PEARSON IT CERTIFICATION

Save 10% on Exam Voucher

See Inside

Practice Tests

Flash Cards

Review Exercises

Study Planner

Cert Guide

Advance your IT career with hands-on learning

CompTIA®

Security+

SY0-601

OMAR SANTOS

RON TAYLOR

JOSEPH MLODZIANOWSKI

Special Offer

Save 80% on Premium Edition eBook and Practice Test

The *CompTIA Security+ SY0-601 Cert Guide Premium Edition eBook and Practice Test* provides three eBook files (PDF, EPUB, and MOBI/Kindle) to read on your preferred device and an enhanced edition of the Pearson Test Prep practice test software. You also receive two additional practice exams with links for every question mapped to the PDF eBook.

See the card insert in the back of the book for your Pearson Test Prep activation code and special offers.

10. What is the uptime measurement of available time for a highly available system, knowing that no system can be up forever?
    a. 95.99 percent
    b. 100 percent
    c. 90 percent
    d. 99.999 percent

11. Your company just had a catastrophic failure, and the entire data center and campus are down. After you establish your chain of command, what is the first service or system you need to bring up in order to start recovery?
    a. Video telecom VTC systems and bridges
    b. Desktop computers in network operations center
    c. Active Directory and LDAP to authorized users
    d. Network connectivity

12. An organization can add many levels of diversity to enhance its systems and network resilience. Which components will help? (Select all that apply.)
    a. Redundant Pathed Facility Power
    b. Network interface card (NIC) teaming
    c. Supplier and supply chain contracts
    d. Security cameras at all junctions

Foundation Topics

Redundancy

When systems are not designed or configured with redundancy in mind, a single point of failure can bring an entire company to its knees. High availability is provided through *redundant* systems and fault tolerance. If the primary method used to store data, transfer information, or perform other operations fails, then a secondary method is used to continue providing services. Redundancy ensures that systems are always available in one way or another, and with minimal downtime. Several key elements are used in the redundancy process:

- ***Geographical dispersal*:** Placing valuable data assets around the city, state, country, or world can provide an extra level of protection.
- ***Disk redundancy*:** Various levels of disk redundancy exist to reduce data loss from disk failure.
- ***Network redundancy*:** Networks deliver data to users during component failures.
- **Power:** Power resilience provides power via multiple sources.

Geographic Dispersal

Key Topic

We live in a connected world, with employees, contractors, servers, data, and connected resources dispersed throughout the world. This dispersal is essential for a company with a global presence (marketing); it also lends itself to resilience. Local issues are likely to be isolated, allowing the rest of the workforce to continue operating as normal. Companies that have workforces dispersed in remote locations throughout the world have adopted and learned to utilize technology to connect people and resources. To implement geographical dispersal within your company for compute resources, you can set up a computing environment, such as a data center or resource center, in another location at least fifty miles away from the main compute campus, data center. You also should consider placing resources closer to groups of users who might be in another city or country. By way of a wide-area network, you can configure resources that are local to the users to respond first, thereby providing resilience and resource dispersal.

Disk Redundancy

Disk redundancy is simply writing data to two or more disks at the same time. Having the same data stored on separate disks enables you to recover the data in the event of a disk failure without resorting to expensive or time-consuming data recovery techniques. Implementing disk redundancy starts with selecting the appropriate level of RAID and a controller that is capable of performing RAID functions at various levels. Newer servers have RAID-capable controllers, thus allowing you to develop a RAID strategy and implement the appropriate level of RAID.

Key Topic

Redundant Array of Inexpensive Disks

Redundant Array of Inexpensive Disks (RAID) 1 and RAID 5 are the most common approaches to disk redundancy. RAID is a group of disks or solid-state drives (SSDs) that increases performance or provides fault tolerance or both. RAID uses two or more physical drives and a RAID controller, which is plugged into a motherboard that does not have RAID circuits. Today, most motherboards have built-in RAID, but not necessarily every RAID level of configuration is available. In the past, RAID was also accomplished through software only, but was much slower. Table 13-2 lists a few of the most common levels of RAID and their safeguards, configurations, and protections. Implementing the appropriate RAID level begins with an overall corporate strategy.

Table 13-2 Common RAID Levels and Characteristics

Raid Level	Description	Data Reliability/Speed	Protection
0—Striping	Disk striping; divides data across a set of hard disks.	Is lower than a single disk; enables writes and reads to be done more quickly.	Very low; if a disk fails, you could lose all your data.
1—Mirroring	Duplicates data from one hard disk to another, typically two hard disks.	Provides redundancy of all data because it is duplicated; is much slower.	A complete copy of data is on both hard disks.
5—Striping with parity	Uses more hard disks and stripes data and parity over all hard disks.	Provides higher performance; quickly writes data to all disks.	Data is protected by the parity writes.
6—Striping with dual parity	Similar to RAID 5; includes a second parity distributed across hard disks.	Provides slower storage system performance; provides best protection.	Dual parity protects against loss of two hard disks.
10—Mirroring a striped set of disks	Similar to RAID 1 and RAID 0 where you start with a striped set of disks and then mirror them.	Provides higher performance and fault tolerance than other RAID levels.	It requires more disks.

TIP To obtain more data on RAID, see https://datatracker.ietf.org/wg/raidmib/about/.

RAID 0 Striping

RAID 0 striping provides no fault tolerance or redundancy; the failure of one drive will cause the entire array to fail. As a result of having data striped across *x* number of disks, the failure will be a total loss of data. RAID 0 is normally used to increase performance and provide additional space (larger than a single drive) because RAID 0 appears as a single disk to the user. To implement RAID 0, you need to boot your computer to the BIOS and look for Disks and Settings for RAID (see Figure 13-1). There, you select the RAID level. In the example shown here, the ASUS motherboard has a built-in RAID controller like most modern motherboards do (see Figure 13-2). In this case, you select Create a RAID Volume, enter the name of the RAID set, select the RAID level, and select the disks to be part of the RAID set. Then you select Create Volume. Be warned that this will delete all data on all disks.

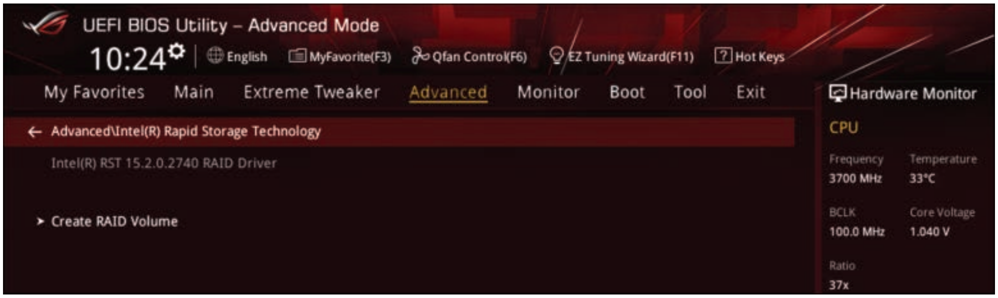

FIGURE 13-1 Bios RAID Controller

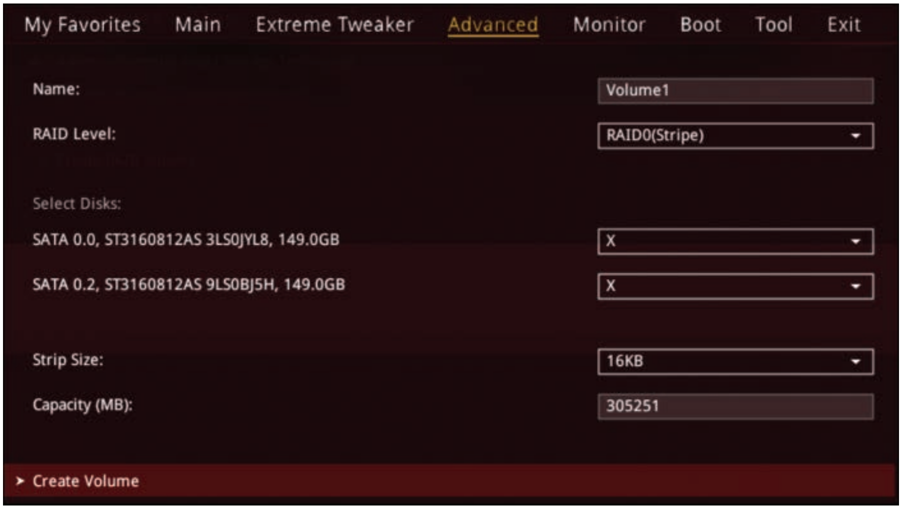

FIGURE 13-2 Creating a Volume

Your motherboard might be different. In fact, different ASUS models are slightly different in layout and configuration, but the steps are virtually the same.

RAID 1—Mirror

RAID 1—Mirror is exactly that—one or more disks with an additional disk to act as a mirror image for each disk. A classic RAID 1 mirrored disk pair contains two disks. The configuration offers no striping, no parity, and no spanning disks across multiple sets. Because all disks are mirrored, you need to match disk sizes to obtain the most benefit. The smallest disk in a mirror set is the largest the array can be. The performance of the RAID 1 array depends on the drive array, controller, slowest disk speed, and I/O load. Typically, you will see slightly slower speeds than a single drive. To configure RAID 1, you follow the steps described for Figure 13-1 and Figure 13-2, and select RAID 1 Mirror for the RAID level.

RAID 5—Striping with Parity

RAID 5—Striping with Parity consists of block-level striping with distributed parity among all drives. RAID 5 requires at least three disks and can suffer only a single drive failure before data becomes corrupted. The disk parity can be distributed in three different ways; only more expensive controllers allow you to select which method. The default sequence of data blocks is written left to right or right to left on the disk array of disks 0–X, also called Left Asynchronous RAID 5. Performance of RAID 5 is increased because all RAID members participate in the servicing of write requests. To configure and implement RAID 5, you follow the steps outlined for Figure 13-1 and Figure 13-2 and select RAID 5 Striping with Parity for the RAID level, of course taking into account your specific motherboard capabilities.

RAID 6—Striping with Double Parity

RAID 6—Striping with Double Parity consists of block-level striping with two parity blocks distributed across all member disks. Similar to RAID 5, there are different capabilities for how data is striped over the disk. None of them provide any greater performance over the other, and RAID 6 takes a performance penalty during read operations but is similar in the write speeds. You can suffer the loss of two disks using RAID 6 and still remain operational. To configure RAID 6, you follow the steps outlined for Figure 13-1 and Figure 13-2 and select RAID 6 for the RAID level. Also, note that higher-end server motherboards and controllers typically only have this RAID level capability.

RAID 10—Stripe and Mirror

RAID 10—Stripe and Mirror is the combination of creating a stripe and then mirroring it. RAID 10 is often called RAID 1 + 0 and array of mirrors, which may be two- or three-way mirrors. This configuration requires a minimum of four drives, and performance-wise it beats all RAID levels with better throughput and lower latency, except for RAID 0. To configure RAID 10, you need to consult your server manufacturer's guide on creating RAID 10 arrays.

Multipath

Key Topic

Storage ***multipath*** involves the establishment of multiple physical routes; for example, multipath I/O defines more than one physical path between the CPU in a computer and its mass-storage devices through the buses, controllers, and bridge devices connecting them. Storage area networks (SANs), covered later in this chapter, are primary examples of storage multipaths. When building a SAN, you create two paths from the server to the SAN switch, providing control from the SAN switch to the storage controller all the way to the disk.

Network Resilience

Network resilience involves more than just redundancy; it enables service operations in the face of outages, faults, and challenges. Threats and challenges for services can range from simple misconfiguration to large-scale natural disasters to targeted attacks. Network availability can be handled from multiple network cards, load balancers, and redundant network hardware. You can implement network resilience in multiple ways, as mentioned previously, such as having multiple network cards in each server, using load balancers to offload traffic to hosts that are less busy, and adding in layers of redundant networking components like switches and routers that are configured in high-availability configurations.

Key Topic

Load Balancers

There are several types of load balancers. Network load balancers automatically distribute incoming traffic across multiple network paths. This capability increases the availability of your network during component outages, misconfigurations, attacks, and hardware failures. Most network load balancers utilize routing protocols to help make path decisions during a network overload or network failure. A load balancer spreads out the network load to various switches, routers, and servers. Server load balancers distribute the traffic load based on specific selected algorithms to each of the servers in a group. Load balancers are typically placed in front of a group of servers to evenly distribute the load among them.