fault tolerance - The capability of a system to continue performing when there is a hardware failure.

redundant arrays of independent drives (RAID) - A fault tolerance technique that records data on multiple disk drives instead of just one to reduce the risk of data loss.

risk of hardware and software failure. The use of redundant components provides **fault tolerance**, which is the ability of a system to continue functioning in the event that a particular component fails. For example, many organizations use **redundant arrays of independent drives (RAID)** instead of just one disk drive. With RAID, data is written to multiple disk drives simultaneously. Thus, if one disk drive fails, the data can be readily accessed from another.

COBIT 2019 management practices DSS01.04 and DSS01.05 address the importance of locating and designing the data centers housing mission-critical servers and databases so as to minimize the risks associated with natural and human-caused disasters. Common design features include the following:

- Raised floors provide protection from damage caused by flooding.
- Fire detection and suppression devices reduce the likelihood of fire damage.
- Adequate air-conditioning systems reduce the likelihood of damage to computer equipment due to overheating or humidity.
- Cables with special plugs that cannot be easily removed reduce the risk of system damage due to accidental unplugging of the device.
- Surge-protection devices provide protection against temporary power fluctuations that might otherwise cause computers and other network equipment to crash.
- An **uninterruptible power supply (UPS)** system provides protection in the event of a prolonged power outage, using battery power to enable the system to operate long enough to back up critical data and safely shut down. (However, it is important to regularly inspect and test the batteries in a UPS to ensure that it will function when needed.)
- Physical access controls reduce the risk of theft or damage.

uninterruptible power supply (UPS) - An alternative power supply device that protects against the loss of power and fluctuations in the power level by using battery power to enable the system to operate long enough to back up critical data and safely shut down.

Training can also reduce the risk of system downtime. Well-trained operators are less likely to make mistakes and will know how to recover, with minimal damage, from errors they do commit. That is why COBIT 2019 management practice DSS01 stresses the importance of defining and documenting operational procedures and ensuring that IT staff understand their responsibilities.

System downtime can also occur because of computer malware (e.g., ransomware can make applications and data inaccessible). Therefore, it is important to install, run, and keep current antivirus and anti-spyware programs. These programs should be automatically invoked not only to scan e-mail, but also any removable computer media (CDs, DVDs, USB drives, etc.) that are brought into the organization. A patch management system provides additional protection by ensuring that vulnerabilities that can be exploited by malware are fixed in a timely manner.

RECOVERY AND RESUMPTION OF NORMAL OPERATIONS

The preventive controls discussed in the preceding section can minimize, but not entirely eliminate, the risk of system downtime. Hardware malfunctions, software problems, or human error can cause data to become inaccessible. For example, RAID devices can experience catastrophic failures, rendering all the drives useless. That's why senior management needs to answer two fundamental questions:

1. How much data are we willing to recreate from source documents (if they exist) or potentially lose (if no source documents exist)?
2. How long can we function without our information system?

recovery point objective (RPO) - The amount of data the organization is willing to reenter or potentially lose.

recovery time objective (RTO) - The maximum tolerable time to restore an organization's information system following a disaster, representing the length of time that the organization is willing to attempt to function without its information system.

Figure 13-1 shows the relationship between these two questions. When a problem occurs, data about everything that has happened since the last backup is lost unless it can be reentered into the system. Thus, management's answer to the first question determines the organization's **recovery point objective (RPO)**, which represents the maximum amount of data that the organization is willing to have to reenter or potentially lose. The RPO is inversely related to the frequency of backups: the smaller the desired RPO, the more frequently backups need to be made. The answer to the second question determines the organization's **recovery time objective (RTO)**, which is the maximum tolerable time to restore an information system after a disaster. Thus, the RTO represents the length of time that the organization is willing to attempt to function without its information system.

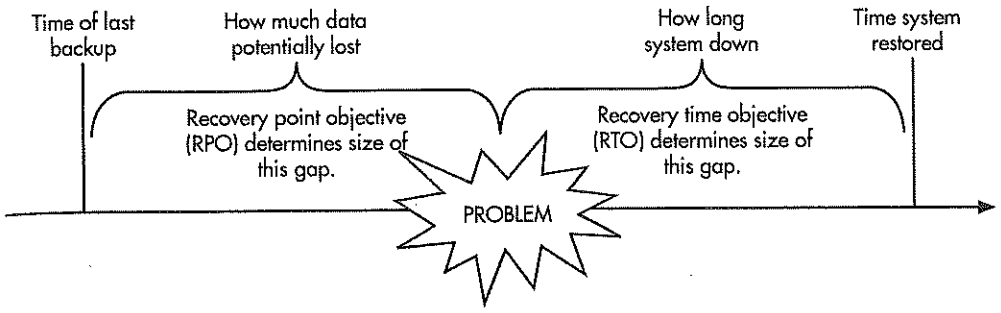

FIGURE 13-1
Relationship of Recovery Point Objective and Recovery Time Objective

DATA BACKUP PROCEDURES Data backup procedures are designed to deal with situations where information is not accessible because the relevant files or databases have become corrupted as a result of hardware failure, software problems, or human error, but the information system itself is still functioning. Several different backup procedures exist. A **full backup** is an exact copy of the entire database. Full backups are time-consuming, so most organizations only do full backups weekly and supplement them with daily partial backups. Figure 13-2 compares the two types of daily partial backups:

full backup - Exact copy of an entire database.

incremental backup - A type of partial backup that involves copying only the data items that have changed since the last *partial* backup. This produces a set of incremental backup files, each containing the results of one day's transactions.

differential backup - A type of partial backup that involves copying all changes made since the last full backup. Thus, each new differential backup file contains the cumulative effects of all activity since the last full backup.

deduplication - A process that uses hashing to identify and backup only those portions of a file or database that have been updated since the last backup.

1. An **incremental backup** involves copying only the data items that have changed since the last *partial* backup. This produces a set of incremental backup files, each containing the results of one day's transactions. Restoration involves first loading the last full backup and then installing each subsequent incremental backup in the proper sequence.
2. A **differential backup** copies all changes made since the last *full* backup. Thus, each new differential backup file contains the cumulative effects of all activity since the last full backup. Consequently, except for the first day following a full backup, daily differential backups take longer than incremental backups. Restoration is simpler, however, because the last full backup needs to be supplemented with only the most recent differential backup, instead of a set of daily incremental backup files.

Deduplication is causing many organizations to eliminate making full backups and to continuously make incremental partial backups instead. **Deduplication** uses hashing to identify

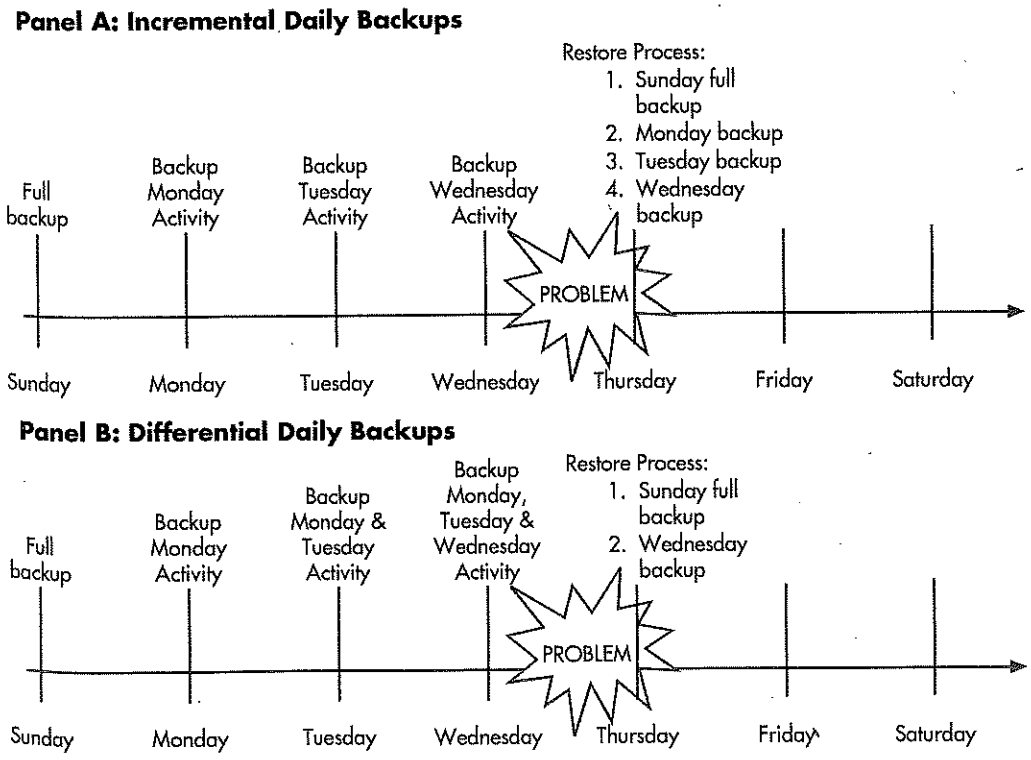

FIGURE 13-2
Comparison of Incremental and Differential Daily Backups