISA buses are based on an arbitrated bus architecture, which enables only one adapter at a time to access the computer's processor or memory. As a result, you can use the following formula to calculate the bus throughput for a MicroChannel bus or an EISA bus:

(the width of the bus in bits multiplied by the speed of the bus in MHz) divided by 8, or the number of bits per byte

Most newer computers include a PCI bus, which is based on a shared bus architecture. This architecture allows multiple buses to access the computer's processor or memory simultaneously, thus increasing the overall throughput. As result, computers can contain multiple PCI buses.

You can use the above equation to calculate the bus throughput for a PCI bus. However, using this equation with a PCI bus may be misleading since many manufacturers do not specify the number of PCI buses in a particular computer. You may have to contact the manufacturer to find out whether or not the computer includes multiple buses. Since there are different types of PCI bus systems, you may also need to contact the manufacturer to find out the type of PCI bus system used in the computer. (For more information about PCI buses, see "Bus Stops.")

## EXAMINING THE BUS

To understand the benefits of multiple PCI buses and the types of PCI bus systems that exist, you should be familiar with the shared bus architecture on which a PCI bus is based. In a shared bus architecture, the PCI bus is isolated from the computer's processor and memory and is connected to these components via a

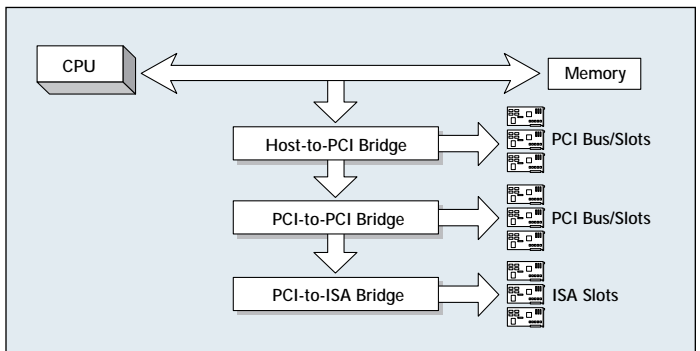


Figure 2. The PCI-to-PCI bridge system

## Bus Stops

If you want more information about Peripheral Component Interconnect (PCI) buses, you can visit the following web sites:

- Novell's PCI hot-plug white paper (<http://www.novell.com/whitepapers/iw/hotplug.html>)
- Compaq's PCI load-balancing white paper (<http://www.compaq.com/support/techpubs/whitepapers/ecg0730398.html>)
- Adaptec's PCI bridge white paper (<http://www.adaptec.co.kr/support/pds/scsi/whitepa.txt>)

- PCI Special Interest Group (<http://www.pcisig.com>)
- Rockwell Network Systems' PCI local bus white paper ([http://icsl.ee.washington.edu/~lyzhou/faq\\_pci.html](http://icsl.ee.washington.edu/~lyzhou/faq_pci.html))

You may also want to read about next-generation I/O (NGIO) technology, which may eventually replace PCI buses:

- AberdeenGroup's NGIO paper ([http://developer.intel.com/design/servers/future\\_server\\_io/documents/aberdeen/Q2991357.htm](http://developer.intel.com/design/servers/future_server_io/documents/aberdeen/Q2991357.htm))
- Intel's take on NGIO (<http://www.intel.it/solutions/archive/issue14/stories/top4.htm>)

host-to-PCI bridge. (The bus includes both the PCI devices in the computer's PCI slots and the embedded PCI devices on the computer's systemboard. See Figure 1.) This connection is similar to a hub that connects a server to a 10Base-T Ethernet network.

Multiple PCI buses are important due to the inherent limitations of a shared bus architecture. One such limitation is that a single PCI bridge cannot support more than 10 loads simultaneously. What is a load? An embedded PCI device on a systemboard counts as one load. A PCI adapter in a PCI slot, on the other hand, counts as two loads.

The load limitation greatly reduces the number of PCI slots that you can use for PCI devices, such as adapters. Because each adapter in a PCI slot counts as two loads, you can install only a limited number of adapters in a computer that includes one PCI host-to-bus bridge.

### PCI-to-PCI Bridge System

The PCI-to-PCI bridge system was developed to alleviate the limitations of a shared bus architecture. A PCI-to-PCI bridge allows PCI buses to be daisy-chained together, thus making room for more PCI devices. (See Figure 2.) In addition, PCI-to-PCI bridges can be embedded on an adapter to free up even more space. By using PCI-to-PCI bridges, one PCI board may include multiple SCSI host adapters and multiple Ethernet adapters.

The PCI-to-PCI bridge system addresses one limitation but introduces another limitation: Daisy-chaining PCI buses through a PCI-to-PCI bridge may hinder performance by increasing the load on the host-to-PCI bridge. This increased load turns the PCI bridge controller into an I/O bottleneck. In addition, when performing I/O operations, the computer's processor slows to the speed of the bus.

### Multiple-Peer System

The multiple-peer system was developed to alleviate the limitations of a shared bus architecture and the PCI-to-PCI bridge system. Most newer server-class computers use the multiple-peer system, which consists of multiple PCI buses. (See Figure 3.) Most newer server-class computers also use a 100 MHz PCI bus clock. Each PCI bus connects directly to the host bus, rather than being daisy-chained through one PCI bridge controller. This connection allows the computer's processor to directly access each PCI bus. As a result, concurrent data paths can exist between the PCI bus and the host bus. This capability increases the overall aggregate throughput between the computer's processor and its PCI devices.

As with the PCI-to-PCI bridge system, the multiple-peer system supports multiple PCI buses, thus maximizing PCI slots. The

main difference between these two bus systems is that the multiple-peer system offers better performance by providing each PCI bus with direct access to the host bus. However, you must ensure that you balance the load generated by the PCI devices across the computer's multiple PCI buses. If you install all of the PCI-based Ethernet adapters and disk controllers in slots on the same PCI bus, you will gain no performance benefit.

## CONCLUSION

You should carefully inspect the computers you are using as servers to ensure that they are server-class computers, providing the level of fault tolerance and performance your company needs. No matter how small or how simple your company's network is, you should use computers that include ECC memory. If your company's network stores mission-critical data, you should use computers that include advanced fault-tolerant components, such as hot-swap hard drive bays and redundant power supplies.

If your company's network is relatively small with minimal traffic, you can probably use a computer with one PCI bus, as long as you realize that this configuration offers limited expansion capabilities. If your company's network includes multiple network segments, you should use a computer with multiple PCI buses. You should also use multiple PCI buses if you want to implement disk-duplexing capabilities across multiple host adapters or if you need multiple host adapters to meet your company's requirements for hard-drive capacity. If you decide to use a computer with multiple PCI buses, you should choose a computer that is based on the multiple-peer system, rather than the PCI-to-PCI bridge system. The multiple-peer system offers higher bus throughput, enabling you to get more bang for your bus.

Mickey Applebaum provides technical support on the Internet for The Forums (<http://theforums.com>) and operates Proactive Team Solutions, a consulting firm that is located in Salt Lake City, Utah.

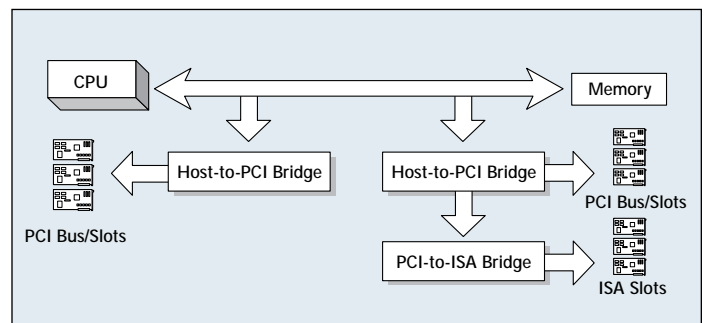


Figure 3. The multiple-peer system