PCI Comparison, 32 Vs. 64-Bit and 33Mhz Vs. 66 MHz

Introduction

The original design for the PCI bus was to move high bandwidth peripherals closer to the CPU for performance gains. Graphics-oriented operating systems (OS) and other high bandwidth functions such as Ultra2 SCSI, Fibre Channel, Fast Ethernet, and 3D graphics are consuming more PCI bus bandwidth. This need for more bandwidth has compelled system vendors to find ways of increasing the throughput not only on the PCI bus, but throughout the entire system.

Today, heavily loaded systems move about 90 MB/sec of data on the 133-MB/sec PCI bus, the difference being overhead. The majority of system vendors have implemented a 32-bit PCI bus at 33 MHz. We'll talk about implementing 64-bit PCI bus extensions and increasing the PCI clock to 66 MHz as a way of increasing system bandwidth and efficiency. These PCI changes will result in part in increasing overall system performance; however, other components like the OS and the CPU also play a major role in increasing system performance. Let's take a look at the benefits of implementing 64-bit and 66-MHz PCI technologies.

PCI Architecture

Since the 32-bit PCI bus is currently able to transfer 133 MB/sec of data less the overhead, what results from changing the bus to 64 bit and 66 MHz? How are PCI bus cycles affected by 64-bit transfers? What are the additions to the PCI bus for 64-bit extensions?

PCI Bus Cycles

Address and data transfers are multiplexed over the same lines on the PCI bus, the address is sent first and then the data. A 32-bit PCI bus has 32 data lines and is able to do 32-bit data transfers and 32-bit memory addressing or 64-bit addressing using two 32-bit PCI cycles known as Dual Address Cycles (DAC).

Memory addressing is not what constrains the PCI bus or system performance. 32-bit addressing allows access to 4 GB of memory--systems such as SMP systems need to address more than this range of memory (see the illustration below). More memory can be addressed with 64-bit addressing in one PCI cycle or two 32-bit cycles using DAC, with the first cycle sending the low address and the second cycle sending the high address.

128 Bytes of Data Being Transferred

	32-bit PCI Bus	64-bit PCI Bus
32-bit Address	A 1 Cycle D 32 Cycles	A 1 Cycle D 16 Cycles
64-bit Address	A 2 Cycles D 32 Cycles	A 1 Cycle D 16 Cycles

In the illustration to the right, notice the number of PCI cycles it takes to send the same 128 bytes of data over a 32-bit PCI bus versus a 64-bit bus, assuming the PCI bus is not interrupted. 64-bit PCI bus transactions are more efficient, both for addressing and data, because the number of PCI cycles is reduced to half.

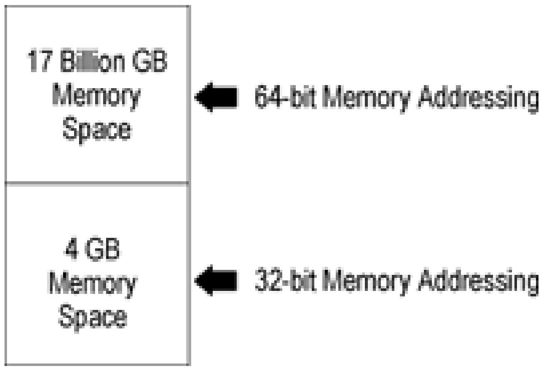

64-Bit Extension Pins

A 64-bit extension to the 32-bit bus architecture requires an additional 39 signal pins which are AD[64::32], C/BE[7::4]#, REQ64#, ACK64#, and PAR64. As we said before, data and addressing are multiplexed over the same pins, either AD[31::00] for 32 bits or AD[31::00] and AD[64::32] for 64 bits. During an address phase, AD[64::32] is used to send the upper 32-bits of a 64-bit address or during a data phase, an additional 32-bits of data. To transfer a 64-bit address in one PCI cycle using the 64-bit bus, you must use the DAC command and assert REQ64#. (Asserting or deasserting a signal means that a particular message is present or missing on the line.) To transfer the additional 32-bits of data on AD[64::32], REQ64# and ACK64# must be asserted.

The Bus Command and Byte Enable Commands are multiplexed over C/BE[7::4]# pins; Bus Commands are transferred during the address phase and Byte Enable Commands during the data phase. An even parity bit, PAR64, protects the AD[63::32] pins from data corruption. PAR64 has the same timing as AD[63::32] but is delayed by one clock cycle.

64-bit transactions are negotiated using a transaction between a master and a target asserting REQ64# and ACK64#. Devices determine if they are connected to a 64-bit bus by asserting or deasserting REQ64# when RST# is deasserted. Only memory commands support 64-bit data transfers.

64-Bit Bus Benefits

Many factors play into overall system performance and affect the industry's progress to 64-bit PCI. The 32-bit PCI bus is not in itself slowing system performance. Peripheral devices such as SCSI, IDE, and Fast Ethernet by themselves do not use the full potential of the current PCI bus. Interaction between the PCI bus, the Host Bridge, DRAM, and the CPU commonly slow down PCI transfers.

Let's also consider the transactions in the system. The CPU communicates with Dynamic Random Access Memory (DRAM) and the Host Bridge: the Host Bridge in turn communicates with the PCI bus and DRAM. Direct Memory Access (DMA) devices transfer data directly to DRAM through the PCI bus as shown here.

ATX Power Supply FAQ

HP/UX Related FAQ

SCSI LVD Related FAQ

The burst rate of data throughput on the PCI bus doubles with 64-bit data transfers. 64-bit DMA devices can move data in 64-bit chunks directly between the PCI bus and DRAM if the PCI bus and DRAM are set up the handle 64-bit transfers. The system with a 64-bit PCI bus is less congested; 64-bit devices in the system get on and off the bus in half the time, making the PCI bus more efficient. In essence, since 64-bit transfers achieve better PCI bus utilization and more devices can be added to the bus before realizing the bus' full bandwidth potential. The more heavily weighted your system becomes with peripheral devices the more it benefits from a 64-bit wide PCI bus.

Notice in the following drawings, the number of components that would be affected by increasing the PCI bus to 64-bits, the host bridge, the DRAM, the CPU and even the OS and driver. Since the majority of transactions are data transfers from the PCI bus to DRAM and from DRAM to the PCI bus, you could consider that increasing the bandwidth of the bus would increase the performance of the system. However, the CPU and the OS may become the bottleneck even if implementing a 64-bit PCI bus.

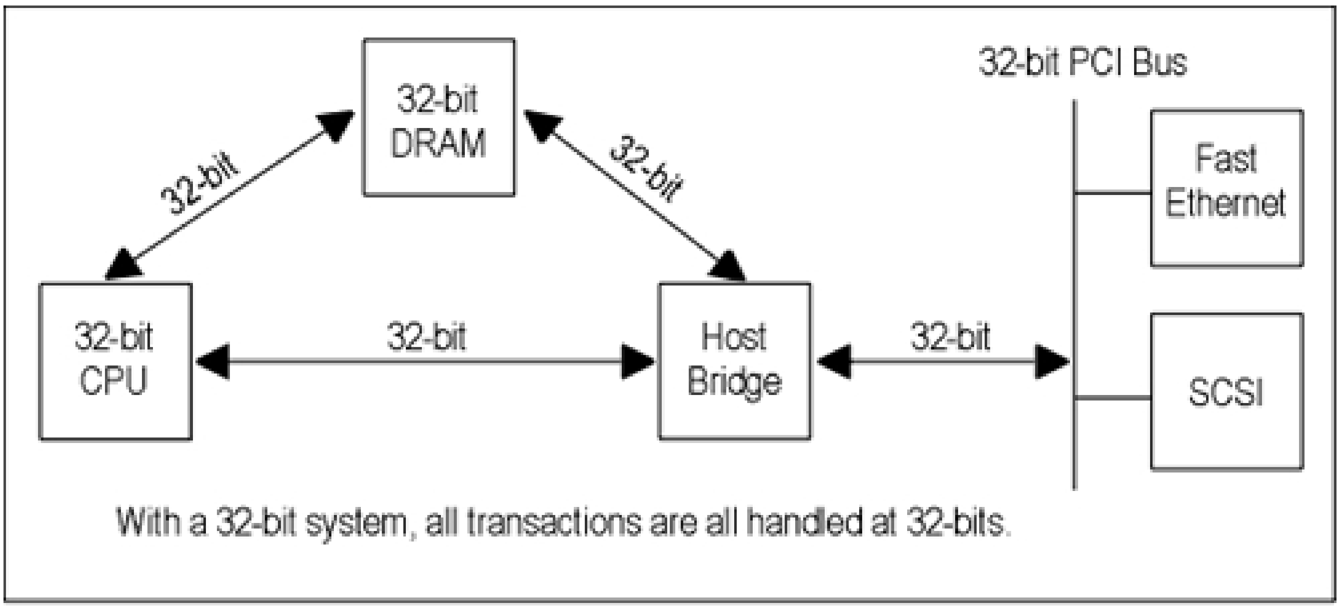

With a 32-bit system, all transactions are all handled at 32-bits.

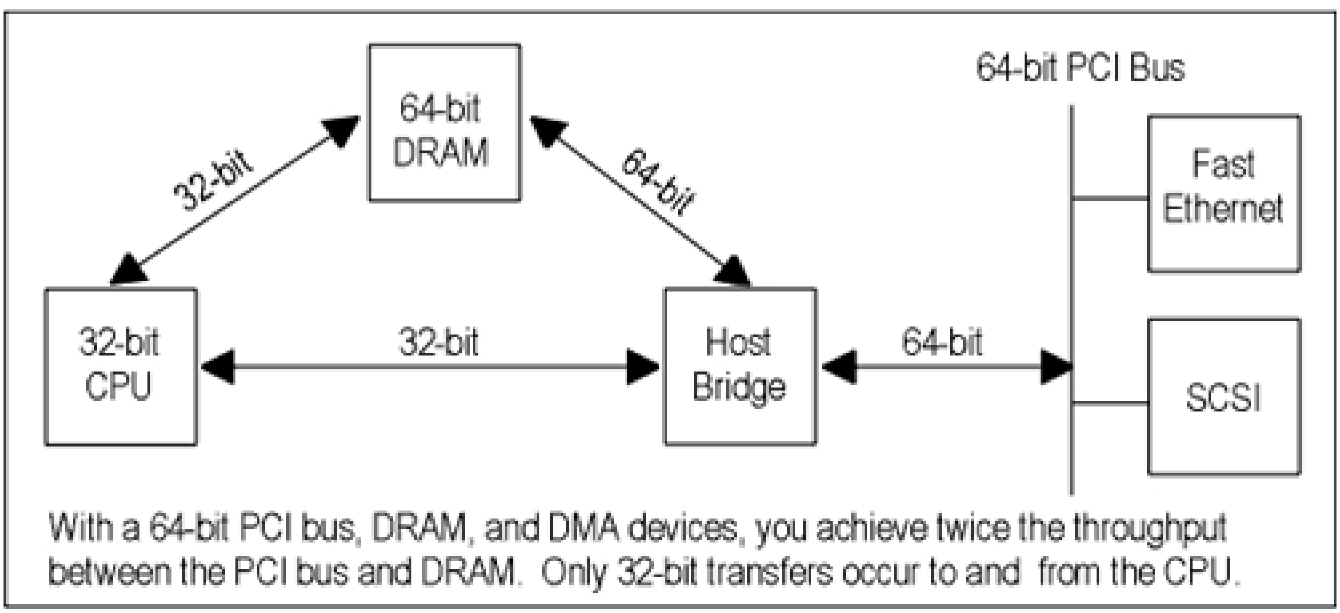

With a 64-bit PCI bus, DRAM, and DMA devices, you achieve twice the throughput between the PCI bus and DRAM. Only 32-bit transfers occur to and from the CPU.

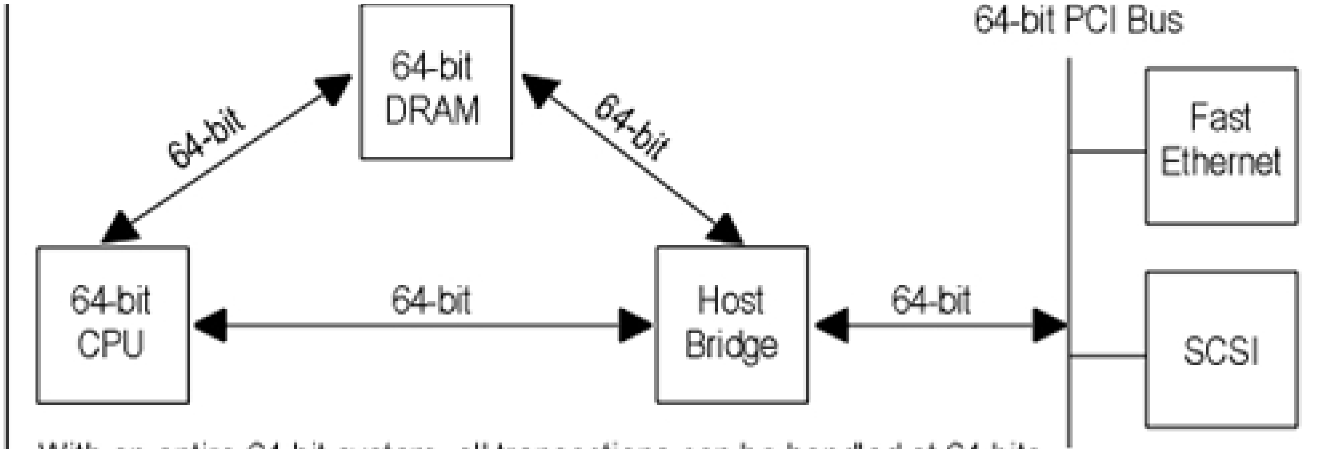

With an entire 64-bit system, all transactions can be handled at 64-bits.

PCI Bus Speeds, 33 MHz Versus 66 MHz

Another way of increasing the throughput of the PCI bus is increasing the PCI clock speed. PCI systems are now at 33 MHz; the PCI spec defines 66 MHz as a way of increasing PCI bandwidth. 66-MHz devices are great for high bandwidth applications and peripherals. Just as 64-bit architecture can double the bus bandwidth, 66 MHz can double the throughput. The diagram to the right illustrates the increase in throughput.

Theoretical Throughput

	32-bit PCI Bus	64-bit PCI Bus
33 Mhz	133 MB/sec	266 MB/sec
66 Mhz	266 MB/sec	532 MB/sec

Some hardware modifications need to be made to the PCI devices and motherboard to allow this increase in clock speed. 66-MHz devices are defined by modifying an existing ground pin to a static signal (M66EN) using a single pullup resistor and adding one bit to the Configuration Status register. 66-MHz PCI requires higher maximum clock frequency and modifying timing parameters. Engineers need to pay close attention to maximum trace lengths, loading of add-in boards, and the maximum pin capacitance of all add-in boards. A 66-MHz bus is capable of operating at 0 to 66 MHz speeds. 33-MHz devices operate at 33 MHz in a 66-MHz bus; likewise 66-MHz devices operate at 33 MHz in a 33-MHz bus.

When designing a 66-MHz PCI bus into the motherboard, electrical problems, chip size, and heat dissipation are some issues to deal with. The loading factor in the 66-MHz bus is significantly reduced from that of the 33 MHz. As the bus speed increases, the total distance data can travel decreases, and the number of slots available on the motherboard decreases. To solve the 66-MHz requirements on a motherboard, 66-MHz PCI bus is typically separated out on another bus using a 66-MHz PCI bridge chip.

64-Bit Market Trends

System vendors are looking for ways of creating the fastest and best systems on the market. Who are some of the players in this market segment? How do the OS, the CPU, and other system components affect system performance?

Vendors Implementing 64-Bit Technologies

CPUs -- Currently, Intel, AMD, and Cyrix CPUs manipulate data and addressing internally in 32 bits. HP, DEC Alpha, Sun Sparc, MIPs, NEC, and IBM RS64 processors operate internally in 64 bits. Intel plans to deliver their new 64-bit CPU, code named Merced, in late 1999. According to an article in SunWorld, October 1997, "HP AND INTEL OFFICIALS REITERATED THEIR CLAIM THAT 64-BIT PROCESSORS WILL BE FULLY BACKWARD-COMPATIBLE WITH 32-BIT APPLICATIONS AND OPERATING SYSTEMS."

PCI Chipsets -- Many of the PCI chipsets on the market support only 32-bit transfers. Some earlier PCI chipsets were only capable of supporting transfers of around 40 MB/sec even though the PCI bus is capable of 133 MB/sec. The sustainable transfer rates on the PCI chipsets should be closer to the theoretical throughput of the PCI bus as shown in the chart above. Newer PCI chipsets provide much more bandwidth and features. The advanced graphics port, for example, on Intel's 440LX PCI chipset is designed to remove the video demand from the PCI bus and improve overall system performance for non-video transactions--3D graphics go directly to memory and bypass the PCI bus. Some vendors will be providing 64-bit PCI chipsets by mid 1998.

OSes -- NetWare versions 3.12 through Moab use 32-bit addressing. Windows NT® 3.51, 4.0, and the new 5.0 all use 32-bit addressing. A native 64-bit Windows NT is planned for in 1999 when Intel releases its 64-bit CPU (It is 2001 now, this has not been released yet). Some UNIX OSes like HP UX 11.x, Solaris, and IBM AIX already provide support for 64-bit addressing while remaining backward compatible with 32-bit applications. SCO UnixWare also plans on releasing a 64-bit version. Linux currently supports 32 bit only, but 64 bit is being working on.

Since the OS plays a huge role in system performance by addressing and transferring data to and from memory, when the OS supports only 32-bits when the hardware supports 64 bits, the system is still somewhat limited. Changing the OS from 32-bit to 64-bit addressing is an enormous endeavor. The memory management schemes need to be rewritten, and data transfers and structures need to be set up differently. Even the device drivers need to be modified for 64-bit addressing.

Systems -- HP PA-RISC L Class and N Class are all equiped with 66Mhz 64 bit PCI Bus and 64 bit CPU and 64 bit OSes. DEC Alpha systems have 64-bit PCI systems with a 64-bit Alpha processor but the OS (Windows NT 4.0) supports 32 bits. Micron supplies 64-bit PCI systems using Intel's current 32-bit CPUs and 32-bit OSes. SGI supplies 64-bit PCI systems with 64-bit MIPS processor; Toshiba supplies their TX49 systems with 64-bit processors. IBM's RS/6000 S70 Enterprise servers offer full 64-bit processing, 64-bit PCI slots, and 64-bit OS support in AIX.

According to the article in SunWorld, "A NUMBER OF SYSTEMS MANUFACTURERS, INCLUDING HP, COMPAQ COMPUTER CORP., SEQUENT COMPUTER SYSTEMS INC. AND NCR, HAVE PLEDGED SUPPORT FOR THE 64-BIT ARCHITECTURE AND EXPECT TO ROLL OUT PRODUCTS WHEN MERCED APPEARS. IBM ALSO PLANS TO INCORPORATE THE 64-BIT PROCESSOR TECHNOLOGY IN ITS IBM INTELLISTATION AND IN ITS NEW IBM NETFINITY LINE OF PC SERVER PRODUCTS."

Modifying Drivers for 64-Bit OSes

64-bit data transfers from the device to the PCI bus will be transparent to the driver. The 64-bit PCI devices should autosense the PCI slot for 32 bit or 64 bit and operate accordingly. Since the driver interfaces with the OS, if the OS supports 32-bit addressing, the driver must support 32-bit addressing. Likewise, if the OS supports 64-bit addressing, the driver must support 64-bit addressing. To create a 64-bit driver, you must modify the DMA descriptors to address 64 bits of memory so the hardware can access the pointers to the data in this memory range.

Gigabit Ethernet Trends

HP, Alteon and Packet Engines are supporting 64-bit Gigabit Ethernet now. However, Gigabit Ethernet is slow in coming to the market because OS, CPU, and protocol support is lacking. We have tested HP's implementation of 64 bit PCI bus for their Gigabit Ethernet on HP/UX 11.0 with L Class box. We are glad to report the result are significant faster than some others' implementation. Gigabit Ethernet Specification is still in development and it's hard to commit to hardware and software development without a finalized specification. For Gigabit Ethernet and many other emerging technologies to realize their full potential, systems need to increase their overall performance and efficiency with 64-bit and 66-MHz PCI, and 64-bit OS and CPU support.

Overall System Benefits

A 64-bit PCI bus provides higher overall throughput for high-performance adapters and better system efficiency by providing the same data in fewer PCI clock cycles. A 66-MHz PCI bus doubles the data throughput over the same amount of time. The benefits of both 64-bit and 66-MHz PCI implementations are better PCI bus utilization, better overall PCI bus efficiency, and a substantial increase in PCI bus performance.

Intel conducted a test in 1999 for the Gigabit ethernet. Their test result shows the full Gigabit bandwidth can only be achieved by implementing the 66MHz 64bit PCI bus, as well as use CPU and memory that fast enough not causing bottleneck effect. We have conducted the test in Compute Aid, Inc.'s lab for the Intel Gigabit ethernet card and switch (based on chips made by HP), they work flawlessly.

Compute Aid Home • Order Now! • ATX Spec

ATX Chassis:

- **Why the computer/fan does not turn on after I flipped the power switch to ON located on the ATX power supply?**
  ATX power supply is different from AT power supply. It depends a logic circuit on the motherboard to turn it on. Once you flipped the power switch on the ATX power supply to on (some ATX power supplies do not even have such a switch), the ATX power supply sends a 5V 720MA current to the motherboard through pin 9 on the power connector. That current is for WOL (Wake-up On Lan) and power on circuits. There is a power-on jumper on the motherboard that connects to the pushbutton located in front of ATX case. When the pushbutton is pressed, it sends a signal to the mothboard, which in turn notifies the ATX power supply to turn on the full power. The location of the power-on jumper on the motherboard is manufacture dependent. You will need to read your motherboard manual to locate that jumper. **At any event, do not try to manually jump-start the power supply without attaching motherboard. Since the power supply is expecting certain sensing circuit feedback to regulate the output voltage, munaully starting it without attaching it to the motherboard could cause damage to the power supply.**

- **How do I test if the power supply is bad?**
  WarningThis test is for those who have extensive electronics training. Do not try this if you do not have extensive electronics training. Make sure there is enough load connected to the power supply before testing (at least motherboard and one hard drive)!

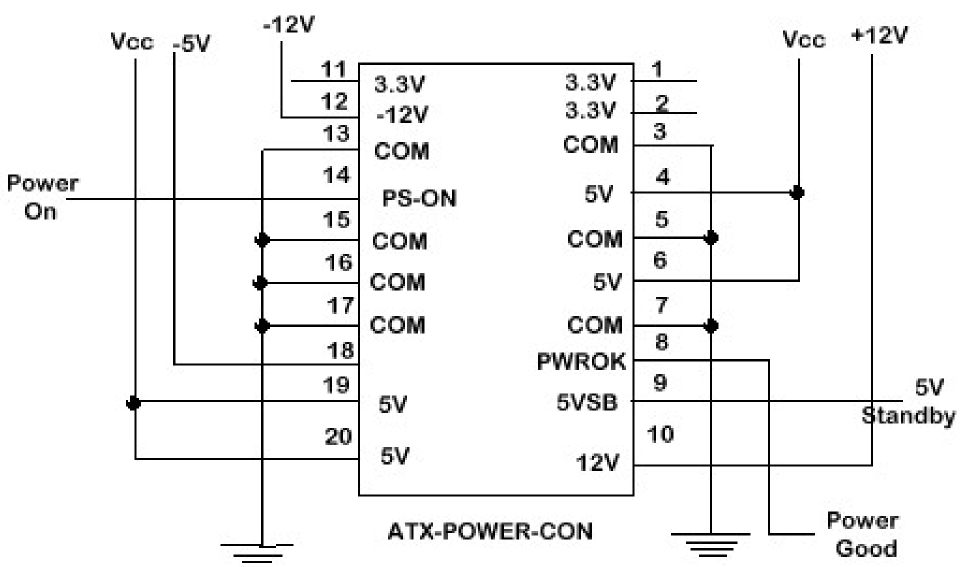

  ATX power connector is layout like above. 5VSB constantly provides 5V power to the connector through pin 9. If you have a voltmeter, while the power supply fan does not turn on, you could identify which pin is pin 9. Normally there is a clip on the connector between pin 15 and 16. If you can identify pin 14, which is power-on pin, you can using a piece of wire short the pin 14 and 15. If you power supply has power connected to it and power switch on the power supply is on, you will see the fan is turnning by now. Otherwise, your power supply may have problem.

- **What kind of power supply do I need to buy for AMD and Pentium 4?**
  To support the faster CPU speed, both AMD and Intel have new requirement to the power supply to be used for their new processors. The most important factor for selecting the correct power supply is the output DC current for the processors and motherboard. For most AMD processors, the 5V supply must provide at least 18A. For the new Intel P4 processors, the 5V must provide at least 20 amperes. Both AMD and Intel P4 also require the +5V standby can provide at least 720mA or more. Without suffcient 5V and 12V current from the power supply, the system may become unstable/unreliable and may experience difficulty to power up the system. The minimum wattage required for a Intel P4

system is 230W.
For server motherboard stability, Intel recently also developed the new power supply requirement for server board. The power supplies for the new Intel P4 server boards recommended to have following connectors:

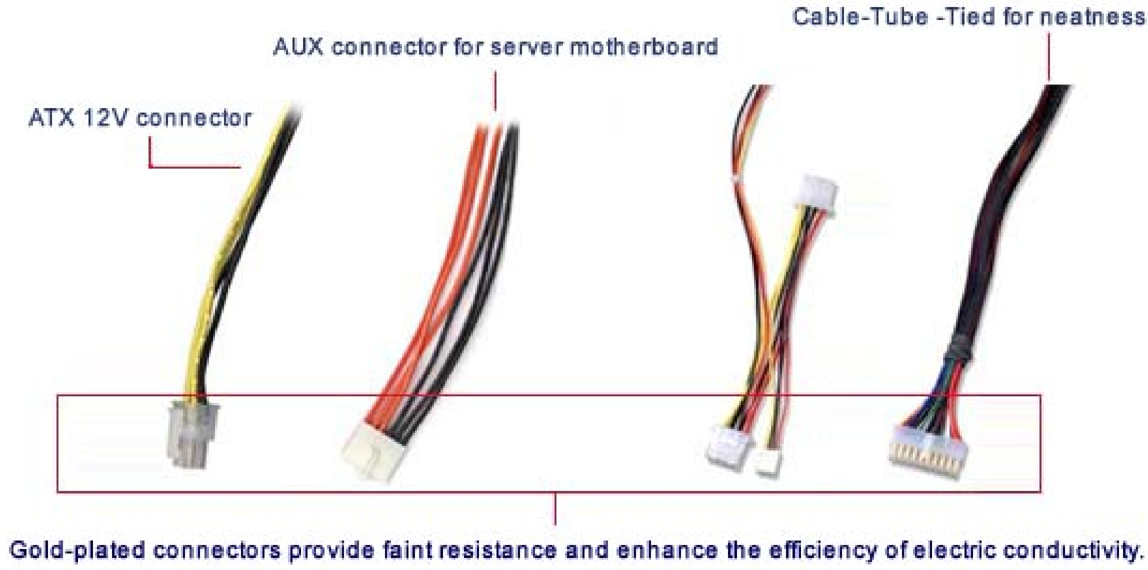

The purpose of those connectors are for the current to go directly from power supply to the point where the current draw the most, thus improve the server stability.

- **How to tell a power supply's real wattage?**
  Many cutomers emailed us asking this question. It is difficult to tell by looking or reading the labels. To be accurate, you will need an instrument to measure it. Roughly, you might be able to tell by the weight. Yes, by the weight! Faked 300W power supply has the similar weight as the 250W power supply. Real 300W poewr supply is much heavier than the 250W power supply, due to the heatsink is larger, and components must be able to handle larger current. To the manufacture, it will cost close to $28 USD marterial to make a 300W power supply, not including the labor and shipping. If someone selling a computer case including a 300W power supply totally less than $30, you know that is not real!

- **What is the difference between the ATX and the microATX (uATX) chassis?**
  Full size ATX motherboards may be integrated with a chassis that complies with the ATX 2.01 or later specification.

  The microATX specification was released in January 1998 and is a derivative of the ATX2.01 specification. The purpose behind drafting this specification was to provide a lower cost platform solution that was backward compatible to standard ATX2.01.

  The microATX motherboard may be integrated into either a full size ATX or a microATX chassis that meets the ATX2.01 or later specification. The full size ATX motherboard may be integrated into a full size ATX chassis and may not be integrated into a microATX chassis. More detailed information on both the ATX and microATX chassis is available on their dedicated web sites at http://www.teleport.com/~atx and http://www.teleport.com/~microatx.

  Note: *A thorough evaluation needs to be performed by the system manufacturer on the ATX or microATX (uATX) chassis chosen. The system manufacturer is responsible for ensuring that there is adequate airflow and cooling for the processor. Our cases are designed with all these possible situations in mind, so the cooling is never a problem. We also recommend all the systems installing the optional cooling fan so that it will never have overheat problem in the hot summer.*

- **Do I need a special chassis for microATX motherboards?**
  The microATX motherboards can be integrated into a micro tower chassis and powered by lower wattage power supplies. The microATX motherboards can also be integrated into the full ATX 2.01 (or later) compliant chassis like the one being sold on our web site.

Power Supply:

- **What special power supply requirements does the ATX motherboard have?**
  The ATX motherboard is designed to operate with at least a 145 Watt power supply for typical system configurations. A higher-wattage supply may be required for heavily-loaded configurations. The power supply must provide +5VSB (voltage

HP/UX related FAQ

SCSI LVD FAQ

stand by) with 720mA of current (see the ATX2.01 or later specification). If the power supply does not supply sufficient current, some system configurations with the motherboard may not power up. Additionally, if the power supply can not supply sufficient current, and does not have overload protection, the power supply may be damaged. That is why we pay close attention to the manufacture process of the power supply in our case so that it would not be a problem of the end users. If you are using AMD CPU chip, it is recommended to use AMD Certified power supply.

- **What power supply requirements does the ATX motherboard have if a WOL (Wake on LAN*) capable Network Interface Card (NIC) is installed?**
  A power supply that provides at least 720mA on 5VSB must be used when building a WOL capable system.

  If your power supply is not capable of meeting the 720 mA current requirement, your system may not power up. Also you could experience damage to your power supply if it does not have any over-current protection. Please contact your system integrator to see if your power supply meets the 720 mA current requirement. All of our power supplies are meet or exceed the ATX 2.01 requirement, so that it is never a problem for our customers.

  If you are using a WOL capable NIC but are not using the WOL capability you can remove the cable from the NIC to the WOL header on the motherboard. This will prevent the NIC from drawing additional current from the 5VSB.

System Airflow

This section explains why our cases are better than many the "fancy" looking ones out there. The heat problem with those fancy ones will kill their motherboards. Our case will allow heat exhausted from the case easier. In addition, we always offer the system fan option to allow the system running in the room temperature without special air conditioning cooling need.

System airflow is determined by:

- Chassis design
- Chassis size
- Location of chassis air intake and intake vents
- Power supply fan capacity and venting
- Location of processor(s) slots
- Placement of add-in cards and cables

System integrators must ensure airflow through the system to allow the fan heatsink to work effectively. Proper attention to airflow when selecting subassemblies and building PCs is important for good thermal management and reliable system operation.

Integrators use three basic motherboard-chassis-power supply form factors for desktop systems: ATX, microATX, and the older Baby AT form factor.

Intel recommends the use of ATX form factor motherboards and chassis for the boxed Pentium II Processor. The ATX form factor simplifies assembly and upgrading of PCs, while improving the consistency of airflow to the processor. With regard to thermal management, ATX components differ from Baby AT components in that the processor slot is located close to the power supply, rather than to the front panel of the chassis. Power supplies that blow air out of the chassis provide proper airflow for active fan heatsinks. The boxed processor's active fan heatsink cools the processor much more effectively when combined with an exhausting power supply fan. Because of this, the airflow in systems using the boxed processor should flow from the front of the chassis, directly across the motherboard and processor, and out of the power supply exhaust vents. Figure 1 shows proper airflow through an ATX system to achieve the most effective cooling for a boxed processor with an active fan heatsink. For the boxed Pentium II Processor, chassis that conform to the ATX Specification Revision 2.01 or later are highly recommended.

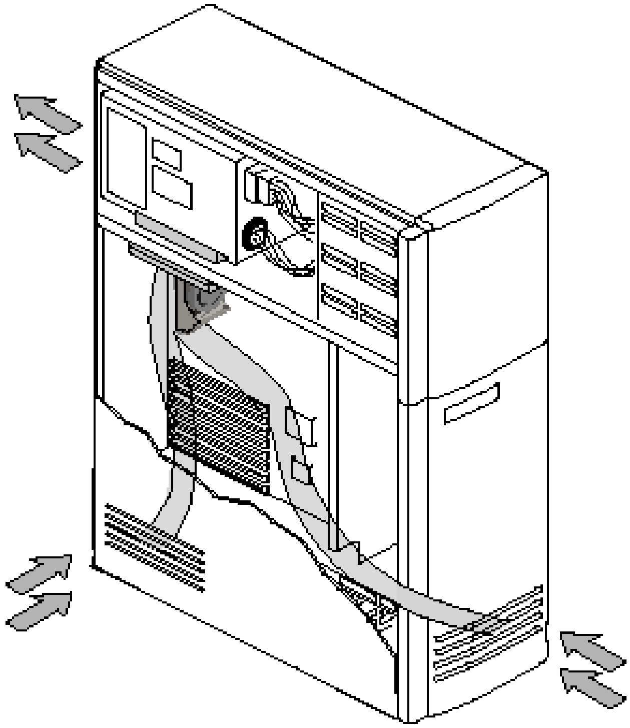

Figure 1. System Airflow through ATX Tower Chassis Optimized For the Boxed Processor With an Active Fan Heatsink

MicroATX chassis differ from most ATX chassis in that the power supply locations vary and they may use SFX or ATX power supplies. Thermal management improvements that apply to ATX chassis will also apply to microATX. Our Mini Tower case fit in the category of microATX and satisfy Intel's spec. Our middle and full ATX case all tested and satisfy the Intel ATX 2.01/2.03 spec.

The following is a list of guidelines to be used when integrating a system. Specific mention of Baby AT, ATX, or microATX components is made where necessary.

- **Provide sufficient air vents**: Systems must have adequate air vents in addition to a fan. Chassis vents must be fully functional. Integrators should be careful not to select chassis that contain cosmetic vents only. Proper location of vents results in a good stream of air flowing over the processor. For Baby AT systems, intake vents on the front of the chassis allow air to flow over the processor. For ATX and microATX systems, exhaust vents in the chassis allow air already forced over the processor (by the power supply) to flow out of the chassis.

- **Power supply air flow direction**: It is important to choose a power supply with a fan that moves air in the proper direction. For Baby AT systems, the power supply fan acts as an exhaust fan, venting system air outside the chassis.

- **Power supply fan strength**: For some chassis that are running too warm, changing to a power supply with a stronger fan can greatly improve airflow. If that is not an option, add an additional system fan.

- **System fan--should it be used?** Some chassis may contain a system fan to assist airflow. A system fan is typically used with passive heatsinks. With fan heatsinks, however, a system fan can have mixed results. Thermal testing both with a system fan and without the fan will reveal which configuration is best for a specific chassis. When a fan heatsink is used on the processor, changing to a power supply with a stronger fan is usually a better choice than adding a system fan.

- **System fan airflow direction:** When using a system fan, ensure that it moves air in the same direction as the overall system airflow. For example, a system fan in a Baby AT system should act as an intake fan, pulling in additional air from the front chassis vents.

- **Protect Against Hot Spots**: A system may have a strong airflow, but still contain "hot spots." Hot spots are areas within the chassis that are significantly warmer than the rest of the chassis air. Such areas can be created by improper positioning of the exhaust fan, adapter cards, cables, or chassis brackets and subassemblies blocking the airflow within the system. To avoid hot spots, place exhaust fans as needed, reposition full-length adapter cards or use half-length cards, reroute and tie cables, and ensure space is provided around and over the processor.

Keep Power On or Off?

People debate a lot on this topic. Lets take a look what is good for turning system off everyday. First, that saves electricity and saves the environment; Secondly, it reduces the wear on the fans, hard disk and other moving parts, reduces the electronics components aging. What is good about keep computer on all the time, then? It can save time, it is on whenever you like to use it.

If your computer stays in a home or an office that maintaining close to constant temperature, preferably under 80 degrees, there is no harm done to turn it off whenever you finish using it. However, if the computer stays in a place that temperature varies a lot, like in a warehouse that could go up to 110 degree during the day and goes down to 60 degree in the night, it is better to keep the computer up and running. In that kind of situation, the temperature creates a lot of tension on the electronics components in the power supply or motherboard, having the computer on all the time helps maintaining the components not being broken apart by the mechanical force(heat expansion and cold contraction).

If you decide to keep your computer on all the time, make sure check all the fans every 3 to six months. Fans tend to wear out quickly in that kind of situation. If fan locked up and not replaced, the power supply or motherboard will overheat and quit working shortly after. It is better to replace the fan before it worn out completely. Choosing ball bearing fans and industrial chassis will also help the system last longer.

Order Online • Server Tip • Browser Security Check

What is Ultra2 SCSI?

Ultra2 SCSI or Low Voltage Differential (LVD) is a highly compatible computer disk drive interface that is faster and more reliable than previous SCSI standards. This upgraded SCSI interface helps meet the growing need for faster data rates and is available in 68- and 80-pin configurations.

Servers, workstations, RAID subsystems, Internet technologies, as well as CAD/CAM, multimedia, video, digital broadcasting, and groupware applications all require a more advanced interface to handle increasing increased data-transfer needs. The Ultra2 SCSI (LVD) interface satisfies these needs by increasing bus bandwidths, allowing greater configuration flexibility, backward compatibility, and faster transfer rates.

ATX Power Supply FAQ

HP/UX Related FAQ

How would users benefit from having an Ultra2 SCSI (LVD) disk drive in their computer systems?

Not so long ago, servers, RAID subsystems, CAD/CAM, and groupware applications were simply buzzwords. Today, these new digital technologies require increased speed, storage capacity, and reliability for critical data applications. By doubling previous SCSI standards for bus data rates and bandwidths, the Ultra2 SCSI (LVD) interface provides a much needed upgrade for high-end computer users.

Ultra2 SCSI advantages

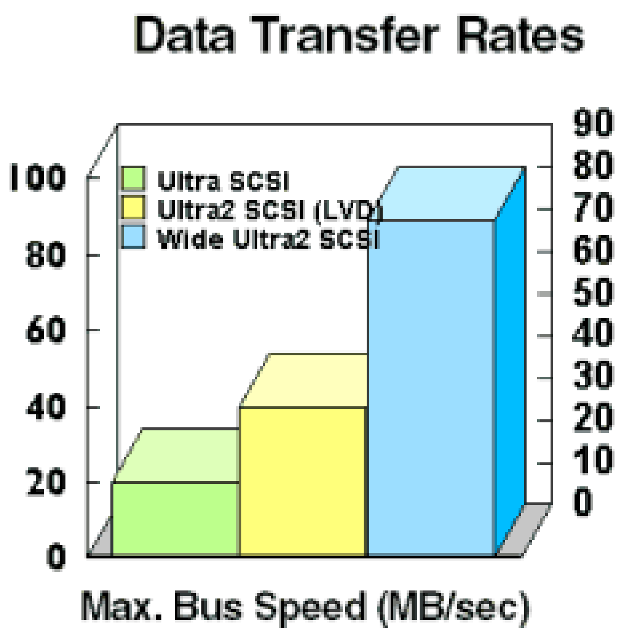

The Ultra2 SCSI (LVD) provides four major advantages:

- Increased bus data rates
- Increased device connectivity
- Increased design flexibility
- Backward compatibility with Ultra SCSI

Faster bus data rates

In the world of digital video and complex RAID environments, information must be channeled more quickly and efficiently. Ultra2 SCSI (LVD) increases the bus data rate to 80 MB/sec in 16-bit SCSI mode while lowering overall power consumption. This advancement allows the bus to accommodate more devices.

Increased device connectivity

Ultra2 SCSI (LVD) allows more drives and devices to be connected to a user's system. Because Ultra2 SCSI (LVD) provides higher bus data rates, it is capable of connecting to higher-quality peripherals. Since the Internet and other server applications require large volumes of data to be transferred and stored, systems need larger and faster hard drives. Ultra2 SCSI (LVD) can connect up to 15 peripherals without creating a bottleneck. Therefore, not only can a user connect more devices, but the speed and reliability of the information also increases.

Increased design flexibility
Providing extended cable lengths of up to 12 meters and transfer rates up to 80 MB/sec, Ultra2 SCSI (LVD) is a more flexible design than previous SCSI interfaces.

Greater compatibility
The Ultra2 SCSI (LVD) interface provides for 15-year backward compatibility with older SCSI controller cards, enabling users to upgrade to new drives. This backward compatibility makes the Ultra2 SCSI (LVD) upgrade low risk and allows peripherals to be added without requiring new controller cards.

Summary

Ultra2 SCSI (LVD) interface technology provides supporting data transfer rates up to 80 MB/sec, greater connectivity of devices, and greater design flexibility. The 15-year backward compatibility of the Ultra2 SCSI (LVD) interface helps minimize upgrade risks and costs. Ultra2 SCSI (LVD) offers a faster, more reliable, disk drive interface satisfying the quality and performance standards of server hard disk drive product line.

The development of the SCSI standard

	SCSI	Ultra SCSI Differential	Wide Ultra SCSI	Wide Ultra SCSI Differential	**Ultra2 SCSI (LVD)**	**Wide Ultra2 SCSI (LVD)**
Data transfer rates max. Bus speed (MB/sec)	20	20	40	40	**40**	**80**
Data Bus width (bits)	8	8	16	16	**8**	**16**
Max cable length	1.5-3	25	1.5-3	25	**12**	**12**
Max device support	8-4	8	8-4	16	**8**	**16**

There is an inverse relationship in all SCSI interfaces between the number of peripherals and the amount of cable length needed for support, i.e., the more peripherals connected, the less cable length required and vice/versa.

Note: Without the presence of an Ultra2 SCSI (LVD) controller, the Ultra2 SCSI (LVD) behaves as a standard SCSI or Ultra SCSI device, limiting its performance benefits.