

# Chapter

# 1

# Server Hardware

---

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

### ✓ 1.1 Given a scenario, install physical hardware.

- Enclosure sizes
- Unit sizes
  - 1U, 2U, 3U, etc.
- Rack layout
  - Cooling management
  - Safety
  - Proper lifting techniques
  - Rack balancing
  - Floor load limitations
  - Power distribution unit (PDU)
  - Keyboard-video-mouse (KVM) placement
  - Rail kits
- Power cabling
  - Redundant power
    - Uninterruptible power supply (UPS)
    - Separate circuits
    - Separate providers
  - Power connector types
  - Cable management
- Network cabling
  - Redundant networking
  - Twisted pair
  - Fiber
    - SC



- LC
  - Single mode
  - Multimode
- Gigabit
- 10 GigE
- Small form factor pluggable (SFP)
- SFP+
- Quad small form factor pluggable (QSFP)
- Cable management
- Server chassis types
  - Tower
  - Rack mount
  - Blade enclosure
- Server components
  - Hardware compatibility list (HCL)
  - Central processing unit (CPU)
  - Graphics processing unit (GPU)
  - Memory
  - Bus types
  - Interface types
  - Expansion cards



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 SK0-005 exam, otherwise known as the CompTIA Server+ exam.

## Racking

Racking is the process of installing rack systems and then populating the rack systems with rack-mounted devices such as routers, switches, servers, intrusion detection system (IDS) devices, and many other types of network gear. In the opening section of this chapter, you'll learn about issues related to racking.

### Enclosure Sizes

Racks come in different sizes, and we refer to the space the rack provides as the enclosure size. For example, the Dell PowerEdge rack enclosures are offered in three height options: 24U (2420), 42U (4220), and 48U (4820). Each U is 1.75 inches (44.45 mm) high.

### Unit Sizes


Rack systems can accommodate a variety of device sizes. Each device is a unit, and in this section you'll learn about standard unit sizes.

#### **1U, 2U, 3U, etc.**

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 Figure 1.1 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 Figure 1.1, 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 Figure 1.1 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.

## Rack Layout

Rack systems have many options to accommodate various power and cooling requirements. Moreover, there are additional components that you need to know about to successfully install and manage rack systems.

## Cooling Management

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 that 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, such as the graphics card, motherboard, and hard drives, in the way of airflow


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 1.3.


**FIGURE 1.3** 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 1.4 they are used to direct the air over components that might block the desired airflow.

**FIGURE 1.4** 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 1.5. 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 1.5** 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 the servers is returned to a central location, and then the process repeats itself. Figure 1.6 shows a server receiving liquid cooling in this way.

### Safety

Equipment can be replaced but not humans. When working with electric equipment, much of it heavy and difficult to move, proper safety practices should be followed. In this section you'll learn about safety issues.

**FIGURE 1.6** Liquid cooling

### Proper Lifting Techniques

An easy way to get hurt is by moving equipment in an unsafe or improper way. Here are some safe lifting techniques to always keep in mind:


- Lift with your legs, not your back. When you have to pick something up, bend at the knees, not at the waist. You want to maintain the natural curve of the back and spine when lifting.
- Be careful to not twist when lifting. Keep the weight on your centerline.
- Keep objects as close to your body as possible and at waist level.
- Where possible, push instead of pull.

### Rack Balancing

The racks in your server room will probably hold significant amounts of weight. The devices that reside there can be fragile, so for both your and the equipment's safety the racks must be stable. Racks that are solidly connected to both the floor and the ceiling will be the least of your worries if they are properly secure at all posts.

Considering that some racks may be on wheels, you may want to consider using rack stabilizers to prevent them from falling over. Then go underneath the rack and secure the rack to the floor, as shown in Figure 1.7.

The goal in lifting should be to reduce the strain on lower back muscles as much as possible since muscles in the lower back aren't nearly as strong as those in the legs or other parts of the body. Some people use a back belt or brace to help maintain the proper position while lifting.

**FIGURE 1.7** Rack stabilizer

### Floor Load Limitations

Whether your server room or datacenter has a raised floor (which is highly advisable) or not, you must ensure that the floor has the ability to support the weight of the equipment. You should determine the total weight of all the equipment and the racks in the room. This includes everything that will be in the room. Then when the raised floor is constructed, you must ensure it has the load capacity. If you are in an area where earthquakes are possible, you need go beyond that and consult load experts who can tell you how much additional load capacity may be required for the intensity of any of these events that occur.

Most server room or datacenter floors have tile panels that can be removed to access things such as cabling that might be under the floor. Your floor load calculations should be done with the tiles installed because they will increase the load capacity when they are all in place.

### SHARP EDGES AND PINCH POINTS

In the process of working with the servers and the racks in which they may be installed, there will be several inescapable hazards you may encounter, but a dose of patience and awareness can help you avoid them. Although manufacturers try to minimize this fact, you may encounter sharp edges that can cut you if you aren't careful. These include metal corners of cases, edges of cards, and subassemblies. If you go slow and exercise care, you can avoid cutting yourself.

There also will be occasions when certain procedures may expose you to a “pinch point,” which is simply an opportunity to pinch your fingers between two objects when performing the procedure. In some cases the documentation describing the procedure will call out this hazard, another good reason to “read the directions.” For example, Figure 1.8 shows a call-out that precedes instructions on removing the power distribution board on a Sun Netra 440 server. Always heed these warnings.

**FIGURE 1.8** Pinch point warning


**Caution** - There are several pinch points on the power supplies, so use caution when removing or installing a power supply in the system. Refer to the caution label on the power supplies for more information.

## Power Distribution Unit (PDU)

A power distribution unit (PDU) is a device that looks much like a simple power strip with multiple outlets, but depending on the model it can be much more than that. Some are large, freestanding devices. Some of the features these devices can provide besides multiple outlets are:

- Power conditioning (evening out sags and surges)
- Surge protection
- Environmental monitoring
- Case alarms

Figure 1.9 shows an example of a rack mount PDU that is installed in a rack cabinet.


**FIGURE 1.9** Rack PDU

## PDU Ratings

PDUs are rated according to the power load they can handle. You should ensure that no single PDU is connected to more devices than it can handle. So a PDU that is rated at 24 amps cannot provide more power at a time to the servers to which it is connected than that. It is also useful to know that PDUs will have two ratings in that regard. The National Electric Code, which is published by the National Fire Protection Association, requires that the continuous current drawn from a branch circuit not exceed 80 percent of the circuit's maximum rating. Therefore, PDUs have a maximum input value and a de-rated value. The de-rated value will be 80 percent of the maximum input value. So a PDU with a maximum value of 30 amps will have a de-rated value of 24 amps.

In cases where every possible effort must be made to keep the servers running, you may decide to deploy multiple rack PDUs in a rack or cabinet, as shown in Figure 1.10, where there are two PDUs.

**FIGURE 1.10** Redundant PDUs




If you do this it for redundancy, you may want to connect each to different circuits so that if one entire circuit goes dead, the rack can still be powered by the second power source.

## Keyboard-Video-Mouse (KVM) Placement

When working with servers locally—that is, standing in the same room with the server—one of the most common ways technicians connect to the server is through a KVM. A keyboard, video, and mouse (KVM) device allows you to plug multiple PCs (usually servers) into the device and to switch easily back and forth from system to system using the same mouse, monitor, and keyboard. The KVM is actually a switch that all of the systems plug into. There is usually no software to install. Just turn off all the systems, plug them all into the switch, and turn them back on; then you can switch from one to another using the same keyboard, monitor, and mouse device connected to the KVM switch. The way in which this switch connects to the devices is shown in Figure 1.11.

**FIGURE 1.11** Standard KVM switch

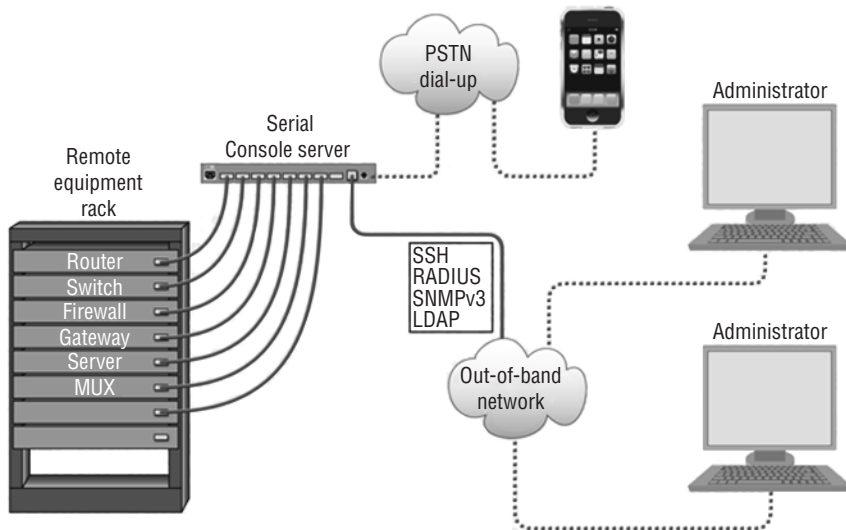


## Serial

Although serial connections have been largely replaced by the use of KVM switches, it is still possible to connect to a server using a serial connection. The issue that arises is that even if a technician's laptop had a serial port (which is unlikely today), there would be at most one. This conundrum led to the development of the serial device server. It provides a

number of serial ports, which are then connected to the serial ports of other equipment, such as servers, routers, or switches. The consoles of the connected devices can then be accessed by connecting to the console server over a serial link such as a modem, or over a network with terminal emulator software such as Telnet or SSH, maintaining survivable connectivity that allows remote users to log in the various consoles without being physically nearby. This arrangement is shown in Figure 1.12. One of the advantages of this is the ability to get to the servers “out of band.” This means even if the network is down, servers can be reached through the serial ports either locally or through the modem.

**FIGURE 1.12** Serial device server



## Rail Kits


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 1.13.

## Power Cabling

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. 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.

**FIGURE 1.13** Rail kit installation



## 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 1.14.

## 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.

**FIGURE 1.14** Voltage switch



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.

## Redundant Power

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.

## Total Potential Power Draw

The power consumption of any individual server will vary. It's unlikely but there may be a point in time when all of your devices require what's called *critical power* at the same time. You need to calculate that value and ensure that if this occurs, you have sufficient power. If the circuit is overloaded, it can cause a complete outage. This means every link in the chain, including the PDU, UPS, and the circuit to these devices, must be capable of the load. When determining that load, use the following guidelines:

- Add up the wattages listed on each device. If the wattage is not listed on the device, multiply the current (amps) by the voltage of the device to get the volt-amperes (VA).
- Multiply the VA by 0.67 to estimate the actual power, in watts, that the critical load will represent.
- Divide the number by 1,000 to establish the kilowatt (kW) load level of the anticipated critical load.
- Don't forget that your needs will not be static, so build in some additional capacity to allow for growth.

Remember, power is the lifeline of the datacenter. One of your goals is to ensure that all systems have a constant, clean source of power. In this section, we'll look at the proper use of uninterruptible power supplies (UPSs) and discuss what power distribution units (PDUs) are. We'll also talk about how to plan to ensure you have sufficient capacity to serve your devices. Finally, we'll explore the use of redundant power supplies and the use of multiple circuits to enhance availability.

## Uninterruptible Power Supply (UPS)

All systems of any importance to the continued functioning of the enterprise should be connected to a UPS. These devices have a battery attached that can provide power to the systems in the case of a power outage. You may also be aware that these systems are designed to provide only short-term power to the systems; that is a length of time sufficient to allow someone to gracefully shut down the devices. In this section, we'll dig a bit deeper and identify some of the features of these devices. We'll also go over best practices with regard to ensuring your UPS solution provides the protection you intended.

### Runtime vs. Capacity

Two important metrics that are related but *not* the same when assessing a UPS are its runtime and its capacity. The *runtime* is the amount of time that the UPS can provide power at a given power level. This means you can't evaluate this metric without knowing the amount of load you will be placing on the UPS. Documentation that comes with the UPS should reveal to you the number of minutes expected at various power levels. So if you doubled the number of like devices attached to the UPS, you should expect the time to be cut in half (actually, it will be cut more than in half in reality because the batteries discharge quicker at higher loads).

Capacity is the maximum amount of power the UPS can supply at any moment in time. So if it has a capacity of 650 VA and you attempt to pull 800 VA from the UPS, it will shut itself down. So both of the values must be considered. You need to know the total amount of power the devices may require (capacity) and, based on that figure, select a UPS that can provide that for the amount of time you will need to shut down all the devices.

One good thing to know is that some UPS vendors can supply expansion packs for existing units that increase their capacity and runtime. That would be a favorable feature to insist on to allow your system to grow.

### Automated Graceful Shutdown of Attached Devices

Today's enterprise-level UPS system tends to offer the ability to shut down a server to which it is attached when the power is lost. If all devices were thus equipped, it could reduce the amount of runtime required and eliminate the race to shut servers down.

There are several approaches that vendors have taken to this. In some cases if you purchase a special network card for the UPS, a single UPS can provide the automatic shutdown to multiple servers. The agent on each server communicates with the network card in the UPS.

Another option is to use a dedicated UPS for each server and attach the server to the UPS using a serial or USB cable. The disadvantage of this approach is that it requires a UPS for each device and you will be faced with the cable length limitations of serial and USB cables.

In either case, using the software that comes with the UPS, you can also have scripts run prior to the shutdown, and you can configure the amount of time to wait for the shutdown so that the script has time to execute, as shown in Figure 1.15. You can set a notification of this event.

**FIGURE 1.15** Automatic shutdown

**APC**  
by Schneider Electric

**Configure Shutdown**  
When PowerChute shuts down your operating system, it will use these settings. Use the Configure the system to shut down.

Yes, I want to notify a user.

Notify all users

Notify only this user:

Send notification again this often (seconds):

Run this command file (full name):

The command file needs this much time to complete (seconds):

### Periodic Testing of Batteries

Just as you would never wait until there is a loss of data to find out if the backup system is working, you should never wait until the power goes out to see whether the UPS does its job. Periodically you should test the batteries to ensure they stand ready to provide the expected runtime.

While the simplest test would be to remove power and see what happens, if you have production servers connected when you do this it could cause a resume generating event (RGE). In most cases, the software that came with the UPS will have the ability to report the current expected runtime based on the current state of the battery, as shown in Figure 1.16.

Even with this information, it is probably advisable to test the units from time to time with devices connected that you don't care about just to make sure the process of switching over to the battery succeeds and the correct runtime is provided.

### Maximum Load

Although the capacity of a UPS is rated in VA, that is not the same as maximum load. The capacity value assumes that all of the attached devices are pulling the maximum amount of power, which they rarely do. As a rule of thumb, if you multiply the VA times 0.8, you will get a rough estimate of the maximum load your UPS may undergo at any particular time. So a UPS that is rated for 650 VA cannot provide more than 520 watts. If either of these values is exceeded during operation, the UPS will fail to provide the power you need.

### Bypass Procedures

Putting a UPS in bypass mode removes the UPS from between the device and the wall output conceptually, without disconnecting it. A *static bypass* is one in which the UPS, either by the administrator invoking the bypass manually or by an inverter failure in the UPS, switches the power path back to the main line and removes itself from the line.


**FIGURE 1.16** Checking the battery level

A *maintenance bypass* is possible when the UPS is augmented with an external appliance called the bypass cabinet. This allows for enabling the bypass and then working with the UPS without concerns about the power being on (although it can be enabled while leaving the power to the UPS on). This concept is shown in Figure 1.17. Notice the two switches on the bypass cabinet that can be opened and shut to accomplish this power segregation.

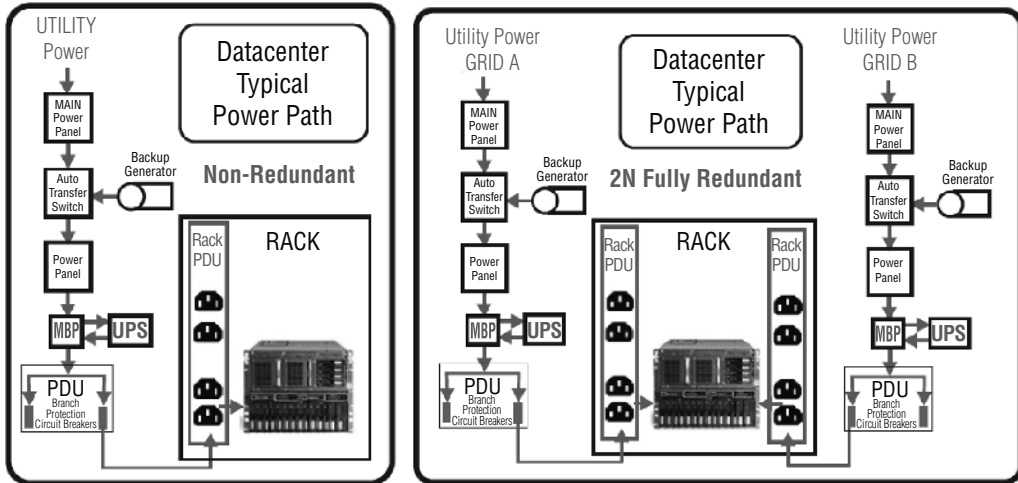
## Separate Circuits

If you have a single power circuit and it fails, you will only be up as long as your batteries last or as long as the generator can run. Many datacenters commission multiple power circuits to prevent this. A comparison of a center with a single circuit to one with two circuits is shown in Figure 1.18. In this case the engineers have gone beyond circuit redundancy and implemented a main power panel, auto transfer switch, power panel, maintenance bypass (MBP), and UPS redundancy as well. An MBP is used to bypass the UPS when either changing the UPS or performing maintenance on it.

**FIGURE 1.17** Maintenance bypass




**FIGURE 1.18** Multiple circuits



**Connect Redundant Power Supplies to Separate PDUs**

If you deploy multiple PDUs for redundancy, you may want to connect each to different circuits so that if one entire circuit goes dead, the rack can still be powered by the second power source. Another approach is to attach two UPSs to the PDU, with one going to the main power and the other to secondary power or to a generator, as shown in Figure 1.19.

**FIGURE 1.19** Redundant UPS with single PDU

## Separate Providers

All this redundancy sounds great, but what if your ISP goes down? Yes, it does happen. Many organizations engage multiple ISPs and maintain connections to all of them; that way, if one of the providers suffers an issue, a quick switchover can solve the issue.

## Power Connector 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 a 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 1.20. Keep in mind the shape of the plug may differ somewhat.

### 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.

**FIGURE 1.20** Edison plug




A sample of this connector for a 6000 W power supply is shown in Figure 1.21.

**FIGURE 1.21** Locking plug




## 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 1.22. 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 1.22** Midplane power

## Cable Management

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 1.23). These arms contain the slack and are designed to follow the server when you slide it out of the bay.

**FIGURE 1.23** Cable management arm

## Cable Ties

Throughout the datacenter or server room, you will need to organize bundles of cable and in some cases attach these bundles to trays and channels. For this operation, you will use cable ties, which come in various sizes and strengths. You should have plenty of these ties in all sizes at all times. In Figure 1.24 a variety of cable ties are shown along with a bundle of cable neatly organized using the cable ties.

**FIGURE 1.24** Cable ties



## Network Cabling

In order for all of these devices in all these racks to talk to one another we have to have network cabling, and in the datacenter there will be lots of this cabling. In this section you'll learn about network cabling and how to manage it.

### Redundant Networking

In the same way that you must build in power redundancy, you must also build in network redundancy. This could be achieved on an endpoint such as a server by installing and connecting multiple network interface cards (NICs) so if one fails additional connections are available.

Multiple connections between switches and routers in the datacenter should also be configured. Study the network diagrams and eliminate any network single points of failure.


## 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.

## Twisted Pair

Copper cabling uses electrical signals to represent the ones and zeroes in a transmission. The most common type of copper cabling in use is twisted pair cabling. There are two primary types of twisted pair cabling: shielded twisted pair (STP) and unshielded twisted pair (UTP). In both cases, the cabling consists of pairs of wires twisted around each other, as shown in Figure 1.25.

**FIGURE 1.25** Twisted pair cable



UTP offers no shielding (hence the name) and is the network cabling type most prone to outside interference. The interference can be from fluorescent light ballast, an electrical motor, or other such source (known as electromagnetic interference [EMI]) or from wires being too close together and signals jumping across them (known as crosstalk). STP adds a foil shield around the twisted wires to protect against EMI.

## Patch Cables


While any twisted pair cable can be called a patch cable, there are actually three types of patch cables (different in the way in which they are wired and in the situation in which they are required), and these cables come in specifications called categories used to describe their capabilities. Let's first look at the three ways in which these cables can be wired.

## Crossover and Straight Through

Two wiring standards are commonly used with twisted pair cabling: T568A and T568B (sometimes referred to simply as 568A and 568B). These are telecommunications standards from TIA and EIA that specify the pin arrangements for the RJ-45 connectors on UTP or STP cables. The number 568 refers to the order in which the wires within the Category 5 cable are terminated and attached to the connector. The signal is identical for both.

T568A was the first standard, released in 1991. Ten years later, in 2001, T568B was released. Figure 1.26 shows the pin number assignments for the 568A and 568B standards. Pin numbers are read left to right, with the connector tab facing down. Notice that the pin-outs stay the same, and the only difference is in the color coding of the wiring.

**FIGURE 1.26** Pin assignments for T568A and T568B

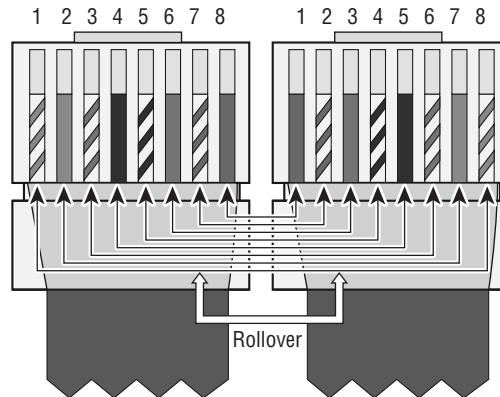


The bottom line here is that if the same standard is used on each end, the cable will be a straight-through cable, and if a different standard is used on either end, it will be a crossover cable.

Crossover cables are used to connect like devices, whereas straight-through cables are used to connect dissimilar devices. For example, to connect a router to another router or a switch to another switch, use a crossover cable. To connect a router to a switch or a computer to a switch, use a straight-through cable. There is one exception: to connect a host to a router, use crossover. Having said all that, most NICs today have the ability to sense the required pin-out pattern for the connection and use that. This function is called Auto-MDI-X.

## Rollover

A rollover cable is a cable using a completely reversed wiring pattern. It is used to connect to a router, switch, or access point console port to manage the device using a HyperTerminal application. The pin-out is shown in Figure 1.27.

**FIGURE 1.27** Rollover pin-out

## CAT5

Category 5 cabling transmits data at speeds up to 100 Mbps and is used with Fast Ethernet (operating at 100 Mbps) with a transmission range of 100 meters. It contains four twisted pairs of copper wire to give the most protection. Although it had its share of popularity (it's used primarily for 10/100 Ethernet networking), it is now an outdated standard. Newer implementations use the 5e standard.

## RJ-45 and RJ-11

Twisted pair cabling uses a connector type called an RJ (registered jack) connector. You are probably familiar with RJ connectors. Most landline phones connect with an RJ-11 connector. The connector used with UTP cable is called RJ-45. The RJ-11 has room for two pairs (four wires), and the RJ-45 has room for four pairs (eight wires). In almost every case, UTP uses RJ connectors; a crimping tool is used to attach an RJ connector to a cable. Figure 1.28 shows an RJ-11 and an RJ-45 connector.

**FIGURE 1.28** RJ-45 and an RJ-11 connector

## Fiber

Because fiber-based media use light transmissions instead of electronic pulses, such problems as EMI and crosstalk become nonissues. Fiber gets around the limitations on almost everything else except cost and is well suited for transferring data, video, and voice transmissions. Since anyone trying to access data signals on a fiber-optic cable must physically tap into the medium, it is the most secure of all cable media. It does have distance limitations based on the mode in use and the cable type.

### SC

The subscriber connector (SC), also sometimes known as a square connector, is shown in Figure 1.29. SCs are latched connectors, making it virtually impossible for you to pull out the connector without releasing its latch, usually by pressing a button or release. SCs work with either single-mode or multimode optical fibers. They aren't as popular as ST connectors for LAN connections.


### ST

The straight tip (ST) fiber-optic connector, developed by AT&T, is probably the most widely used fiber-optic connector. It uses a BNC attachment mechanism that makes connections and disconnections fairly easy. The ease of use of the ST is one of the attributes that makes this connector so popular. Figure 1.29 shows an ST connector along with an SC and a local connector.

### LC

The local connector (LC), which was developed by Lucent Technologies, is a mini form factor (MFF) connector, especially popular for use with Fibre Channel adapters, fast storage area networks, and Gigabit Ethernet adapters (see Figure 1.29).

**FIGURE 1.29** Fiber connectors ST, SC, and LC



## Single Mode/Multimode

Two types of fiber-optic cable are available: single-mode and multimode. As the name implies, single-mode uses a single direct beam of light, thus allowing for greater distances and increased transfer speeds. With multimode, a lot of light beams travel through the cable,

bouncing off the cable walls; this weakens the signal, reducing the length that the data signal can travel.

The most common types of fiber-optic cable include the following:

- 8.3 micron core/125 micron cladding single mode
- 50 micron core/125 micron cladding multimode
- 62.5 micron core/125 micron cladding multimode

### Speed and Transmission Limitations

Table 1.1 lists the speed and transmission limitations for the most common fiber-optic implementations.

**TABLE 1.1** Fiber speeds and limitations

Characteristic	100BaseFX	1000BaseSX	1000BaseLX	10GBaseER
Speed	100 Mbps	1000 Mbps	1000 Mbps	10,000 Mbps
Distance (multimode)	412 meters	220 to 550 meters	550 meters	(not used)
Distance (single mode)	10,000 meters	(not used)	5 km	40 km

## Gigabit

CAT5e cabling transmits data at speeds up to 1 Gbps (1000 Mbps). Category 5e cabling can be used up to 100 meters, depending on the implementation and standard used, and it provides a minimum of 100 MHz of bandwidth. It also contains four twisted pairs of copper wire, but they're physically separated and contain more twists per foot than Category 5 to provide maximum interference protection.

## 10 GigE

CAT6 cabling transmits data at speed up to 10 Gbps, has a minimum of 250 MHz of bandwidth, and specifies cable lengths up to 100 meters (using CAT6a). It contains four twisted pairs of copper wire and is used in 10GBaseT networks. Category 6 cable typically consists of four twisted pairs of copper wire, but its capabilities far exceed those of other cable types. Category 6 twisted pair uses a longitudinal separator, which separates each of the four pairs of wires from each other and reduces the amount of crosstalk possible.

## Small Form-Factor Pluggable (SFP)

The small form-factor pluggable (SFP) is a compact, hot-pluggable transceiver that, though not standardized by any governing body like the IEEE, was created through a multisource agreement (MSA) between competing manufacturers. For this reason, you may find that there is not full compatibility among these from various sources.

These devices allow for adding functionality to a device. For example, you plug in a fiber SFP into an open SFP slot in a device and add a fiber connection where there was none. SFP sockets are found in Ethernet switches, routers, firewalls, and network interface cards. You will find them in storage devices as well. An example of a fiber SFP is shown in Figure 1.30. The example shows a fiber cable plugged into the SFP module, which is then plugged into the SFP slot. The SFP slot is shown removed from a generic slot on a Cisco device.

**FIGURE 1.30** Fiber SFP




SFP cards can be added to servers if you need to add a connection type currently not present. These cards can be added to a PCI-Express slot. An example of one of these cards is shown in Figure 1.31. This particular model accepts two SFP+ connectors, requiring either Direct Attach Copper Cable (DAC) for copper environments, or fiber transceivers supporting short-haul (SR) optics plus fiber cables for fiber-optic environments.

## SFP+

In the previous paragraph we made passing mention of SFP+. So what's the difference between SFP and SFP+? Well, they look the same but at introduction, typical speeds were 1 Gbit/s for Ethernet SFPs and up to 4 Gbit/s for Fibre Channel SFP modules. In 2006, SFP+ specification brought speeds up to 10 Gbit/s, and the SFP28 iteration is designed for speeds of 25 Gbit/s.

**FIGURE 1.31** HP two-port server adapter



## Quad Small Form-Factor Pluggable (QSFP)

The quad small form-factor pluggable (QSFP) transceiver is a version of SFP with four lanes, allowing for speeds 4 times their corresponding SFP. In 2014, the QSFP28 variant was published, allowing speeds up to 100 Gbit/s. In 2019, the closely related QSFP56 was standardized, doubling the top speeds to 200 Gbit/s.

## Cable Management

It can be time-consuming to tie cables up, run them in channels, and snake them through walls, but it is time well spent when it keeps one person from harm. It is all too easy to get tangled in a cable or trip over one that is run across the floor. Take the extra time to manage cables, and it will increase your safety as well as that of others who work in that environment. These final sections will discuss cable management in the server room or datacenter.

## Cable Channels

Cable channels are used to route cables across floors and other surfaces. They enclose the cables and protect them from damage while also preventing someone from tripping over them. In some cases, these trays may be integrated into the floor of the datacenter as well. An example of a cable channel is shown in Figure 1.32.

**FIGURE 1.32** Cable channel



## Cable Management Trays



While cable channels are good for floors, in some cases cable must be run up and over racks of equipment. In such instances you can use cable trays, which come in two types: vertical and horizontal.

### Vertical

Vertical trays are used to route a group of cables up or down. For example, you may need to get cable from the back of a rack to the ceiling, where it may intersect with horizontal trays that lead the cabling out of the room or to another rack in the room. With vertical installation, cables on cable trays must be fixed by clips or suitable binding materials. An example of a vertical tray is shown in Figure 1.33.

### Horizontal

Horizontal trays are used to route cable across the room on the floor, the ceiling, or anywhere in between. A horizontal tray is shown in Figure 1.34.

**FIGURE 1.33** Attaching cable to vertical trays**FIGURE 1.34** Horizontal cable tray

## Server Chassis Types

As you learned earlier, servers can come in different form factors, also called chassis types. 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. In this section you'll also learn about chassis security.

### Tower

A 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 1.35.

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 1.36.

**FIGURE 1.35** Tower server




**FIGURE 1.36** Tower server in a rack



## Rack Mount

Rack-mounted 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 1.37.

**FIGURE 1.37** Server in a rack



Rack-mounted servers may come with locks that prevent removing the server from the rack without opening the lock. Whereas many of these locks are on cabinets that enclose the servers, as discussed in the previous section, others are a part of the server case itself, such as the one shown in Figure 1.38.

**FIGURE 1.38** Rack mount lock



### Server

You may also have servers you need to physically secure from tampering that are not located in racks or lockable server cabinets. Perhaps you are still using some tower servers. These servers can be secured using a lockable rack such as the one you see in Figure 1.39. A tower server lock connects to a cable that you secure to an immovable object as you would secure a laptop.

**FIGURE 1.39** Tower server lock



You also can secure a tower in a lockable cabinet made for just such a situation, as shown in Figure 1.40. Usually these come with a ventilation system to keep the server cool.


**FIGURE 1.40** Tower cabinet



## 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 1.41.


**FIGURE 1.41** 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 1.42, a blade server is shown being inserted into an enclosure.

**FIGURE 1.42** Blade enclosure



## 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 1.43.

## Chassis Locks/Intrusion Detection


Although securing the BIOS can keep the server from being accessed by booting to another operating system, it cannot prevent theft of memory or hard drives. You should put locks on the case to prevent it from being opened. You should also use settings in the BIOS to alert

you when the case has been opened. These settings are shown in Figure 1.44. Here the open case warning has *not* yet been enabled.

**FIGURE 1.43** Inserting a blade server and filler



**FIGURE 1.44** Open case warning in the BIOS



## 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.

### Hardware Compatibility List (HCL)

When deploying a new server or when adding devices and applications to a server, it can be highly beneficial to ensure compatibility between the new addition and the server before spending money. Vendors of both software and hardware create compatibility lists that can be used to ensure that a potential piece of software or hardware will work with the server. These compatibility lists fall into three categories, as covered in the following sections.

#### Operating Systems

Major operating system vendors issue several types of compatibility lists. Some list the hardware requirements of each operating system. Others list hardware devices that have been tested and are known to work with the operating system. Microsoft calls its list the Windows Compatible Products List. You'll find it here:

<https://docs.microsoft.com/en-us/windows-hardware/drivers/dashboard/windows-certified-products-list>

#### Hardware

Vendors of the hardware we often connect to our servers also create compatibility lists that describe the operating systems and other pieces of hardware with which their devices are compatible.

#### Applications

Finally, vendors of software applications also issue compatibility lists that describe the operating systems on which their software will run and the hardware requirements of the system on which the software will be installed.

### Central Processing Unit (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 1.2.

**TABLE 1.2** Server socket types

Socket name	CPU families supported	Package	Pin count	Bus speed
LGA 771/ Socket J	Intel Xeon	LGA	771	1600 MHz
LGA 1366/ Socket B	Intel Core i7 (900 series) Intel Xeon (35xx, 36xx, 55xx, 56xx series), Intel Celeron	LGA	1366	4.8–6.4 GT/s (gigatransfers 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
LGA 2011/ Socket R	Intel Core i7 3xxx Sandy Bridge-E Intel Core i7 4xxx Ivy Bridge-E Intel Xeon E5 2xxx/4xxx [Sandy Bridge EP] (2/4S) Intel Xeon E5-2xxx/4xxx v2 [Ivy Bridge EP] (2/4S)	LGA	2011	4.8–6.4 GT/s
Socket F	AMD Opteron 13xs, 2200, 2300, 2400, 8200, 8300, 8400, AMD 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, Sempron, Opteron 3xxx	PGA-ZIF	942	3.2 GHz

Notice in Table 1.2 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 1.45.

**FIGURE 1.45** 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 1.2) 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** 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 four) holding the heatsink, and then lift it off 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).
  10. Replace the server node cover.
  11. Reinstall the server node.
  12. Replace power cords and then power on the system.
- 

## Graphics Processing Unit (GPU)

The graphics processing unit (GPU) is a processor dedicated to working on behalf of generating graphic images. They do a better job of this than the system CPU, and they offload some of the work from the CPU, allowing it to operate more efficiently as well.

## Memory

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 1.3.

**TABLE 1.3** Selected memory details

Module standard	Speed	Format
DDR-500	4,000 MBps	PC4-000
DDR-533	4,266 MBps	PC4-200
DDR2-667	5,333 MBps	PC2-5300
DDR2-750	6,000 MBps	PC2-6000
DDR2-800	6,400 MBps	PC2-6400
DDR3-800	6,400 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 (not shown in Figure 1.46) banks, as shown in Figure 1.46, and other times they use the same colors. Consult your documentation.

**FIGURE 1.46** 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 to 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, whereas 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.

## Replacing RAM

Replacing RAM in a server is not all that different from doing so in a workstation. The box looks different but otherwise the basic steps are the same. In Exercise 1.2 you'll use an IBM blade server.




This exercise applies to an IBM blade server. The procedure for your server may vary, so consult the documentation.

### EXERCISE 1.2

#### Changing RAM in an IBM Blade Server

1. If the blade server is installed in a BladeCenter unit, remove it from the BladeCenter unit.
2. Remove the blade server cover.
3. If an optional expansion unit is installed, remove the expansion unit. This may require an extraction device (thumbscrews or levers) and it may not. If one is not provided, using the blade server cover releases on each side, lift the expansion unit from the blade server, as shown here:




**EXERCISE 1.2 (continued)**

4. Locate the DIMM connectors. Determine which DIMM you want to remove from the blade server.



To avoid breaking the retaining clips or damaging the DIMM connectors, handle the clips gently.

5. Move the retaining clips on the ends of the DIMM connector to the open position by pressing the retaining clips away from the center of the DIMM connector. To access DIMM connectors 7 through 12, use your fingers to lift the DIMM access door, as shown here:



6. Using your fingers, pull the DIMM out of the connector.
7. Install a DIMM or DIMM filler in each empty DIMM connector.

Note that DIMMs or DIMM fillers must occupy DIMM connectors 1, 2, 13, 14, 15, and 16 for proper cooling on the IBM HS22 blade server.

## Bus Types

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. You will learn more about PCI slots in the next section.


## Interface Types

In this section we'll look at some internal and external interface types.

### 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 1.47 shows some PCI slots.

**FIGURE 1.47** PCI slots



PCI cards that are 32 bit with 33 MHz operate up to 133 MBps, whereas 32-bit cards with 64 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 that 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 six versions of PCIe: version 1 is up to 8 GBps, version 2 is up to 16 GBps, version 3 is up to 32 GBps, version 4 is up to 64 GBps, version 5 is up to 128 GBps, and version 6 is up to 256 GBps. Figure 1.48 shows the slots discussed so far in this section, and Table 1.4 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 1.48** Comparison of PCI slot types



**TABLE 1.4** 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. 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 1.49.

**FIGURE 1.49** Traditional and CNA



## Expansion Cards

Expansion cards allow you to add additional device support to the system. Let's look at some types of these cards.

### 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-mounted 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 1.50.


**FIGURE 1.50** 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 1.51.

**FIGURE 1.51** RAID controller

## 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-mounted 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. Outline
2. Which of the following is used to make physical 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.
  - A. True
  - B. False
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 of 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 the network card for the server?
- A. Auto-MDIX
  - B. Full-duplex
  - C. Converged
  - D. HBA
18. Which of the following supplies power connections to various components?
- A. Cable management arm
  - B. Front plane
  - C. Midplane
  - D. MBP
19. What is the height of a 2U system?
- A. 1.75"
  - B. 3.5"
  - C. 5.25"
  - D. 7"

20. Which of the following 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?
- A. DLP
  - B. RAID
  - C. UEFUI
  - D. SFP