

# Chapter

# 1

# Server Hardware

---

## COMPTIA SERVER+ EXAM OBJECTIVES COVERED IN THIS CHAPTER:

### ✓ 1.1 Explain the purpose and function of server form factors

- Rack mount (dimensions [1U, 2U, 4U], cable management arms, rail kits)
- Tower
- Blade technology (blade enclosure [backplane/midplane, power supply sockets, network modules/switches, management modules], blade server)

### ✓ 1.2 Given a scenario, install, configure and maintain server components

- CPU (Multiprocessor vs. multicore, socket type, cache levels: L1, L2, L3, speeds [core, bus, multiplier], CPU stepping, architecture [x86, x64, ARM])
- RAM (ECC vs. non-ECC, DDR2, DDR3, number of pins, static vs. dynamic, module placement, CAS latency, timing, memory pairing)
- Bus types, bus channels and expansion slots (height differences and bit rate differences, PCI, PCIe, PCI-X)
- NICs
- Hard drives
- Riser cards
- RAID controllers
- BIOS/UEFI (CMOS battery)
- Firmware
- USB interface/port
- Hotswap vs. non-hotswap components



### ✓ 1.3 Compare and contrast power and cooling components

- Power (voltage [110V vs. 220V vs. -48V, 208V vs. 440V/460V/480V], wattage, consumption, redundancy, 1-phase vs. 3-phase power, plug types [NEMA, Edison, twist lock])
- Cooling (airflow, thermal dissipation, baffles/shrouds, fans, liquid cooling)



While servers and workstations have many of the same hardware components and in many cases use the same or similar operating systems, their roles in the network and therefore the requirements placed upon them are quite different. For this reason, CompTIA has developed the Server+ certification to validate the skills and knowledge required to design, install, and maintain server systems in the enterprise. Although many of the skills required to maintain workstations are transferable to maintaining servers, there are certainly enough differences both in the devices themselves and in the environment in which they operate to warrant such a certification. This book is designed to prepare you for the SK0-004 exam, otherwise known as the CompTIA Server+ exam.


## Server Form Factors

When we use the term *form factor* when discussing any computing device or component, we are talking about its size, appearance, or dimensions. Form factor is typically used to differentiate one physical implementation of the same device or component from another. In the case of servers, we are talking about the size and dimensions of the enclosure in which the server exists.

In this section we'll look at the major server form factors: the rack mount, the tower, and the blade. Each has its own unique characteristics and considerations you need to take into account when deploying.

### Rack Mount

Rack mount servers are those that are designed to be bolted into a framework called a rack and thus are designed to fit one of several standard size rack slots, or *bays*. They also require *rail kits*, which when implemented allow you to slide the server out of the rack for maintenance. One of the benefits of using racks to hold servers, routers, switches, and other hardware appliances is that a rack gets the equipment off the floor, while also making more efficient use of the space in the server room and maintaining good air circulation. A rack with a server and other devices installed is shown in Figure 1.1.

**FIGURE 1.1** Server in a rack

## Dimensions

As you may have noticed in Figure 1.1, there are several items in the rack and they take up various amounts of space in the rack. While both 19 and 23 inch wide racks are used, this is a 19 inch wide rack. Each module has a front panel that is 19 inches (482.6 mm) wide. The dimension where the devices or modules differ is in their height. This dimension is measured in *rack units*, or U for short. Each U is 1.75 inches (44.45 mm) high. While in the diagram the Liquid Crystal Display (LCD) takes up 7U, there are four standard sizes for servers:


**1U** These are for very small appliances or servers that are only 1.75 inches high. In the diagram, there is a KVM switch (which provides a common keyboard, mouse, and monitor to use for all devices in the rack) and an Ethernet switch or hub that uses a 1U bay.

**2U** This is the middle of the most common sizes. In the diagram there is a server in the bottom of the rack that is using a 2U bay.

**3U** While not as common, 3U servers are also available.

**4U** Although there are no devices in the rack shown in Figure 1.1 that use 4U, this is a common bay size for servers. A 4U server is shown in Figure 1.2. For comparison, this server has twice the height of the 2U server in Figure 1.1.




**FIGURE 1.2** A 4U server

It is also worth knowing that there are enclosures for blade servers that can be 10U in size. The typical rack provides 42U of space.

### Cable Management Arms


One of the challenges in the server room is to keep control of all the cables. When you consider the fact that servers use rail kits to allow you to slide the servers out for maintenance, there must be enough slack in both the power cable and the data cable(s) to permit this. On the other hand, you don't want a bunch of slack hanging down on the back of the rack for each device. To provide the slack required and to keep the cables from blanketing the back of the rack and causing overheating, you can use *cable management arms* (see Figure 1.3). These arms contain the slack and are designed to follow the server when you slide it out of the bay.

**FIGURE 1.3** Cable management arm

## Rail Kits

You already know that rail kits are used to provide a mechanism for sliding the server out of the rack. The rail kits have an inner rack and an outer rack. The inner rack attaches to the server, whereas the outer one attaches to the rack. The inner rack is designed to fit inside the outer rack and then it “rides” or slides on the outer rack. The installation steps are shown in Figure 1.4.

**FIGURE 1.4** Rail kit installation



## Tower

A second form factor with which you are likely to be familiar is the *tower server*. This type bears the most resemblance to the workstations you are used to working with. When many of these devices are used in a server room, they reside not in the rack but on shelves. They are upright in appearance, as shown in Figure 1.5.

**FIGURE 1.5** Tower server



It is also possible to place a tower server in a rack by using a conversion kit. The issue with this approach is that it wastes some space in the rack. A tower server using a conversion kit is shown in Figure 1.6.


**FIGURE 1.6** Tower server in a rack



## Blade Technology

Finally, servers may also come in blade form. This technology consists of a server chassis housing multiple thin, modular circuit boards, known as *server blades*. Each blade (or card) contains processors, memory, integrated network controllers, and other input/output (I/O) ports. Servers can experience as much as an 85 percent reduction in cabling for blade installations over conventional 1U or tower servers. Blade technology also uses much less space, as shown in a comparison of a blade system and a rack system in Figure 1.7.


**FIGURE 1.7** Rack vs. blade



## Blade Enclosure

A blade enclosure is a system that houses multiple blade servers. The chassis of the enclosure provides power and cooling to the blade servers. In Figure 1.8, a blade server is shown being inserted into an enclosure.


**FIGURE 1.8** Blade enclosure



## Backplane/Midplane

The backplane provides a connection point for the blade servers in the blade enclosure. Some backplanes are constructed with slots on both sides, and in that case, they are located in the middle of the enclosure and are called *midplanes*. In other cases, servers will be connected on one side, and power supplies and cooling might be connected on the other side. This arrangement is shown in Figure 1.9. The component labeled 3 is the midplane.


**FIGURE 1.9** Midplane



**Power Supply Sockets**

The midplane or backplane also supplies power connections to various components. When a midplane is in use, connections are provided on the back side for power modules. The power connectors on an IBM midplane are shown in Figure 1.10. The blade power connector is where the blade servers get their power and the power module connector is for the cable that plugs into the power sockets.

**FIGURE 1.10** Midplane power



**Network Modules/Switches**


Blade enclosures can accept several types of modules in addition to blade servers. At least one and probably two switch modules will be present to provide networking for the servers. This switch module is typically Ethernet but not always.

**Management Modules**

Finally, there will be a management module that allows for configuring and managing the entire enclosure of blade servers. This includes things like the IP addresses of the individual blade servers and connecting to and managing the storage. For redundancy's sake, there may be multiple management modules. The typical location of the management module is shown in Figure 1.11.




**FIGURE 1.11** Advanced management module



## Blade Server

The blade servers are individual cards, each of which acts as a separate physical server. There will be a number of these—for example, 8, 16, or 24. Any blade slots that are not in use should have the blade filler in place. The insertion of both a blade server and a blade filler is shown in Figure 1.12.

**FIGURE 1.12** Inserting a blade server and filler



# Installing and Configuring Server Components

Just as an A+ technician needs to be familiar with all of the possible components that may exist inside the box and how to install, maintain, and repair those components, as a Server + technician, you must know the same with regard to servers. Servers have all the same components that are found in workstations, but due to the high workloads they experience as a result of their roles in the network, the components must be more robust. This section explores server versions of key components.

## CPU

The central processing unit (CPU) in servers must be capable of handling high workloads without overheating. In many cases, this requires the use of both multiple-core processors and multiple CPUs. A multiple-core processor is one with multiple cores, each of which can operate as a separate CPU. In this section we'll look at the types of sockets server CPUs use, the way they use memory, the possible architectures you may encounter, and the various speed values you may see and their meaning. We'll also introduce the concept of CPU stepping.

## Socket Type

CPUs are connected to the motherboard via a socket on the board. The most common socket types are listed in Table 1.1.


**TABLE 1.1** Server socket types

| Socket name           | CPU families supported   | Package | Pin count | Bus speed                                     |
|-----------------------|--|---------|-----------|---|
| LGA 771/<br>Socket J  | Intel Xeon   | LGA     | 771       | 1600 MHz                                      |
| LGA 1366/<br>Socket B | Intel Core i7 (900 series)<br>Intel Xeon (35xx, 36xx, 55xx,<br>56xx series), Intel Celeron | LGA     | 1366      | 4.8–6.4 GT/s<br>(gigatransfers<br>per second) |
| LGA 1248              | Intel Itanium 9300-series  | LGA     | 1248      | 4.8 GT/s                                      |
| LGA 1567              | Intel Xeon 6500/7500-series  | LGA     | 1567      | 4.8–6.4 GT/s                                  |

**TABLE 1.1** Server socket types (*continued*)

| Socket name           | CPU families supported  | Package | Pin count | Bus speed    |
|-----------------------|---|---------|-----------|--------------|
| LGA 2011/<br>Socket R | Intel Core i7 3xxx Sandy Bridge-E<br>Intel Core i7 4xxx Ivy Bridge-E<br>Intel Xeon E5 2xxx/4xxx [Sandy<br>Bridge EP] (2/4S)<br>Intel Xeon E5-2xxx/4xxx v2 [Ivy<br>Bridge EP] (2/4S) | LGA     | 2011      | 4.8–6.4 GT/s |
| Socket F              | AMD Opteron 13xs, 2200, 2300,<br>2400, 8200, 8300, 8400, AMD<br>Athlon 64 FX  | LGA     | 1207      | 200 MHz      |
| Socket 940            | Opteron 100, 200 and 800  | PGA-ZIF | 940       | 800 MHz      |
| G34                   | AMD Opteron 6000  | LGA     | 1974      | 3.2 GHz      |
| AM3+                  | AMD Phenom II, Athlon 2,<br>Sempron, Opteron 3xxx   | PGA-ZIF | 942       | 3.2 GHz      |

Notice in Table 1.1 that most of the processors use the land grid array (LGA) package. These types of sockets don't have pins on the chip. Instead, they have bare gold-plated copper that touches pins that protrude from the CPU that goes in the socket. The LGA 2011/Socket R, however, uses a version of pin grid array (PGA), an alternative design in which the socket has the pins and they fit into the CPU when it is placed in the socket. A comparison of PGA (on the left) and LGA sockets is shown in Figure 1.13.

**FIGURE 1.13** PGA and LGA

LGA-compatible sockets have a lid that closes over the CPU and is locked in place by an L-shaped arm that borders two of the socket's edges. The nonlocking leg of the arm has a bend in the middle that latches the lid closed when the other leg of the arm is secured.

For CPUs based on the PGA concept, zero insertion force (ZIF) sockets are used. ZIF sockets use a plastic or metal lever on one of the two lateral edges to lock or release the mechanism that secures the CPU's pins in the socket. The CPU rides on the mobile top portion of the socket, and the socket's contacts that mate with the CPU's pins are in the fixed bottom portion of the socket.

## Cache Levels: L1, L2, L3

CPUs in servers use system memory in the server, but like most workstation CPUs they also contain their own memory, which is called *cache*. Using this memory to store recently acquired data allows the CPU to retrieve that data much faster in the event it is needed again. Cache memory can be located in several places, and in each instance it is used for a different purpose.

The Level 1 (L1) cache holds data that is waiting to enter the CPU. On modern systems, the L1 cache is built into the CPU. The Level 2 (L2) cache holds data that is exiting the CPU and is waiting to return to RAM. On modern systems, the L2 cache is in the same packaging as the CPU but on a separate chip. On older systems, the L2 cache was on a separate circuit board installed in the motherboard and was sometimes called cache on a stick (COAS<sub>t</sub>).

On some CPUs, the L2 cache operates at the same speed as the CPU; on others, the cache speed is only half the CPU speed. Chips with full-speed L2 caches have better performance. Some newer systems also have an L3 cache, which is external to the CPU die but not necessarily the CPU package.

The distance of the cache from the CPU affects both the amount of cache and the speed with which the CPU can access the information in that cache. The order of distance, with the closest first, is L1, L2, and L3. The closer to the CPU, the smaller the cache capacity, but the faster the CPU can access that cache type.

## Speeds

When measuring the speed of a CPU, the values are typically expressed in megahertz (MHz) and gigahertz (GHz). You may sometimes see it (as in Table 1.1) expressed in gigatransfers per second (GT/s). When expressed in GT/s, to calculate the data transmission rate, you must multiply the transfer rate by the bus width.

However, there are two speeds involved when comparing CPUs: core and bus.

**Core** Processors can have one or more *cores*. Each core operates as an individual CPU and each has an internal speed, which is the maximum speed at which the CPU can perform its internal operations, and is expressed in either MHz or GHz.

**Bus** The bus speed is the speed at which the motherboard communicates with the CPU. It's determined by the motherboard, and its cadence is set by a quartz crystal (the system crystal) that generates regular electrical pulses.

## Multiplier

The internal speed may be the same as the motherboard's speed (the external or bus speed), but it's more likely to be a multiple of it. For example, a CPU may have an internal speed of 1.3 GHz but an external speed of 133 MHz. That means for every tick of the system crystal's clock, the CPU has 10 internal ticks of its own clock.

## CPU Performance

*CPU time* refers to the amount of time the CPU takes to accomplish a task for either the operating system or for an application, and it is measured in clock *ticks* or seconds. The *CPU usage* is the total capacity of the CPU to perform work. The CPU time will be a subset of the usage and is usually represented as a percentage.

CPU usage values can be used to assess the overall workload of the server. When CPU usage is high—say 70 percent—there might be a slowing or lag in the system. CPU time values for a specific application or program, on the other hand, represent the relative amount of CPU usage attributable to the application.

We can also monitor CPU usage in terms of which component in the system is being served and in which security domain it is taking place. There are two main security domains in which the CPU operates: user mode and kernel mode. In user mode, it is working on behalf of an application and does not directly access the hardware. In kernel mode, it is working for the operating system and has more privileges.

When you are monitoring CPU performance, the following are common metrics and their meanings you'll encounter:

**User Time** Time the CPU was busy executing code in user space.

**System Time** Time the CPU was busy executing code in kernel space.

**Idle Time** Time the CPU was not busy; measures unused CPU capacity.

**Steal Time (Virtualized Hardware)** Time the operating system wanted to execute but was not allowed to by the hypervisor because it was not the CPU's turn for a time slot.

## CPU Stepping

When CPUs undergo revisions, the revisions are called *stepping levels*. When a manufacturer invests money to do a stepping, that means they have found bugs in the logic or have made improvements to the design that allow for faster processing. Integrated circuits have two primary classes of *mask sets* (mask sets are used to make the changes): *base layers* that are used to build the structures that make up the logic, such as transistors, and *metal layers* that connect the logic together. A base layer update is more difficult and time consuming than one for a metal layer. Therefore, you might think of metal layer updates as software versioning. Stepping levels are indicated by an alphabetic letter followed by a numeric number—for example, C-4. Usually, the letter indicates the revision level of a chip's base layers, and the number indicates the revision level of the metal layers. As an example, the first version of a processor is always A-0.



## Architecture

Some processors operate on 32 bits of information at a time, and others operate on 64 bits at a time. Operating on 64 bits of information is more efficient, but is only available in processors that support it and when coupled with operating systems that support it. A 64-bit processor can support 32-bit and 64-bit applications and operating systems, whereas a 32-bit processor can only support a 32-bit operating system and applications. This is what is being described when we discuss the *architecture* of the CPU. There are three main architectures of CPUs.

**x86** Processors that operate on 32 bits of information at a time use an architecture called x86. It derives its name from the first series of CPUs for computers (8086, which was only 16 bits, 286, 386, and 486).

**x64** Processors that operate on 64 bits of information at a time use an architecture called x64. It supports larger amounts of virtual memory and physical memory than is possible on its 32-bit predecessors, allowing programs to store larger amounts of data in memory.

**ARM** Advanced RISC Machine (ARM) is a family of reduced instruction set computing (RISC) instruction set architectures developed by British company ARM Holdings. Since its initial development, both ARM and third parties have developed CPUs on this architecture. It is one that requires fewer resources than either x86 or x64. In that regard, ARM CPUs are suitable for tablets, smartphones, and other smaller devices.

In Exercise 1.1, you'll replace a CPU in a server.

### EXERCISE 1.1

#### Replacing a CPU in a Server

1. Shut down and remove power from the entire system.
2. Remove the server node from the system. (Many systems have multiple servers in bays. To get to each server, you must remove the bay.)
3. Remove the server node cover (follow any instructions included with the documentation).
4. Use the proper type and size of screwdriver (usually a Number 2 Phillips-head) to loosen the screws (usually 4) holding the heatsink, and then lift it off of the CPU. (Yours may require a different type of screwdriver.)
5. Unclip the first CPU retaining latch and then unclip any remaining (usually there are two) latches.
6. Open the hinged CPU cover plate.
7. Remove the old CPU.
8. Insert the new CPU.
9. Install the heatsink (don't forget to put thermal grease between the CPU and the heatsink).

**EXERCISE 1.1 (continued)**

10. Replace the server node cover.
  11. Reinstall the server node.
  12. Replace power cords and then power on the system.
- 

## RAM

Like any computing device, servers require memory, and servers in particular require lots of it. In this section we will discuss the types of memory chips that are used in servers and describe some of the characteristics that differentiate them.

### ECC vs. Non-ECC

When data is moved to and from RAM, the transfer does not always go smoothly. Memory chips have error detection features and in some cases error correction functions. A type of RAM error correction is error correction code (ECC). RAM with ECC can detect and correct errors. To achieve this, additional information needs to be stored and more processing needs to be done, making ECC RAM more expensive and a little slower than non-ECC RAM.

In ECC, an algorithm is performed on the data and its check bits whenever the memory is accessed. If the result of the algorithm is all zeroes, then the data is deemed valid and processing continues. ECC can detect single- and double-bit errors and actually correct single-bit errors. This is a now a rarely used type of parity RAM. Most RAM today is non-ECC.

### DDR2 and DDR3

Double data rate (DDR) is clock-doubled SDRAM (covered later in this section). The memory chip can perform reads and writes on both sides of any clock cycle (the up, or start, and the down, or ending), thus doubling the effective memory executions per second. So, if you're using DDR SDRAM with a 100 MHz memory bus, the memory will execute reads and writes at 200 MHz and transfer the data to the processor at 100 MHz. The advantage of DDR over regular SDRAM is increased throughput and thus increased overall system speed.

#### DDR2

The next generation of DDR SDRAM is DDR2 (double data rate 2). This allows for two memory accesses for each rising and falling clock and effectively doubles the speed of DDR. DDR2-667 chips work with speeds of 667 MHz and PC2-5300 modules.

#### DDR3

The primary benefit of DDR3 over DDR2 is that it transfers data at twice the rate of DDR2 (eight times the speed of its internal memory arrays), enabling higher bandwidth

or peak data rates. By performing two transfers per cycle of a quadrupled clock, a 64-bit wide DDR3 module may achieve a transfer rate of up to 64 times the memory clock speed in megabytes per second (MBps). In addition, the DDR3 standard permits chip capacities of up to 8 GB. Selected memory standards, speeds, and formats are shown in Table 1.2.

**TABLE 1.2** Selected memory details

| Module standard | Speed       | Format    |
|-----------------|-------------|-----------|
| DDR-500         | 4000 MBps   | PC4-000   |
| DDR-533         | 4266 MBps   | PC4-200   |
| DDR2-667        | 5333 MBps   | PC2-5300  |
| DDR2-750        | 6000 MBps   | PC2-6000  |
| DDR2-800        | 6400 MBps   | PC2-6400  |
| DDR3-800        | 6400 MBps   | PC3-6400  |
| DDR3-1600       | 12,800 MBps | PC3-12800 |

## Number of Pins

Memory modules have pins that connect them to the motherboard slot in which they reside. Dual inline memory modules (DIMMs) have two rows of pins and twice the contact with the motherboard, creating a larger interface with it and resulting in a wider data path than older single inline memory modules (SIMMs). DIMMs differ in the number of conductors, or pins, that each particular physical form factor uses. Some common examples are 168-pin (SDR RAM), 184-pin (DDR, DDR2), and 240-pin (DDR3) configurations.

## Static vs. Dynamic

RAM can be either static or dynamic. Dynamic RAM requires a refresh signal whereas static RAM does not. This results in better performance for static RAM. A static RAM cell, on the other hand, requires more space on the chip than a dynamic RAM cell, resulting in less memory on the chip. This results in static RAM being more expensive when trying to provide the same number of cells.


In summary, static RAM is more expensive but faster, whereas dynamic RAM is slower but cheaper. The two types are often both used, however, due to their differing strengths and weaknesses. Static RAM is used to create the CPU's speed-sensitive cache, and dynamic RAM forms the larger system RAM space.

## Module Placement

Utilizing multiple channels between the RAM and the memory controller increases the transfer speed between these two components. Single-channel RAM does not take advantage of this concept, but dual-channel memory does and creates two 64-bit data channels. Do *not* confuse this with DDR. DDR doubles the rate by accessing the memory module twice per clock cycle.

Using dual channels requires a motherboard that supports dual channels and two or more memory modules. Sometimes the modules go in separate color-coded banks, as shown in Figure 1.14, and other times they use the same colors. Consult your documentation.

**FIGURE 1.14** Dual-channel memory slots



Memory runs in banks with two slots comprising a bank. The board should indicate which two slots are in the same bank by the color coding. It could be orange and yellow, or it might be some other combination of two colors. When installing the memory, install the same size modules in the same bank. If you don't, the modules will not operate in dual channel mode. This will impair the performance of the bank.

## CAS Latency

Another characteristic that can be used to differentiate memory modules is their CAS latency value. Column access strobe (CAS) latency is the amount of time taken to access a memory module and make that data available on the module's pins.

The lower the CL value, the better. In asynchronous DRAM, the delay value is measured in nanoseconds and the value is constant, while in synchronous DRAM, it is measured in clock cycles and will vary based on the clock rate.

## Timing

Memory timing measures the performance of RAM and consists of four components:

**CAS Latency** The time to access an address column if the correct row is already open

**Row Address to Column Address Delay** The time to read the first bit of memory without an active row

**Row Precharge Time** The time to access an address column if the wrong row is open

**Row Active Time** The time needed to internally refresh a row

Memory timings are listed in units of clock cycles; therefore, when translating these values to time, remember that for DDR memory, this will be half the speed of the transfer rate. It is also useful to note that memory timing is only part of the performance picture. The memory *bandwidth* is the throughput of the memory. Although advances in bandwidth technology (DDR2, DDR3) may have a negative effect on latency from timing, DDR2 and DDR3 can be clocked faster, resulting in a net gain in performance.

## Memory Pairing

Each motherboard supports memory based on the speed of the front-side bus (FSB) and the memory's form factor. If you install memory that is rated at a lower speed than the FSB, the memory will operate at that lower speed, if it works at all. In their documentation, most motherboard manufacturers list which type(s) of memory they support as well as maximum speeds and required pairings.

With regard to adding and upgrading memory, faster memory can be added to a server with slower memory installed, but the system will operate only at the speed of the slowest module present.

Moreover, although you can mix speeds, you cannot mix memory types. For example, you cannot use SDRAM with DDR, and DDR cannot be mixed with DDR2. When looking at the name of the memory, the larger the number, the faster the speed. For example, DDR2-800 is faster than DDR2-533.

Finally, memory pairing also refers to installing matched pairs of RAM in a dual-channel memory architecture.

## Bus Types, Bus Channels, and Expansion Slots

The motherboard provides the platform to which all components are attached and provides pathways for communication called *buses*. A bus is a common collection of signal pathways over which related devices communicate within the computer system. Expansion buses incorporate slots at certain points in the bus to allow insertion of external devices. In this section, we'll look at common server bus types and their characteristics.


### Height Differences and Bit Rate Differences

Two major differentiating characteristics of bus types are their bit rates and the form factor of the slot and adapter to which it mates. The dominant bus types in servers are forms of the Peripheral Component Interconnect (PCI) expansion bus. In the following sections, the three major types of PCI buses are covered, with attention given to both form factor and bit rate.

### PCI

The Peripheral Component Interconnect (PCI) bus is a 33 MHz wide (32-bit or 64-bit) expansion bus that was a modern standard in motherboards for general-purpose expansion devices. Its slots are typically white. You may see two PCI slots, but most motherboards have gone to newer standards. Figure 1.15 shows some PCI slots.



**FIGURE 1.15** PCI slots


PCI cards that are 32 bit with 33 MHz operate up to 133 MBps, whereas 32-bit cards with 66 MHz operate up to 266 MBps. PCI cards that are 64 bit with 33 MHz operate up to 266 MBps, whereas 64-bit cards with 66 MHz operate up to 533 MBps.

## PCI-X

PCI-eXtended (PCI-X) is a double-wide version of the 32-bit PCI local bus. It runs at up to four times the clock speed, achieving higher bandwidth, but otherwise it uses the same protocol and a similar electrical implementation. It has been replaced by the PCI Express (see the next section), which uses a different connector and a different logical design. There is also a 64-bit PCI specification that is electrically different but has the same connector as PCI-X. There are two versions of PCI-X: version 1 gets up to 1.06 GBps, and version 2 gets up to 4.26 GBps.

## PCIe

PCI Express (PCIE, PCI-E, or PCIe) uses a network of serial interconnects that operate at high speed. It's based on the PCI system; you can convert a PCIe slot to PCI using an adapter plug-in card, but you cannot convert a PCI slot to PCIe. Intended as a replacement for the Advanced Graphics Processor (AGP was an interim solution for graphics) and PCI, PCIe has the capability of being faster than AGP while maintaining the flexibility of PCI. There are four versions of PCIe: version 1 is up to 8 GBps, version 2 is up to 16 GBps, version 3 is up to 32 GBps, and as of this writing final specifications for version 4 are still being developed. Figure 1.16 shows the slots discussed so far in this section. Table 1.3 lists the speeds of each. The PCIe speeds shown are per lane. So a 4-lane version of PCIe 2 would operate at 20 GBps.

**FIGURE 1.16** Comparison of PCI slot types**TABLE 1.3** PCI and PCIe slot speeds

| Type           | Data transfer rate |
|----------------|--------------------|
| PCI 33, 32-bit | 133 MBps           |
| PCI 33, 64-bit | 266 MBps           |
| PCI 66, 32-bit | 266 MBps           |
| PCI 66, 64-bit | 533 MBps           |
| PCIe version 1 | 2 GBps             |
| PCIe version 2 | 5 GBps             |
| PCIe version 3 | 8 GBps             |
| PCIe version 4 | 16 GBps            |

## NICs

Network cards do exactly what you would think; they provide a connection for the server to a network. In general, network interface cards (NICs) are added via an expansion slot or they are integrated into the motherboard, but they may also be added through USB. The most common issue that prevents network connectivity is a bad or unplugged patch cable.

Network cards are made for various access methods (Ethernet, token ring) and for various media types (fiber optic, copper, wireless) connections. The network card you use must support both the access method and the media type in use.


The most obvious difference in network cards is the speed of which they are capable. Most networks today operate at 100 Mbps or 1 Gbps. Regardless of other components, the server will operate at the speed of the slowest component, so if the card is capable of 1 Gbps but the cable is only capable of 100 MBps, the server will transmit only at 100 Mbps.

Another significant feature to be aware of is the card's ability to perform auto-sensing. This feature allows the card to sense whether the connection is capable of full duplex and to operate in that manner with no action required.

There is another type of auto-sensing, in which the card is capable of detecting what type of device is on the other end and changing the use of the wire pairs accordingly. For example, normally a PC connected to another PC requires a crossover cable, but if both ends can perform this sensing, that is not required. These types of cards are called auto-MDIX.

In today's servers you will most likely be seeing 10 Gb cards and you may even see 40 Gb or 100 Gb cards. Moreover, many servers attach to storage networks and may run converged network adapters (CNAs), which act both as a host bus adapter (HBA) for the storage area network (SAN) and as the network card for the server. This concept is shown in Figure 1.17. In the graphic, FC stands for Fiber Channel and NC stands for Network Card.

**FIGURE 1.17** Traditional and CNA




## Hard Drives


Servers can contain three different types of hard drive architectures. In this section we'll look at each type.

### Magnetic Hard Drives

Magnetic drives were once the main type of hard drive used. The drive itself is a mechanical device that spins a number of disks or platters and uses a magnetic head to read and write data to the surface of the disks. One of the advantages of solid-state drives (discussed in the next section) is the absence of mechanical parts that can malfunction. The parts of a magnetic hard drive are shown in Figure 1.18.

**FIGURE 1.18** Magnetic hard drive

The basic hard disk geometry consists of three components: the number of sectors that each track contains, the number of read/write heads in the disk assembly, and the number of cylinders in the assembly. This set of values is known as CHS (for cylinders/heads/sectors). A *cylinder* is the set of tracks of the same number on all the writable surfaces of the assembly. It is called a cylinder because the collection of all same-number tracks on all writable surfaces of the hard disk assembly looks like a geometric cylinder when connected together vertically. Therefore, cylinder 1, for instance, on an assembly that contains three platters consists of six tracks (one on each side of each platter), each labeled track 1 on its respective surface. Figure 1.19 illustrates the key terms presented in this discussion.

**FIGURE 1.19** CHS

### 5400 rpm

The rotational speed of the disk or platter has a direct influence on how quickly the drive can locate any specific disk sector on the drive. This locational delay is called latency and is

measured in milliseconds (ms). The faster the rotation, the smaller the delay will be. A drive operating at 5400 rpms will experience about 5.5 ms of this delay.

### **7200 rpm**

Drives that operate at 7200 rpm will experience about 4.2 ms of latency. As of 2015, a typical 7200 rpm desktop hard drive has a sustained data transfer rate up to 750 Mbps. This rate depends on the track location, so it will be higher for data on the outer tracks and lower toward the inner tracks.

### **10,000 rpm**

At 10,000 rpm, the latency will decrease to about 3 ms. Data transfer rates (about 1.5 Gb) also generally go up with a higher rotational speed but are influenced by the density of the disk (the number of tracks and sectors present in a given area).

### **15,000 rpm**

Drives that operate at 15,000 rpm are higher-end drives and suffer only 2 ms of latency. They operate at just under 2 Gb. These drives also generate more heat, requiring more cooling to the case. They also offer faster data transfer rates for the same areal density (areal density refers to the amount of bits that can be stored in a given amount of space.).

### **Hot-Swappable Drives**

If a drive can be attached to the server without shutting down, then it is a hot-swappable drive. Drive types that are hot-swappable include USB, FireWire, SATA, and those that connect through Ethernet. You should always check the documentation to ensure that your drive supports this feature.

### **Solid-State Drives**

Solid-state drives (SSDs) retain data in nonvolatile memory chips and contain no moving parts. Compared to electromechanical hard disk drives (HDDs), SSDs are typically less susceptible to physical shock, are silent, have lower access time and latency, but are more expensive per gigabyte.

### **Hybrid Drives**

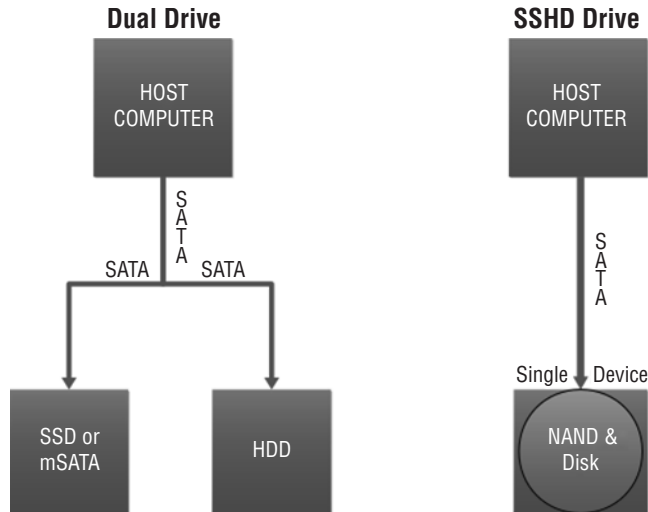
A hybrid drive is one in which both technologies, solid-state and traditional mechanical drives, are combined. This is done to take advantage of the speed of solid-state drives while maintaining the cost effectiveness of mechanical drives.

There are two main approaches to this: dual-drive hybrid and solid-state hybrid. Dual-drive systems contain both types of drives in the same machine, and performance is optimized by the user placing more frequently used information on the solid-state drive and less frequently accessed data on the mechanical drive—or in some cases by the operating system creating hybrid volumes using space in both drives.

A solid-state hybrid drive (SSHD), on the other hand, is a single storage device that includes solid-state flash memory in a traditional hard drive. Data that is most related to the performance of the machine is stored in the flash memory, resulting in improved performance. Figure 1.20 shows the two approaches to hybrid drives. In the graphic mSATA

refers to a smaller form of the SATA drive and NAND disk refers to a type of flash memory named after the NAND logic gate.

**FIGURE 1.20** Hybrid approaches



## Riser Cards

Riser cards allow you to add expansion cards to a system. You may already be familiar with their use in low-profile cases where the height of the case doesn't allow for a perpendicular placement of the full-height expansion card. They are also used in rack mount and blade servers to allow you to add feature cards in a horizontal position (instead of a standard vertical position).

Typically, a 1U system uses a 1U single-slot riser card whereas a 2U system uses a 2U three-slot riser card. An example of a riser card in a rack server is shown in Figure 1.21.

**FIGURE 1.21** Riser card in rack server




## RAID Controllers

Redundant Array of Independent Disks (RAID) is a multiple disk technology that either increases performance or allows for the automatic recovery of data from a failed hard drive by simply replacing the failed drive. There are several types of RAID that provide varying degrees of increased performance and/or fault tolerance. All of these techniques involve two or more hard drives operating together in some fashion.

RAID can be implemented using software or hardware. The highest levels of protection are provided by using hardware RAID, which requires that the system have a RAID *controller*. This hardware device is used to manage the disks in the storage array so they work as a logical unit. This is a card that fits into a PCI express slot to which the drives in the array are connected. This concept is shown in Figure 1.22.

**FIGURE 1.22** RAID controller



## BIOS/UEFI

Servers also contain firmware that provides low-level instructions to the device even in the absence of an operating system. This firmware, called either the Basic Input/Output System (BIOS) or the Unified Extensible Firmware Interface (UEFI), contains settings that can be manipulated and diagnostic utilities that can be used to monitor the device.

UEFI is a standard firmware interface for servers and PCs designed to replace BIOS. Some advantages of UEFI firmware are

- Better security; protects the preboot process
- Faster startup times and resuming from hibernation



- Support for drives larger than 2.2 terabytes (TB)
- Support for 64-bit firmware device drivers
- Capability to use BIOS with UEFI hardware


At startup, the BIOS or UEFI will attempt to detect the devices and components at its disposal. The information that it gathers, along with the current state of the components, will be available for review in the BIOS settings. Some of the components and the types of information available with respect to these devices and components are covered in this section.

You can view and adjust a server's base-level settings through the CMOS Setup program, which you access by pressing a certain key at startup, such as F1 or Delete (depending on the system). Complementary metal oxide semiconductor (CMOS) refers to the battery type that maintains power to the BIOS settings (also referred to as BIOS Setup). The most common settings you can adjust in CMOS are port settings (parallel, serial, USB), drive types, boot sequence, date and time, and virus/security protections. The variable settings that are made through the CMOS setup program are stored in nonvolatile random access memory (NVRAM), whereas the base instructions that cannot be changed (the BIOS) are stored on an EEPROM (Electrically Erasable Programmable Read-Only Memory) chip.

## CMOS Battery

The CMOS chip must have a constant source of power to keep its settings. To prevent the loss of data, motherboard manufacturers include a small battery to power the CMOS memory. On modern systems, this is a coin-style battery, about the diameter of a U.S. dime and about as thick. One of these is shown in Figure 1.23.

**FIGURE 1.23** CMOS battery in a server



When the server is not keeping correct time or date when turned off, it is usually a CMOS battery issue and a warning that the battery is soon going to die. In the absence of the server receiving time and date updates from a time server such as a Network Time Protocol (NTP) server, the time kept in the CMOS is the time source for the computer.

## Firmware

Firmware includes any type of instruction for the server that is stored in nonvolatile memory devices such as ROM, EPROM, or flash memory. BIOS and UEFI code is the most common example for firmware. Computer BIOSs don't go bad; they just become out of date or contain bugs. In the case of a bug, an upgrade will correct the problem. An upgrade may also be indicated when the BIOS doesn't support some component that you would like to install—a larger hard drive or a different type of processor, for instance.

Most of today's BIOSs are written to an EEPROM chip and can be updated through the use of software. Each manufacturer has its own method for accomplishing this. Check out the documentation for complete details. Regardless of the exact procedure, the process is referred to as *flashing* the BIOS. It means the old instructions are erased from the EEPROM chip and the new instructions are written to the chip.

Firmware can be updated by using an update utility from the motherboard vendor. In many cases, the steps are as follows:

1. Download the update file to a flash drive.
2. Insert the flash drive and reboot the machine.
3. Use the specified key sequence to enter the CMOS setup.
4. If necessary, disable Secure Boot.
5. Save the changes and reboot again.
6. Reenter the CMOS settings.
7. Choose boot options and then boot from the flash drive.
8. Follow the specific directions with the update to locate the upgrade file on the flash drive.
9. Execute the file (usually by typing **flash**).
10. While the update is completing, ensure you maintain power to the device.

## USB Interface/Ports

Like other computing devices, servers will probably have USB ports. There will probably be at least two, one for a mouse and keyboard respectively (although you will probably use a KVM switch for this when the servers are rack mounted or blade). These will probably be on the front of the server, although there may be additional ones on the back of the server.

Some specialized server products are able to operate as USB servers in that they allow network devices to access shared USB devices that are attached to the server. This typically

requires some sort of USB management software. In this case it may be necessary to use a USB hub connected to one of the USB ports on the server if there are not enough ports provided by the server.

## Hot-Swappable Component

A hot-swappable component is one that can be changed without shutting down the server. This is a desirable feature because for many server roles, shutting down the server is something to be minimized. However, just because a component is hot swappable doesn't mean changing the component doesn't require some administrative work before you make that change.

For example, to change a hot-swappable hard drive, in most cases you must prevent any applications from accessing the hard drive and remove the logical software links. Moreover, in many cases drives cannot be hot-plugged if the hard drive provides the operating system and the operating system is not mirrored on another drive. It also cannot be done if the hard drive cannot be logically isolated from the online operations of the server module. Nevertheless, it is still a great feature. In some high-end servers, it is even possible to hot-swap memory and CPU.

# Maintaining Power and Cooling Requirements

Computing equipment of any kind, including servers, require a certain level of power and an environment that is cool enough to keep the devices from overheating. In this section we'll discuss both power and cooling requirements and issues you should be aware of relating to those issues.

## Power

When discussing power it is helpful to define some terms that relate to power. In this section we'll do that, and we'll also look at power consumption and power redundancy. Finally, we'll explore power plug types you may encounter when dealing with servers in the enterprise.

## Voltage

Two terms that are thrown about and often confused when discussing power are voltage and amperage. *Voltage* is the pressure or force of the electricity, whereas *amperage* is the amount of electricity. They together describe the *wattage* supplied. The watts required by a device are the amps multiplied by the voltage.

Amps multiplied by the volts give you the wattage (watts), a measure of the work that electricity does per second.

Power supplies that come in servers (and in all computers for that matter) must be set to accept the voltage that is being supplied by the power outlet to which it is connected. This voltage is standardized but the standard is different in different countries. Almost all IT power supplies are now autosensing and universal voltage-capable (100-250 V) to allow the same product to operate worldwide. Those that do not will provide a switch on the outside of the case that allows you to change the type of power the supply is expecting, as shown in Figure 1.24.

**FIGURE 1.24** Voltage switch



### Single-Phase vs. Three-Phase Power

There are two types of power delivery systems: single-phase and three-phase. Single-phase power refers to a two-wire alternating current (AC) power circuit. Typically there is one power wire and one neutral wire. In the United States, 120V is the standard single-phase voltage,

with one 120V power wire and one neutral wire. In some countries, 230V is the standard single-phase voltage, with one 230V power wire and one neutral wire. Power flows between the power wire (through the load) and the neutral wire.

Three-phase power refers to three-wire AC power circuits. Typically there are three (phase A, phase B, phase C) power wires (120 degrees out of phase with one another) and one neutral wire. For example, a three-phase, four-wire 208V/120V power circuit provides three 120V single-phase power circuits and one 208V three-phase power circuit. Installing three-phase systems in datacenters helps to consolidate the power distribution in one place, reducing the costs associated with installing multiple distribution units.

Single-phase is what most homes have whereas three-phase is more typically found in industrial settings.

### **110V vs. 220V vs. 48V**

Although 110V is used in some parts of the world and 220V in others, the two systems have advantages and disadvantages. While 220V is more efficient in that it suffers less transmission loss (and it can use wiring rated for less current), 110V is safer if someone is electrocuted. Some datacenters deliver power to a rack at 220V and then use a transformer to step it down to 110V to the equipment if required.

Some equipment also is made for -48V DC power rather than 110/220 AC power. 48V is the common power scheme used in central offices and many datacenters. Many telcos can deliver 48V DC power to the facility and many are currently doing so. The advantage of using it is heat output. You no longer have the AC/DC conversion inside each device—just a DC/DC conversion. Less heat output means less (smaller) HVAC equipment. You will, however, need a rectifier, which is a small device that receives the 48V power and makes it -48V.

### **120/208V vs. 277/480V**

Earlier you learned that systems can be one-phase or three-phase. Most commercial systems use one of two versions of three-phase. The first we mentioned earlier: 120/208V. To review, that power circuit provides three 120V single-phase power circuits and one 208V three-phase power circuit.

The 277/480V circuit provides two 277V single-phase power circuits and one 480V three-phase power circuit. Server power supplies that operate directly from 480/277V power distribution circuits can reduce the total cost of ownership (TCO) for a high-performance cluster by reducing both infrastructure and operating cost. The trade-off is that 277/480V systems are inherently more dangerous.

## **Wattage**

Earlier you learned that voltage is the pressure or force of the electricity, whereas amperage is the amount of electricity. They together describe the wattage supplied. Amps multiplied by the volts give you the wattage (watts), a measure of the work that

electricity does per second. The power supply must be able to provide the wattage requirements of the server and any devices that are also attached and dependent on the supply for power.

## Consumption

Servers vary in their total consumption of power. However, there have been studies over the years that can give you an idea of what a server and some of its components draw in power. The following can be used as a rough guideline for planning:

- 1U rack mount x86: 300 W–350 W
- 2U rack mount, 2-socket x86: 350 W–400 W
- 4U rack mount, 4-socket x86: average 600 W, heavy configurations, 1000 W
- Blades: average chassis uses 4500 W; divide by number of blades per chassis (example: 14 per chassis, so about 320 per blade server)

Keep in mind that these are values for the server only. In a datacenter, much additional power is spent on cooling and other requirements. A value called power usage effectiveness (PUE) is used to measure the efficiency of the datacenter. It is a number that describes the relationship between the amount of power used by the entire datacenter and the power used by the server only. For example, a value of 3 means that the datacenter needs three times the power required by the servers. A lower value is better. Although this is changing, the general rule of thumb is that PUE is usually 2.0, which means a datacenter needs twice the power required by the servers.

## Redundancy

Datacenters usually deploy redundant power sources to maintain constant power.

Redundancy can be provided in several ways:

- Parallel redundancy or the N+1 option describes an architecture where there is always a single extra UPS available (that's the +1) and the N simply indicates the total number of UPSs required for the datacenter. Because the system runs in two feeds and there is only one redundant UPS, this system can still suffer failures.
- 2N redundancy means the datacenter provides double the power it requires. This ensures that the system is fully redundant.

Redundancy also refers to using redundant power supplies on the devices. Many servers come with two supplies, and you can buy additional power supplies as well. Always ensure that the power supply you buy can accommodate all the needs of the server. As you saw earlier in the section “Consumption,” many 4U rack and blade servers use a lot of power.

## Plug Types

You'll encounter several types of power plugs with servers. Let's examine each.


## NEMA

Power plugs that conform to the U.S. National Electrical Manufacturers Association (NEMA) standards are called NEMA plugs. There are many types of these plugs, and they differ in the orientation of the plugs and their shape. The two basic classifications of NEMA device are straight-blade and locking.

### Edison

The term *Edison plug* refers to the standard three-prong grounded or two-prong ungrounded plugs with which we are all familiar. Both are shown in Figure 1.25. Keep in mind the shape of the plug may differ somewhat.

**FIGURE 1.25** Edison plug



### Twist Lock

Twist-locking connectors refer to NEMA locking connectors manufactured by any company, although “Twist-Lock” remains a registered trademark of Hubbell Inc. The term is applied generically to locking connectors that use curved blades. The plug is pushed into the receptacle and turned, causing the now-rotated blades to latch.

A sample of this connector for a 6000 W power supply is shown in Figure 1.26.

**FIGURE 1.26** Locking plug





## Cooling

When all power considerations have been satisfied, your attention should turn to ensuring that the servers do not overheat. The CPUs in a server produce a lot of heat, and this heat needs to be dealt with. In this section, we'll look at the sources of heat in a server room or datacenter and approaches used to control this heat so it doesn't cause issues such as reboots (or worse).

### Airflow

Airflow, both within the server and in the server room or datacenter in general, must be maintained and any obstructions to this flow must be eliminated if possible. Inside the server case, if you add any fans, avoid making the following common mistakes:

- Placing intake and exhaust in close proximity on the same side of the chassis, which causes exhausted warm air to flow back into the chassis, lowering overall cooling performance
- Installing panels and components in the way of airflow, such as the graphics card, motherboard, and hard drives


You must also consider the airflow around the rack of servers and, in some cases, around the rows of racks in a large datacenter. We'll look at some approaches to that in the “Baffles/Shrouds” section later in this chapter.

### Thermal Dissipation

Heat is generated by electronic devices and must be dissipated. There are a number of techniques to accomplish this. Heatsinks are one approach with which you are probably already familiar. Although heatsinks may pull the heat out of the CPU or the motherboard, we still have to get the heat out of the case, and we do that with fans. Finally, we need to get the collected heat from all of the servers out of the server room, or at least create a flow in the room that keeps the hot air from reentering the devices.

One of the ways to do that is through the use of hot and cold aisle arrangements. The goal of a hot aisle/cold aisle configuration is to conserve energy and lower cooling costs by managing airflow. It involves lining up server racks in alternating rows with cold air intakes facing one way and hot air exhausts facing the other. The cold aisles face air conditioner output ducts. The hot aisles face air conditioner return ducts. This arrangement is shown in Figure 1.27.


**FIGURE 1.27** Hot aisle/cold aisle configuration



## Baffles/Shrouds

Another technique used both inside the case and in the server room is deploying baffles or shrouds to direct and alter the flow of air. Inside the case they are used to channel the air in the desired direction. For example, in Figure 1.28 they are used to direct the air over components that might block the desired airflow.

**FIGURE 1.28** Baffles




In the server room or datacenter, baffles may be deployed to channel the air in a desirable fashion as well. Here they are usually used to cover open rack slots, and in some cases, they are used under the raised floor to close holes there as well. Closing off these holes improves the airflow. You may have learned that open slots on the back of a tower computer should be closed with spacers. That recommendation is made for the same reason: improved airflow.

## Fans

Fans are used in several places inside the server case. There may be one on top of the heatsink used to assist the heatsink in removing the heat from the CPU. However, there will also be at least one, if not two, case fans used to move the hot air out of the case.

In server rooms and datacenters, the racks in which servers reside will probably also have multiple fans to pull the air out of the rack. An example of the fans in the back of a rack system is shown in Figure 1.29. In this instance the fans are located in an external unit that can be bought and placed on the back of a rack that either has no fans or has insufficient fans.

**FIGURE 1.29** Rack fans

## Liquid Cooling

In cases where passive heat removal is insufficient, liquid cooling may be deployed inside the case. In large datacenters this may be delivered from outside the case to the chips that need cooling. When done in this fashion, each server receives cool water from a main source, the heated water from all of the servers is returned to a central location, and then the process repeats itself. Figure 1.30 shows a server receiving liquid cooling in this way.

**FIGURE 1.30** Liquid cooling



## Summary

In this chapter we covered hardware in a server, including the topics in Objective 1 of the exam. This included a discussion of form factors such as the tower, rack, and blade server. We also discussed configuring and maintaining server components such as CPU, memory, NICs, hard drives, riser cards, and RAID controllers. We ended the chapter by exploring methods of satisfying the power and cooling requirements of servers and of the server rooms and datacenters in which they live.

## Exam Essentials

**Differentiate the server form factors.** These include tower servers; 1U, 2U, 3U, and 4U rack mount servers; and blade servers. The U in the rack server notation indicates the number of units in the rack that the servers use.

**Describe the components found inside the server.** Inside the server case you will find all of the same components you might find in a workstation, but they will be more robust and there may be more of them. These include CPU, memory, NICs, hard drives, riser cards, and RAID controllers.

**Understand the power requirements of servers.** Servers can require from 350 W (for a 1U rack mount) to 4500 W for a chassis with 14 blades in it.

**Identify and mitigate cooling issues.** Explain how to use heatsinks, fans, and baffles inside the case to eliminate the heat created by servers. In the server room or datacenter, understand how to deploy baffles and hot/cold aisles to remove heat from the room.

# Review Questions

You can find the answers in the Appendix.

1. Which term refers to the size, appearance, or dimensions of a server?
  - A. Form factor
  - B. Footprint
  - C. Physical reference
  - D. U measure
2. Which of the following is used to make maintenance easier with a rack server?
  - A. KVM
  - B. Rail kits
  - C. Baffles
  - D. Rack slot
3. How large is each U in a rack?
  - A. 19 inches
  - B. 4.445 inches
  - C. 1.75 inches
  - D. It depends on the rack.
4. What technology consists of a server chassis housing multiple thin, modular circuit boards, each of which acts as a server?
  - A. Rack servers
  - B. Towers
  - C. KVM
  - D. Blade technology
5. What type of CPU cache holds data that is waiting to enter the CPU?
  - A. L1
  - B. L2
  - C. L3
  - D. L4
6. What term describes the relationship between the internal speed of the CPU and the speed of the system bus?
  - A. CPU time
  - B. Multiplier
  - C. Differential
  - D. Coefficient

7. What term describes the time the CPU was executing in kernel mode?
  - A. User time
  - B. Steal time
  - C. System time
  - D. Idle time
8. What are revisions in CPUs called?
  - A. Service packs
  - B. Hot fixes
  - C. Base layers
  - D. Stepping levels
9. Which CPU architecture was designed for a tablet?
  - A. ARM
  - B. x86
  - C. x64
  - D. LGA
10. DDR3 memory is \_\_\_\_\_ as fast as DDR2.
  - A. Three times
  - B. Twice
  - C. Half
  - D. One-third
11. True/False: DDR doubles the rate by accessing the memory module twice per clock cycle.
12. What statement is true with regard to dual-channel memory?
  - A. Installing different size modules in the same bank will result in the modules operating in single-channel mode.
  - B. Installing different size modules in the same bank will result in the modules operating in dual-channel mode.
  - C. Installing equal size modules in the same bank will result in the modules operating in single-channel mode.
  - D. Installing different size modules in the same bank will increase the performance of the bank.
13. Which if the following is the time to access a memory address column if the correct row is already open?
  - A. CAS Latency
  - B. Row Address to Column Address Delay
  - C. Row Precharge Time
  - D. Row Active Time



14. Which of the following can be mixed when installing memory? Choose two.
- A. Different speeds
  - B. Different types
  - C. Different form factors
  - D. Different manufacturers
15. Which of the following is a double-wide version of the 32-bit PCI local bus?
- A. PCI
  - B. PCI-X
  - C. PCIe
  - D. PCI/2
16. Which type of NIC detects the type of device on the other end and changes the use of the wire pairs accordingly?
- A. Auto-MDIX
  - B. Full-duplex
  - C. Converged
  - D. HBA
17. What type of NIC acts as both a host bus adapter (HBA) for the SAN and also as the network card for the server?
- A. Auto-MDIX
  - B. Full-duplex
  - C. Converged
  - D. HBA
18. What are the two implementations of hybrid drives?
- A. Dual-drive
  - B. Single-drive
  - C. Solid-state
  - D. Dual-state
19. What is the height of a 2U system?
- A. 1.75"
  - B. 3.5"
  - C. 5.25"
  - D. 7"
20. Which statement is false with regard to UEFI?
- A. It protects the preboot process.
  - B. It has a slower startup time than BIOS.
  - C. It supports 64-bit firmware device drivers.
  - D. It supports drives larger than 2.2 terabytes (TB).

