Glossary of Terms

- AMI Searchable Knowledgebase
- Search by Product
- Document Library
- Download Center
- Support Request Form
- AMI Contact Phone Numbers
- Warranty Information
- Glossary of Terms

- RMA Request Form
- Product Registration

HOME : Tech Support : Glossary of Terms

Glossary of Terms

0-9 A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

Access Time
The required time to read or write data to RAM or other storage device. Since the operating environment and varying conditions affect access time, this is usually given as an average.

ACPI (Advanced Configuration and Power Interface)
An industry specification for the efficient handling of power consumption in desktop and mobile computers. ACPI specifies how a computer's BIOS, operating system and peripherals communicate with each other about power usage. ACPI allows the following capabilities (assuming the operating system supports them):

- The user can specify at what time a device, such as a display monitor, is to turn off or on.
- The user of a notebook computer can specify a lower-level of power consumption when the battery starts running low so that essential applications can still be used while other, less important applications are allowed to become inactive.
- The operating system can lower the clock speed during times when applications don't require the full processor clock speed.
- The operating system can reduce motherboard and peripheral device power needs by not activating devices until they are needed.
- The computer can enter a stand-by mode when no one is using it, but with modem power left on to receive incoming faxes.
- Devices can be plug-and-play. As soon as plugged in, they can be controlled by ACPI.

ACPI must be supported by the computer's motherboard, BIOS and the operating system. One of several power schemes can be chosen. Within a power scheme, the user can control the power to individual devices. In order for ACPI to work on your computer, your BIOS must include the ACPI software and the operating system must be ACPI-compatible. ACPI is designed to work with Windows 98 and with Windows 2000.

ACK
In serial port modem communications, A control code (06h) sent to a sending station or computer by the receiving unit to acknowledge either that the receiver is ready to accept transmissions or that transmitted data arrived without error. The ability to receive and send acknowledgment signals is built into the hardware and software. For example, the serial ports send and receive ACK commands.

Adapter
The term used to describe expansion cards that are inserted into bus expansion slots.

Adapter Card

A flat rectangular fiberglass board with electronic circuitry. Inserted in an expansion slot on the computer main bus, it provides additional system functions, such as device controllers or video adapters. Also called an Add-In card.

Address
Every memory location is numbered consecutively. This number is the address of the memory location. An address can be a label, number, or name that identifies a register, memory location, or a location on a disk drive or external device accessed via an I/O port.

AGP (Advanced Graphics Port)
An interface specification that enables 3-D graphics to display quickly on ordinary personal computers. AGP is an interface designed to convey 3-D images (for example, from Web sites or CD-ROMs) much more quickly and smoothly than is possible today on any computer other than an expensive graphics workstation. The interface uses your computer's main storage (RAM) for refreshing the monitor image and to support the texture mapping, z-buffering and alpha blending required for 3-D image display. The AGP main memory use is dynamic, meaning that when not being used for accelerated graphics, main memory is restored for use by the operating system or other applications.

Analog
A term used to describe any device that represents values by a continuously varied physical property, such as voltage.

ANSI - American National Standards Institute
A standards-setting, non-government organization that develops and publishes standards for voluntary use in the USA. API Attachment Packet Interface. A standard hard disk drive specification used for Integrated Drive Electronics (IDE) devices.

Arbitration
A process where devices compete for possession of the channel on a prioritized basis.

Architecture
Refers to the way a system is designed and how the components are connected with each other. There are computer architectures, network architectures and software architectures.

Array
An array of disk drives combines the storage space on the disk drives into a single segment of contiguous storage space. MegaRAID can group disk drives on one or more SCSI channels into an array. A hot spare drive does not participate in an array.

Array Management Software
Software that provides common control and management for a disk array. Array Management Software most often executes in a disk adapter or intelligent host bus adapter, but can also execute in a host server. When it executes in a disk adapter or adapter, Array Management Software is often called firmware.

Array Spanning
Array spanning by a logical drive combines storage space in two arrays of disk drives into a single, contiguous storage space in a logical drive. MegaRAID logical drives can span consecutively numbered arrays that each consist of the same number of disk drives. Array spanning promotes RAID levels 1, 3, and 5 to RAID levels 10, 30 and 50, respectively. ASCII American Standard Code for Information Interchange. An industry standard 7-bit code consisting of control, character, and graphic codes (8 bits if the parity bit is included).

Asynchronous Data Transfer
Data transfer (usually at a low rate and independent of any external timing constraints) performed by a SCSI device involving the interlocking of a signal to the initiator (REQ) and a signal to the target (ACK) such that each step of

the data transfer protocol must occur before the next step can begin.

Asynchronous Event Notification
A process when a SCSI target can send unsolicited sense information to an initiator using the SCSI SEND command.

Asynchronous Operations
Those operations that bear no relationship to each other in time and can overlap. The concept of asynchronous I/O operations is central to independent access arrays in throughput-intensive applications.

AT
Advanced Technology. Specifically refers to the IBM PC AT incorporating the Intel 80286 processor. It is also used as a reference of BIOS compatibility. AT refers to the original IBM PC/AT computer architecture, more commonly known as ISA now.

ATA - AT Attachment
A disk drive implementation that integrates the controller on the disk drive itself. There are several versions of ATA, all developed by the Small Form Factor (SFF) Committee:

- **ATA:** Known also as IDE, supports one or two hard drives, a 16-bit interface and PIO modes 0, 1 and 2.
- **ATA-2:** Supports faster PIO modes (3 and 4) and multiword DMA modes (1 and 2). Also supports logical block addressing (LBA) and block transfers. ATA-2 is marketed as Fast ATA and Enhanced IDE (EIDE).
- **ATA-3:** Minor revision to ATA-2.
- **Ultra-ATA:** Also called Ultra-DMA, ATA-33, and DMA-33, supports multiword DMA mode 3 running at 33 MBps.
- **ATA/66:** Also called Ultra-DMA/66 or UDMA-66. A new version of ATA proposed by Quantum Corporation, and supported by Intel, that doubles ATA's throughput to 66 MBps.
- **ATA/100:** Also called Ultra-DMA/100 or UDMA-100. A new version of ATA proposed by Quantum Corporation, and supported by Intel, that doubles ATA's throughput to 100 MBps.

ATAPI (AT Attachment Packet Interface)
An interface between your computer and attached CD-ROM drives and tape backup drives. Most of today's PC computers use the standard IDE (Integrated Drive Electronics) interface to address hard disk drives. ATAPI provides the additional commands needed for controlling a CD-ROM player or tape backup so that your computer can use the IDE interface and controllers to control these relatively newer device types. ATAPI is part of the Enhanced IDE (EIDE) interface (also known as ATA-2).

ATASPI - ATA Software Programming Interface
A specification for writing enhanced IDE drivers. It specifies a standard low-level interface for all enhanced IDE functions. The ATASPI driver acts as an I/O manager and will provide arbitration and handshaking that will allow two different types of peripheral devices to be attached to a single enhanced IDE connector.

ATX
ATX is an industry-wide open specification for a desktop computer's motherboard. ATX improves the motherboard design by taking the small AT motherboard that has been an industry standard and rotating by 90 degrees the layout of the microprocessor and expansion slots. This allows space for more full-length add-in cards. A double-height aperture is specified for the rear of the chassis, allowing more possible I/O arrangements for a variety of devices such as TV input and output, LAN connection and so forth. The new layout is also intended to be less costly to manufacture. Fewer cables will be needed. The power supply has a side-mounted fan, allowing direct cooling of

the processor and cards, making a secondary fan unnecessary. Version 2.0 incorporates improvements suggested by chassis and power supply vendors.

Baud
A unit of measurement of the discrete number of signal elements that can be transmitted per second by a device. It is not an exact measure of the amount of information being transmitted and is not the same as Bits Per Second.

Binary
The base two numbering system, where the only digits are 0 and 1. It is used by all computers.

Binary Coded Decimal
A method of encoding decimal digits into four binary bits.

BIOS (Basic Input/output system)
The program a personal computer's microprocessor uses to get the computer system started after you turn it on. It also manages data flow between the computer's OS and attached devices such as the hard disk, video adapter, keyboard, mouse and printer.

BIOS is an integral part of the computer and comes with it when you bring it home. (In contrast, the OS can either be preinstalled by the manufacturer or vendor or installed by the user.) BIOS is made accessible to the microprocessor on an erasable programmable read-only memory (EPROM) chip. When you turn on your computer, the microprocessor passes control to the BIOS program, which is always located at the same place on EPROM.

When BIOS boots up the computer, it first determines whether all of the attachments are in place and operational and then it loads the OS (or key parts of it) into a computer's random access memory (RAM) from a hard disk or diskette drive.

With BIOS, an OS and its applications are□ freed from having to understand exact details (such as hardware addresses) about the attached input/output devices. When device details change, only the BIOS program needs to be changed. Sometimes this change can be made during system setup. In any case, either an OS or any applications used needs to be changed.

Although BIOS is theoretically always the intermediary between the microprocessor and I/O device control information and data flow, in some cases, BIOS can arrange for data to flow directly to memory from devices (such as video cards) that require faster data flow to be effective.

Bit
A binary digit that can take either the value 0 or 1. A bit is the smallest unit of information that a computer can process.

Block
An amount of data moved or addressed as a single unit; the smallest amount of data that can be read or written at a time. Blocks are separated by physical gaps, or identified by their track/sector addresses or logical addresses.

Boot
The process of initializing, testing and configuring a computer system blding higher-level services on top of lower-level primitive services.

BPS - Bits per Second
The number of binary digits that can be transmitted in one second. Generally, modem speeds are given in BPS, not baud rate. Neither baud rate nor BPS take into account the gaps between transmissions, so neither baud rate nor BPS accurately express the amount of information being transferred.

Buffer
An area of memory or storage that is temporarily reserved for I/O processing.

Burst Mode
A method of data transfer that allows a device to remain inactive for long periods of time and then send large amounts of data in a short time without

interruption. Can be used for DMA transfers on the EISA bus.

Bus
(1) A collection of wires through which data is transmitted from one part of a computer to another. You can think of a bus as a highway on which data travels within a computer. When used in reference to personal computers, the term bus usually refers to internal bus. This is a bus that connects all the internal computer components to the CPU and main memory. There's also an expansion bus that enables expansion boards to access the CPU and memory. All buses consist of two parts -- an address bus and a data bus. The data bus transfers actual data whereas the address bus transfers information about where the data should go. The size of a bus, known as its width, is important because it determines how much data can be transmitted at one time. For example, a 16-bit bus can transmit 16 bits of data, whereas a 32-bit bus can transmit 32 bits of data. Every bus has a clock speed measured in MHz. A fast bus allows data to be transferred faster, which makes applications run faster. On PCs, the old ISA bus is being replaced by faster buses such as PCI. Nearly all PCs made today include a local bus for data that requires especially fast transfer speeds, such as video data. The local bus is a high-speed pathway that connects directly to the processor. Several different types of buses are used on Apple Macintosh computers. Older Macs use a bus called NuBus, but newer ones use PCI.

(2) In networking, a bus is a central cable that connects all devices on a local-area network (LAN). It is also called the backbone.

Bus Free Phase
The phase when no SCSI device is actively using the SCSI bus and the bus is available for use.

Bus Mastering
Refers to a feature supported by some bus architectures that enables a controller connected to the bus to communicate directly with other devices on the bus without going through the CPU. Most modern bus architectures, including PCI, support bus mastering because it improves performance.

Byte
A unit of data made up of eight contiguous bits. A byte is usually the smallest addressable unit of memory.

Cache
Pronounced cash, a special high-speed storage mechanism. It can be either a reserved section of main memory or an independent high-speed storage device. Two types of caching are commonly used in personal computers: memory caching and disk caching. A memory cache, sometimes called a cache store or RAM cache, is a portion of memory made of high-speed static RAM (SRAM) instead of the slower and cheaper dynamic RAM (DRAM) used for main memory. Memory caching is effective because most programs access the same data or instructions over and over. By keeping as much of this information as possible in SRAM, the computer avoids accessing the slower DRAM. Some memory caches are built into the architecture of microprocessors. The Intel 80486 microprocessor, for example, contains an 8K memory cache, and the Pentium has a 16K cache. Such internal caches are often called Level 1 (L1) caches. Most modern PCs also come with external cache memory, called Level 2 (L2) caches. These caches sit between the CPU and the DRAM. Like L1 caches, L2 caches are composed of SRAM but they are much larger. Disk caching works under the same principle as memory caching, but instead of using high-speed SRAM, a disk cache uses conventional main memory. The most recently accessed data from the disk (as well as adjacent sectors) is stored in a memory buffer. When a program needs to access data from the disk, it first checks the disk cache to see if the data is there. Disk caching can dramatically improve the performance of applications, because accessing a byte of data in RAM can be thousands of times faster than accessing a byte on a hard disk. When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate. Many cache systems use a technique known as smart caching, in which the system can recognize certain types of frequently used data. The strategies for determining which information should be kept in the cache constitute some of the more interesting problems in computer science.

Coercion of Drives
The ability to coerce a drive into taking on a specific size. At AMI, drive coercion refers to the ability of our controllers to recognize the sizes of the drives connected, and to then coerce the larger drives to use only that amount of space which the smallest has available. So, if there is a 10GB drive connected with a 15GB drive, and they are combined in a RAID configuration, the 15GB drive will act as a 10GB drive, and will only have 10GB worth of usable space.

Controller
A device that controls the transfer of data from a computer to a peripheral device and vice versa. For example, disk drives, display screens, keyboards, and printers all require controllers. In personal computers, the controllers are often single chips. When you purchase a computer, it comes with all the necessary controllers for standard components, such as the display screen, keyboard, and disk drives. If you attach additional devices, however, you may need to insert new controllers that come on expansion boards. Controllers must be designed to communicate with the computer's expansion bus. There are three standard bus architectures for PCs -- the AT bus, PCI (Peripheral Component Interconnect), and SCSI. When you purchase a controller, therefore, you must ensure that it conforms to the bus architecture that your computer uses.

Clustering
Connecting two or more computers together in such a way that they behave like a single computer. Clustering is used for parallel processing, for load balancing and for fault tolerance. Clustering is a popular strategy for implementing parallel processing applications because it enables companies to leverage the investment already made in PCs and workstations. In addition, it's relatively easy to add new CPUs simply by adding a new PC to the network. Microsoft's clustering solution for Windows NT systems is called MSCS.

CPU (Central Processing Unit)
Is an older term for processor and microprocessor, the central unit in a computer containing the logic circuitry that performs the instructions of a computer's programs.

DIMM (Dual In-Line Memory Module)
A double SIMM (Single In-Line Memory Module). Like a SIMM, it's a module containing one or several random access memory (RAM) chips on a small circuit board with pins that connect it to the computer motherboard. A SIMM typically has a 32 data bit (36 bits counting parity bits) path to the computer that requires a 72-pin connector. For synchronous dynamic RAM (SDRAM) chips, which have a 64 data bit connection to the computer, SIMMs must be installed in in-line pairs (since each supports a 32 bit path). A single DIMM can be used instead. A DIMM has a 168-pin connector and supports 64-bit data transfer. It is considered likely that future computers will standardize on the DIMM.

Disk Mirroring
Also known as RAID Level 1. A form of RAID in which the Array Management Function maintains two or more identical copies of data on separate disks. Also known as RAID Level 1 and disk shadowing.

Disk Spanning
Disk spanning allows multiple disk drives to function like one big drive. Spanning not only overcomes disk space shortage; it also simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 400 MB disk drives can be combined to appear to the operating system as one single 1600 MB drive. Spanning alone does not provide reliability or performance enhancements. Spanned logical drives must have the same stripe size and must be contiguous. For example Logical Drives 1 and 2 can be spanned; Logical Drives 1 and 3 cannot.

Disk Striping
Also known as RAID Level 0. A mapping technique in which fixed-size consecutive ranges of virtual disk data addresses are mapped to successive

array members in a cyclical pattern.

DMI (Desktop Management Interface)
An industry framework for managing and keeping track of hardware and software components in a system of personal computers from a central location. DMI was created by the Desktop Management Task Force (DMTF) to automate system management and is particularly beneficial in a network computing environment where dozens or more computers are managed. DMI is hardware and operating system-independent, independent of specific management protocols, easy for vendors to adopt, mappable to existing management protocols such as the Simple Network Management Protocol (SNMP) and used on network and non-network computers. DMI consists of four components:

- **Management Information Format (MIF):**A text file that contains specific information about the hardware and software being used on a computer. Consists of one or more groups containing attributes that describe each component. By default, each MIF file contains the standard component ID group. This group contains the product name, version, serial number and the time and date of the last installation. The ID number is assigned based on when the component was installed in relation to other components. Manufacturers can create their own MIFs specific to a component.
- **Service layer:**A memory-resident code that acts as a mediator for the management interface and the component interface and allows management and component software to access MIF files in the MIF database. Is available as an operating system add-on and is a shared resource for all programs. Because the service layer must run all the time, it is designed not to use a lot of memory. The service layer also includes a common interface called the local agent, which is used to manage individual components.
- **Component interface (CI):**An application program interface (API) that sends status information to the appropriate MIF file via the service layer. Commands include the Get and Set command that modifies the MIF as needed and the Event command that notifies management software of critical events.
- **Management interface (MI):**The management software communicates with the service layer using the MI application program interface. The MI allows administrators to issue the Get and Set command and the List command that lists all the DMI-manageable devices.

Drive Roaming
When a drive is moved from one placement on a bus to another, it is said to have roamed.

Driver
A program that extends the capabilities of a computer by enabling the computer to operate peripheral devices, such as WORM drives, CD-ROM drives, or tape drives.

ECC (Error Correction [or Correcting] Code; Error Checking and Correcting)
Allows data that is being read or transmitted to be checked for errors and, when necessary, corrected "on the fly." It differs from parity-checking in that errors are not only detected but also corrected. ECC is increasingly being designed into data storage and transmission hardware as data rates (and therefore error rates) increase.

EDO - Extended Data Output Dynamic Random Access Memory
A type of DRAM that is faster than conventional DRAM. Unlike conventional DRAM which can only access one block of data at a time, EDO RAM can start fetching the next block of memory at the same time that it sends the previous block to the CPU.

EISA - Extended Industry Standard Architecture
A bus architecture designed for PCs using an Intel 80386, 80486, or Pentium microprocessor. EISA buses are 32 bits wide and support multiprocessing. The EISA bus was designed by nine IBM competitors (sometimes called the Gang of Nine): AST Research, Compaq Computer, Epson, Hewlett-Packard, NEC, Olivetti, Tandy, WYSE, and Zenith Data Systems. They designed the architecture to compete with IBM's own high-speed bus architecture called the Micro Channel architecture (MCA). The principal difference between EISA and MCA is that EISA is backward compatible with the ISA bus (see above), while MCA is not. This means that computers with an EISA bus can use new EISA expansion cards as well as old AT expansion cards. Computers with an MCA bus can use only MCA expansion cards. EISA and MCA are not compatible with each other. This means that the type of bus in your computer determines which expansion cards you can install. Neither EISA nor MCA has been very successful. Instead, a new technology called local bus (PCI) is being used in combination with the old ISA bus.

Fibre / Fibre channel
A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers and now being standardized by ANSI. The most prominent Fibre Channel standard is Fibre Channel Arbitrated Loop (FC-AL). FC-AL was designed for new mass storage devices and other peripheral devices that require very high bandwidth. Using optical fiber to connect devices, FC-AL supports full-duplex data transfer rates of 100MBps. FC-AL is compatible with, and is expected to eventually replace, SCSI for high-performance storage systems.

Firmware
Programming inserted into programmable read-only memory (PROM), thus becoming a permanent part of a computing device. Firmware is created and tested like software (using microcode simulation). When ready, it can be distributed like other software and, using a special user interface, installed in the programmable read-only memory by the user. Is sometimes distributed for printers, modems and other computer devices.

Hot Plug (Hot Swap)
The ability to add and remove devices to a computer while the computer is running and have the operating system automatically recognize the change. Two new external bus standards -- Universal Serial Bus (USB) and IEEE 1394 -- support hot plugging. This is also a feature of PCMCIA. Hot plugging is also called hot swapping.

I20 - Intelligent I/O
A fairly new I/O architecture developed by a consortium of computer companies called the I2O special Interest Group (SIG). I2O is designed to eliminate I/O bottlenecks by utilizing special I/O processors (IOPs) that handle the nitty gritty details of interrupt handling, buffering and data transfer. In addition, an I2O driver consists of an OS-specific module (OSM) that deals with higher-level operating system details (such as accessing files) and a hardware device module (HDM), that understands how to communicate with specific devices. Because the OSM and HDM are autonomous, they can perform a number of tasks independently, without sending data over the I/O bus. I2O is designed to work with PCI.

IA-64
Merced is the code name for a new 64-bit microprocessor from Intel that will begin to appear in new workstations and enterprise servers over the next few years. It's the first of Intel's IA-64 series and, because of its greatly increased I/O bandwidth relative to today's 32-bit microprocessors, it will make possible visual computing or the ability to interact dynamically with visual (and therefore high bandwidth) images as models of work objects.

IDE (Integrated Drive Electronics)
A standard electronic interface used between a computer motherboard's data paths or bus and the computer's disk storage devices. The IDE interface is based on the IBM PC ISA 16-bit bus standard, but it is also used in computers that use other bus standards. Most computers sold today use an enhanced version of IDE called EIDE. IDE gets its name because the disk drive controller is built into the logic board in the disk drive. IDE was adopted as a

standard by ANSI in November 1990. The ANSI name for IDE is Advanced Technology Attachment (ATA). The IDE (ATA) standard is one of several related standards maintained by the T10 Committee.

IDE RAID
IDE disks with Ultra DMA can now run at up to 33Mb/s, but because they can not be daisy-chained in the same way as SCSI and Fibre drives, they require a special RAID controller. The RAID controllers uses Ultra-Wide SCSI for a host channel, because IDE has a very limited cable length and less bandwidth, so although IDE disks are used, the RAID appears to the host as a normal SCSI drive.

IEEE 1394
A new, very fast external bus * standard that supports data transfer rates of up to 400 Mbps (400 million bits per second). Products supporting the 1394 standard go under different names, depending on the company. Apple, which originally developed the technology, uses the trademarked name FireWire. Other companies use other names, such as i.link and Lynx, to describe their 1394 products. A single 1394 port can be used to connect up 63 external devices. In addition t/o its high speed, 1394 also supports isochronous data -- delivering data at a guaranteed rate. This makes it ideal for devices that need to transfer high levels of data in real-time, such as video devices. Although extremely fast and flexible, 1394 is also expensive. Like USB, 1394 supports both Plug-and-Play and hot plugging, and also provides power to peripheral devices. The main difference between 1394 and USB is that 1394 supports faster data transfer rates and is more expensive. For this reason, it is expected to be used mostly for devices that require large throughputs , such as video cameras, whereas USB will be used to connect most other peripheral devices.

Initialization
(1) With hardware, initializing a disk means formatting it. (2) In programming, initialize means to assign a starting value to a variable. (3) Initialize can refer to the process of starting up a program or system.

I/O (Input/Output)
Describes any operation, program or device that transfers data to or from a computer. A typical I/O device includes a printer, hard disk, keyboard and mouse. Some devices are basically input-only devices (keyboard and mouse); others are primarily output-only devices (printer) and others provide both input and output of data (hard disk, diskette, writeable CD-ROMs).

IPMI
Intelligent Platform Management Interface (IPMI) is a hardware level interface specification that defines a common, abstracted, message-based interface to platform monitoring and control functions. As a hardware-level interface, it sits at the bottom of a typical management software stack. Thus, IPMI is "management software neutral." It can be exposed through any standard management software interface, such as WMI, CIM, SNMP or DMI.

ISA (Industry Standard Architecture)
A standard bus (computer interconnection) architecture that is associated with the IBM AT motherboard. It allows 16 bits at a time to flow between the motherboard circuitry and an expansion slot card and its associated device(s).

LAN (Local Area Network)
A network of interconnected workstations sharing the resources of a single processor or server within a relatively small geographic area. Typically, this might be within the area of a small office building. However, FDDI extends a local area network over a much wider area. Usually, the server has applications and data storage that are shared in common by multiple workstation users. A local area network may serve as few as four or five users or, in the case of FDDI, may serve several thousand. The main LAN technologies are:

- Ethernet
- Token ring

- ARCNET
- FDDI (Fiber Distributed Data Interface)

Typically, a suite of application programs can be kept on the LAN server. Users who need an application frequently can download it once and then run it from their local hard disk. Users can order printing and other services as needed through applications run on the LAN server. A user can share files with others at the LAN server; read and write access is maintained by a LAN administrator. A LAN server may also be used as a Web server if safeguards are taken to secure internal applications and data from outside access.

LVD / LVDS - Low Voltage Differential Signaling
LVD is a low noise, low power, low amplitude method for high-speed (gigabits per second) data transmission over copper wire. LVD differs from normal input/output (I/O) in a few ways: · Normal digital I/O works with 5 volts as a high (binary 1) and 0 volts as a low (binary 0). When you use a differential, you add a third option (-5 volts), which provides an extra level with which to encode and results in a higher maximum data transfer rate. · A higher data transfer rate means fewer wires are required, as in UW (Ultra Wide) and UW-2/3 SCSI hard disks, which use only 68 wires. These devices require a high transfer rate over short distances. Using standard I/O transfer, SCSI hard drives would require a lot more than 68 wires. · Low voltage means that the standard 5 volts is replaced by either 3.3 volts or 1.5 volts. · LVD uses a dual wire system, running 180 degrees of each other. This enables noise to travel at the same level, which in turn can get filtered more easily and effectively. · With standard I/0 signaling, data storage is contingent upon the actual voltage level. Voltage level can be affected by wire length (longer wires increase resistance, which lowers voltage). But with LVD, data storage is distinguished only by positive and negative voltage values, not the voltage level. Therefore, data can travel over greater lengths of wire while maintaining a clear and consistent data stream.

MDRAM (Multibank Dynamic RAM)
A type of video RAM, developed by MoSys, that divides memory into multiple 32 KB parts or "banks" that can be accessed individually. Traditional video RAM is monolithic; the entire frame buffer is accessed at one time. Having individual memory banks allows accesses to be interleaved concurrently, increasing overall performance. It's also cheaper since, unlike other forms of video RAM, cards can be manufactured with just the right amount of RAM for a given resolution capability instead of requiring it to be in multiples of megabytes.

Mirroring
Short for disk mirroring; Also known as RAID Level 1. A form of RAID in which the Array Management Function maintains two or more identical copies of data on separate disks. Also known as RAID Level 1 and disk shadowing.

Motherboard
The physical arrangement in a computer that contains the computer's basic circuitry and components. On the typical motherboard, the circuitry is imprinted or affixed to the surface of a firm planar surface and usually manufactured in a single step. The most common motherboard design in desktop computers today is the AT, based on the IBM AT motherboard. A more recent motherboard specification, ATX, improves on the AT design. In both the AT and ATX designs, the computer components included in the motherboard are:

- The microprocessor
- (Optionally) coprocessors
- Memory
- BIOS
- Expansion slots
- Interconnecting circuitry

Additional components can be added to a motherboard through its expansion slots. The electronic interface between the motherboard and the smaller boards or cards in the expansion slots is called the bus.

NVRAM - Non-Volatile Random Access Memory
A type of memory that retains its contents when power is turned off. One type of NVRAM is SRAM that is made non-volatile by connecting it to a constant power source such as a battery. Another type of NVRAM uses EEPROM chips to save its contents when power is turned off. In this case, NVRAM is composed of a combination of SRAM and EEPROM chips.

OLCE - Online Capacity Expansion
Provides the capability of adding drive space to existing arrays without the need for rebooting the system. When using OLCE, the system dynamically reallocates the existing information, while realigning the array to provide for the additional drive(s).

Parity
The quality of being either odd or even. The fact that all numbers have a parity is commonly used in data communications to ensure the validity of data. This is called parity checking.

PC99
A collection of PC system definitions and bus and device design requirements and recommendations for 1999-2000. The Design Guide is for engineers who build personal computers, expansion cards and peripheral devices that will be used with the Windows* NT* and Windows 98 operating systems. The goal of the design guide is to give clarity to the industry, advance the platform, and ensure availability of hardware and software drivers that provide a good experience for PC users with the Windows NT and Windows 98 operating systems. In addition, PC99 System Design Guide addresses major PC platform initiatives, including Instantly Available PC, USB, and Accelerated Graphics Port (AGP).

PC100
PC100 compliant SDRAM is the latest memory standard. This new memory is a new standard for SDRAM, capable of providing memory access time by following the new JEDEC standard of SDRAMs and is fully backward compatible with existing memory systems. This new specification is the latest module design to fully support the Intel 440BX AGPset and a new generation of mainboard series power, PC100 SDRAM modules is strongly recommended. Due to the strict timing issues involved when operating at 100MHz frequency, using non-compliant memory modules may cause systems to fail to boot. Even if the system does boot, changes to the operating environment such as temperature or certain hardware applications will severely impact memory reliability. PC100 compliant modules must be embedded with all PC100 compliant SDRAM chips. Although most memory modules shipped today are specified or marked as 10ns, which are often called 100MHz, they don't actually perform at 100MHz all the time and, most importantly, are not PC100 compliant. Some chip venders have been developing special versions of PC100 compliant 8 ns modules, but only a select group of them with very sophisticated test capabilities will be able to meet the critical requirements.

PCI (Peripheral Component Interconnect)
An interconnection system between a microprocessor and attached devices in which expansion slots are spaced closely for high-speed operation. Using PCI, a computer can support both new PCI cards while continuing to support ISA expansion cards, currently the most common kind of expansion card. Designed by Intel, the original PCI was similar to the VESA Local Bus. However, PCI 2.0 is no longer a local bus and is designed to be independent of microprocessor design. PCI is designed to be synchronized with the clock speed of the microprocessor, in the range of 20 to 33 MHz.

PCI is now installed on most new desktop computers, not only those based on Intel's Pentium processor but also those based on the PowerPC. PCI transmits 32 bits at a time in a 124-pin connection (the extra pins are for power supply and grounding) and 64 bits in a 188-pin connection in an expanded implementation. PCI uses all active paths to transmit both address and data signals, sending the address on one clock cycle and data on the next. Burst

data can be sent starting with an address on the first cycle and a sequence of data transmissions on a certain number of successive cycles.

PCMCIA (Personal Computer Memory Card International Association)
An industry group organized in 1989 to promote standards for a credit card-size memory or I/O device that would fit into a personal computer, usually a notebook or laptop computer. The PCMCIA 2.1 Standard was published in 1993. As a result, PC users can be assured of standard attachments for any peripheral device that follows the standard. The initial standard and its subsequent releases describe a standard product, the PC Card.

RAID (Redundant Array of Independent Disks)
A way of storing the same data in different places (thus, redundantly) on multiple hard disks. By placing data on multiple disks, I/O operations can overlap in a balanced way, improving performance. Since multiple disks increases the mean time between failure (MTBF), storing data redundantly also increases fault-tolerance.

A RAID appears to the OS to be a single logical hard disk. RAID employs the technique of striping, which involves partitioning each drive's storage space into units ranging from a sector (512 bytes) up to several megabytes. The stripes of all the disks are interleaved and addressed in order.

In a single-user system where large records, such as medical or other scientific images, are stored, the stripes are typically set up to be small (perhaps 512 bytes) so that a single record spans all disks and can be accessed quickly by reading all disks at the same time.

In a multi-user system, better performance requires establishing a stripe wide enough to hold the typical or maximum size record. This allows overlapped disk I/O across drives. There are at least nine types of RAID plus a non-redundant array (RAID-0):

- RAID-0. Has striping but no redundancy of data. Offers the best performance but no fault-tolerance.
- RAID-1. Also known as disk mirroring and consists of at least two drives that duplicate the storage of data. There is no striping. Read performance is improved since either disk can be read at the same time. Write performance is the same as for single disk storage. Provides the best performance and the best fault-tolerance in a multi-user system.
- RAID-2. Uses striping across disks with some disks storing error checking and correcting (ECC) information. It has no advantage over RAID-3.
- RAID-3. This type uses striping and dedicates one drive to storing parity information. The ECC information detects errors. Data recovery is accomplished by calculating the exclusive OR (XOR) of the information recorded on the other drives. Since an I/O operation addresses all drives at the same time, RAID-3 cannot overlap I/O. For this reason, RAID-3 is best for single-user systems with long record applications.
- RAID-4. Uses large stripes, which means you can read records from any single drive. This allows you to take advantage of overlapped I/O for read operations. Since all write operations have to update the parity drive, no I/O overlapping is possible. RAID-4 offers no advantage over RAID-5.
- RAID-5. Includes a rotating parity array, thus addressing the write limitation in RAID-4. Thus, all read and write operations can be overlapped. RAID-5 stores parity information but not redundant data (but parity information can be used to reconstruct data). RAID-5 requires at least three and usually five disks for the array. It's best for multi-user systems in which performance is not critical or which do few write operations.

- RAID-6. Similar to RAID-5 but includes a second parity scheme distributed across different drives and thus offers extremely high fault- and drive-failure tolerance. There are few or no commercial examples currently.
- RAID-7. Includes a real-time embedded operating system as a controller, caching via a high-speed bus and other characteristics of a stand-alone computer. One vendor offers this system.
- RAID-10. Offers an array of stripes in which each stripe is a RAID-1 array of drives. Offers higher performance than RAID-1 but at much higher cost.
- RAID-53. Offers an array of stripes in which each stripe is a RAID-3 array of disks. This offers higher performance than RAID-3 but at much higher cost.

RAM
Random Access Memory. Any byte of RAM can be accessed directly in a single memory cycle. Information can be read from and written to RAM. RAM is volatile — when power is turned off, it loses its memory.

RDRAM - Rambus DRAM
A type of memory (DRAM) developed by Rambus, Inc. Whereas the fastest current memory technologies used by PCs (SDRAM) can deliver data at a maximum speed of about 100 MHz, RDRAM transfers data at up to 600 MHz. In 1997, Intel announced that it would license the Rambus technology for use on its future motherboards, thus making it the likely de facto standard for memory architectures. However, a consortium of computer vendors is working on an alternative memory architecture called SyncLink DRAM (SLDRAM). RDRAM is already being used in place of VRAM in some graphics accelerator boards, but it is not expected to be used commonly for the main memory of PCs until late 2000 or 2001. Intel and Rambus are also working a new version of RDRAM, called nDRAM, that will support data transfer speeds at up to 1,600 MHz (1.6GHz).

RISC - Reduced Instruction Set Computer (Pronounced risk)
A type of microprocessor that recognizes a relatively limited number of instructions. Until the mid-1980s, the tendency among computer manufacturers was to build increasingly complex CPUs that had ever-larger sets of instructions. At that time, however, a number of computer manufacturers decided to reverse this trend by building CPUs capable of executing only a very limited set of instructions. One advantage of reduced instruction set computers is that they can execute their instructions very fast because the instructions are so simple. Another, perhaps more important advantage, is that RISC chips require fewer transistors, which makes them cheaper to design and produce. Since the emergence of RISC computers, conventional computers have been referred to as CISCs (Complex Instruction Set Computers). There is still considerable controversy among experts about the ultimate value of RISC architectures. Its proponents argue that RISC machines are both cheaper and faster, and are therefore the machines of the future. Skeptics note that by making the hardware simpler, RISC architectures put a greater burden on the software. They argue that this is not worth the trouble because conventional microprocessors are becoming increasingly fast and cheap anyway. To some extent, the argument is becoming moot because CISC and RISC implementations are becoming more and more alike. Many of today's RISC chips support as many instructions as yesterday's CISC chips. And today's CISC chips use many techniques formerly associated with RISC chips.

RLM - RAID Level Migration
Allows a user to change RAID levels. Dynamic RLM allows the user to do this without restarting the system.

SAF-TE - SCSI Accessed Fault-Tolerant Enclosure
Industry Standard for enclosure management support. SAF-TE uses a common interface to continuously monitor elements such as temperature, drive, power and fan status. In addition, SAF-TE supported enclosures can provide status updates to LAN administrators via any SAF-TE-compliant

device

SCSI (Small Computer System Interface)
A set of evolving ANSI standard electronic interfaces that allow personal computers to communicate with peripheral hardware such as disk drives, tape drives, CD-ROM drives, printers and scanners faster and more flexibly than previous interfaces. Developed at Apple Computer and still used in the Macintosh, the present set of SCSIs are parallel interfaces. SCSI ports are built into most personal computers today and supported by all major OSes.

In addition to faster data rates, SCSI is more flexible than earlier parallel data transfer interfaces. The latest SCSI standard, Ultra-2 SCSI for a 16-bit bus can transfer data at up to 80 megabytes per second (MBps). SCSI allows up to seven or 15 devices (depending on the bus width) to be connected to a single SCSI port in daisy-chain fashion. This allows one circuit board or card to accommodate all the peripherals, rather than having a separate card for each device, making it an ideal interface for use with portable and notebook computers. A single host adapter, in the form of a PC Card, can serve as a SCSI interface for a "laptop," freeing up the parallel and serial ports for use with an external modem and printer while allowing other devices to be used in addition.

Although not all devices support all levels of SCSI, the evolving SCSI standards are generally backward-compatible. That is, if you attach an older device to a newer computer with support for a later standard, the older device will work at the older and slower data rate.

The original SCSI, now known as SCSI-1, evolved into SCSI-2, known as "plain SCSI." as it became widely supported. SCSI-3 consists of a set of primary commands and additional specialized command sets to meet the needs of specific device types. The collection of SCSI-3 command sets is used not only for the SCSI-3 parallel interface but for additional parallel and serial protocols, including Fibre Channel, Serial Bus Protocol (used with the IEEE 1394 Firewire physical protocol) and the Serial Storage Protocol (SSP).

A widely implemented SCSI standard is Ultra-2 (sometimes spelled "Ultra2") which uses a 40 MHz clock rate to get maximum data transfer rates up to 80 MBps. It provides a longer possible cabling distance (up to 12 meters) by using Low Voltage Differential (LVD) signaling. Earlier forms of SCSIs use a single wire that ends in a terminator with a ground. Ultra-2 SCSI sends the signal over two wires with the data represented as the difference in voltage between the two wires. This allows support for longer cables. A low voltage differential reduces power requirements and manufacturing costs.

The latest SCSI standard is Ultra-3 (sometimes spelled "Ultra3")which increases the maximum burst rate from 80 Mbps to 160 Mbps by being able to operate at the full clock rate rather than the half-clock rate of Ultra-2. The standard is also sometimes referred to as Ultra160/m. New disk drives supporting Ultra160/m will offer much faster data transfer rates. Ultra160/m also includes cyclical redundancy checking (CRC) for ensuring the integrity of transferred data and domain validation for testing the SCSI network.

SDRAM (Synchronous Dynamic Random Access Memory)
A generic name for various kinds of DRAM that are synchronized with the clock speed that the microprocessor is optimized for. This tends to increase the number of instructions that the processor can perform in a given time. The speed of SDRAM is rated in MHz rather than in nanoseconds (ns). This makes it easier to compare the bus speed and the RAM chip speed. You can convert the RAM clock speed to nanoseconds by dividing the chip speed into 1 billion ns (which is one second). For example, an 83 MHz RAM would be equivalent to 12

SIMM - Single In-Line Memory Module
A small circuit board that can hold a group of memory chips. Typically, SIMMs hold up to 8 (on Macintoshes) or 9 (on PCs) RAM chips. On PCs, the ninth chip is often used for parity error checking. Unlike memory chips, SIMMs are measured in bytes rather than bits. SIMMs are easier to install than individual memory chips. The bus from a SIMM to the actual memory chips is 32 bits wide. A newer technology, called dual in-line memory module (DIMM),

provides a 64-bit bus. For modern Pentium microprocessors that have a 64-bit bus, you must use either DIMMs or pairs of SIMMs.

SMBIOS - System Management BIOS
Allows manufacturers to develop structures to access attributes that are known by the system BIOS, but have no standard interface to management software.

SNMP (Simple Network Management Protocol)
The protocol governing network management and the monitoring of network devices and their functions.

Striped Array
Distributes application data across two or more members disks in a regular pattern.

Striping
Short for disk striping; also known as RAID Level 0. A mapping technique in which fixed-size consecutive ranges of virtual disk data addresses are mapped to successive array members in a cyclical pattern.

UDMA (Ultra DMA or Ultra DMA/33)
A protocol for transferring data between a hard disk drive through the computer's data paths (or bus) to the computer's random access memory (RAM). The Ultra DMA/33 protocol transfers data in burst mode at a rate of 33.3 MBps (megabytes per second), twice as fast as the previous Direct Memory Access (DMA) interface. Ultra DMA was developed as a proposed industry standard by the Quantum Corporation, makers of hard disk drives and Intel, makers of chip sets that support computer bus technology.

Ultra DMA support in your computer means that it will boot (start) and open new applications more quickly. It will also help users of graphics-intensive and other applications that require large amounts of access to data on the hard drive. Ultra DMA uses Cyclical Redundancy Checking (CRC), offering a new level of data protection. Because the Ultra DMA protocol is designed to work with legacy PIO and DMA protocols, it can be added to many existing computers by installing an Ultra DMA/33 PCI adapter card. Ultra DMA uses the same 40-pin IDE interface cable as PIO and DMA.

ULTRASCSI (Ultra160 / Ultra320)
A method that enables very fast data transfer rate on the SCSI bus. The maximum UltraSCSI data transfer rates are 20 MBytes/second (160Mbits/second) and 40 MBytes/second(320Mbits/second) for Wide SCSI host adapters.

USB
A "plug-and-play" interface between a computer and add-on devices (such as audio players, joysticks, keyboards, telephones, scanners and printers). With USB, a new device can be added to your computer without having to add an adapter card or even having to turn the computer off. The USB peripheral bus standard was developed by Compaq, IBM, DEC, Intel, Microsoft, NEC and Northern Telecom and the technology is available without charge for all computer and device vendors.

USB supports a data speed of 12 megabits per second. This speed will accommodate a wide range of devices, including MPEG-2 video devices, data gloves and digitizers. It is anticipated that USB will easily accommodate plug-in telephones that use ISDN and digital PBXs.

Si dnce October 1996, the Windows operating systems have been equipped with USB drivers or special software designed to work with specific I/O device types. USB is integrated into Windows 98. Today, most new computers and peripheral devices are equipped with USB. A different "plug-and-play" standard, Firewire/IEEE 1394, is designed to support much higher data rates and devices such as video camcorders and digital videodisk (DVD) players. Both standards are expected to exist together, serving different device types.

WAN (Wide Area Network)
A geographically dispersed telecommunications network and the term

distinguishes a broader telecommunication structure from a local area network (LAN). A wide area network may be privately owned or rented, but the term usually connotes the inclusion of public (shared user) networks. An intermediate form of network in terms of geography is a metropolitan area network (MAN).

Write-Through / Write-Back Cache
When the processor writes to main memory, the data is first written to cache memory, assuming that the processor will probably read this data again soon. In write-through cache, data is written to main memory at the same time it is written to cache memory. In write-back cache, data is only written to main memory when it is forced out of cache memory. Write-through caching is simpler than write-back because an entry to cache memory that must be replaced can be overwritten in cache memory because it will already have been copied to main memory. Write-back requires cache memory to initiate a main memory write of the flushed entry followed (for a processor read) by a main memory read. However, write-back is more efficient because an entry can be written many times to cache memory without a main memory access.

XOR - eXclusive OR
A Boolean operator that returns a value of TRUE only if just one of its operands is TRUE. In contrast, an inclusive OR operator returns a value of TRUE if either or both of its operands are TRUE.

Copyright ® 2000 - American Megatrends Inc.