Chapter 11
Introduction to Business Continuity

KEY CONCEPTS

- Business Continuity
- Information Availability
- Disaster Recovery
- Disaster Restart
- BC Planning
- Business Impact Analysis

Continuous access to information is a must for the smooth functioning of business operations today, as the cost of business disruption could be catastrophic. There are many threats to information availability, such as natural disasters (e.g., flood, fire, earthquake), unplanned occurrences (e.g., cybercrime, human error, network and computer failure), and planned occurrences (e.g., upgrades, backup, restore) that result in the inaccessibility of information. It is critical for businesses to define appropriate plans that can help them overcome these crises. Business continuity is an important process to define and implement these plans.

Business continuity (BC) is an integrated and enterprisewide process that includes all activities (internal and external to IT) that a business must perform to mitigate the impact of planned and unplanned downtime. BC entails preparing for, responding to, and recovering from a system outage that adversely affects business operations. It involves proactive measures, such as business impact analysis and risk assessments, data protection, and security, and reactive countermeasures, such as disaster recovery and restart, to be invoked in the event of a failure. The goal of a business continuity solution is to ensure the "information availability" required to conduct vital business operations.

This chapter describes the factors that affect information availability. It also explains how to create an effective BC plan and design fault-tolerant mechanisms to protect against single points of failure.

11.1 Information Availability

Information availability (IA) refers to the ability of the infrastructure to function according to business expectations during its specified time of operation. Information availability ensures that people (employees, customers, suppliers, and partners) can access information whenever they need it. Information availability can be defined with the help of reliability, accessibility and timeliness.

- **Reliability:** This reflects a component's ability to function without failure, under stated conditions, for a specified amount of time.
- **Accessibility:** This is the state within which the required information is accessible at the right place, to the right user. The period of time during which the system is in an accessible state is termed *system uptime;* when it is not accessible it is termed *system downtime*.
- **Timeliness:** Defines the exact moment or the time window (a particular time of the day, week, month, and/or year as specified) during which information must be accessible. For example, if online access to an application is required between 8:00 AM and 10:00 PM each day, any disruptions to data availability outside of this time slot are not considered to affect timeliness.

11.1.1 Causes of Information Unavailability

Various planned and unplanned incidents result in data unavailability. *Planned outages* include installation/integration/maintenance of new hardware, software upgrades or patches, taking backups, application and data restores, facility operations (renovation and construction), and refresh/migration of the testing to the production environment. *Unplanned outages* include failure caused by database corruption, component failure, and human errors.

Another type of incident that may cause data unavailability is natural or man-made disasters such as flood, fire, earthquake, and contamination. As illustrated in Figure 11-1, the majority of outages are planned. Planned outages are expected and scheduled, but still cause data to be unavailable. Statistically, less than 1 percent is likely to be the result of an unforeseen disaster.

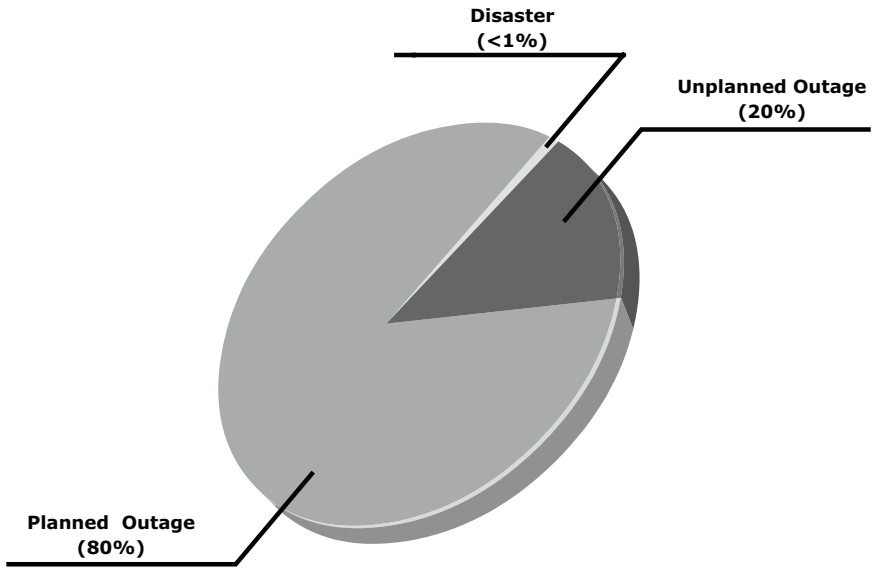

Figure 11-1: Disruptors of data availability

11.1.2 Measuring Information Availability

Information availability relies on the availability of the hardware and software components of a data center. Failure of these components might disrupt information availability. A failure is the termination of a component's ability to perform a required function. The component's ability can be restored by performing an external corrective action, such as a manual reboot, a repair, or replacement of the failed component(s). Repair involves restoring a component to a condition that enables it to perform a required function within a specified time by using procedures and resources. Proactive risk analysis performed as part of the BC planning process considers the component failure rate and average repair time, which are measured by MTBF and MTTR:

- **Mean Time Between Failure (MTBF):** It is the average time available for a system or component to perform its normal operations between failures.
- **Mean Time To Repair (MTTR):** It is the average time required to repair a failed component. While calculating MTTR, it is assumed that the fault responsible for the failure is correctly identified and that the required spares and personnel are available. Note that a fault is a physical defect

at the component level, which may result in data unavailability. MTTR includes the time required to do the following: detect the fault, mobilize the maintenance team, diagnose the fault, obtain the spare parts, repair, test, and resume normal operations.

IA is the fraction of a time period that a system is in a condition to perform its intended function upon demand. It can be expressed in terms of system uptime and downtime and measured as the amount or percentage of system uptime:

IA = system uptime / (system uptime + system downtime)

In terms of MTBF and MTTR, IA could also be expressed as

IA = MTBF / (MTBF + MTTR)

Uptime per year is based on the exact timeliness requirements of the service, this calculation leads to the number of "9s" representation for availability metrics. Table 11-1 lists the approximate amount of downtime allowed for a service to achieve certain levels of 9s availability.

For example, a service that is said to be "five 9s available" is available for 99.999 percent of the scheduled time in a year (24 × 7 × 365).

Table 11-1: Availability Percentage and Allowable Downtime

UPTIME (%)	DOWNTIME (%)	DOWNTIME PER YEAR	DOWNTIME PER WEEK
98	2	7.3 days	3 hr 22 minutes
99	1	3.65 days	1 hr 41 minutes
99.8	0.2	17 hr 31 minutes	20 minutes 10 sec
99.9	0.1	8 hr 45 minutes	10 minutes 5 sec
99.99	0.01	52.5 minutes	1 minute
99.999	0.001	5.25 minutes	6 sec
99.9999	0.0001	31.5 sec	0.6 sec

11.1.3 Consequences of Downtime

Data unavailability, or downtime, results in loss of productivity, loss of revenue, poor financial performance, and damages to reputation. Loss of productivity reduces the output per unit of labor, equipment, and capital. Loss of revenue includes direct loss, compensatory payments, future revenue losses, billing losses, and investment losses. Poor financial performance affects revenue

recognition, cash flow, discounts, payment guarantees, credit rating, and stock price. Damages to reputation may result in a loss of confidence or credibility with customers, suppliers, financial markets, banks, and business partners. Other possible consequences of downtime include the cost of additional equipment rental, overtime, and extra shipping.

The business impact of downtime is the sum of all losses sustained as a result of a given disruption. An important metric, *average cost of downtime per hour*, provides a key estimate in determining the appropriate BC solutions. It is calculated as follows:

> Average cost of downtime per hour = average productivity loss per hour + average revenue loss per hour
>
> Where:
>
> Productivity loss per hour = (total salaries and benefits of all employees per week) / (average number of working hours per week)
>
> Average revenue loss per hour = (total revenue of an organization per week) / (average number of hours per week that an organization is open for business)

The average downtime cost per hour may also include estimates of projected revenue loss due to other consequences such as damaged reputations and the additional cost of repairing the system.

11.2 BC Terminology

This section introduces and defines common terms related to BC operations and are used in the next few chapters to explain advanced concepts:

- **Disaster recovery:** This is the coordinated process of restoring systems, data, and the infrastructure required to support key ongoing business operations in the event of a disaster. It is the process of restoring a previous copy of the data and applying logs or other necessary processes to that copy to bring it to a known point of consistency. Once all recoveries are completed, the data is validated to ensure that it is correct.
- **Disaster restart:** This is the process of restarting business operations with mirrored consistent copies of data and applications.
- **Recovery-Point Objective (RPO):** This is the point in time to which systems and data must be recovered after an outage. It defines the amount of data loss that a business can endure. A large RPO signifies high tolerance to information loss in a business. Based on the RPO, organizations plan for the minimum frequency with which a backup or replica must be made. For

example, if the RPO is six hours, backups or replicas must be made at least once in 6 hours. Figure 11-2 shows various RPOs and their corresponding ideal recovery strategies. An organization can plan for an appropriate BC technology solution on the basis of the RPO it sets. For example:

- **RPO of 24 hours:** This ensures that backups are created on an offsite tape drive every midnight. The corresponding recovery strategy is to restore data from the set of last backup tapes.
- **RPO of 1 hour:** This ships database logs to the remote site every hour. The corresponding recovery strategy is to recover the database at the point of the last log shipment.
- **RPO of zero:** This mirrors mission-critical data synchronously to a remote site.

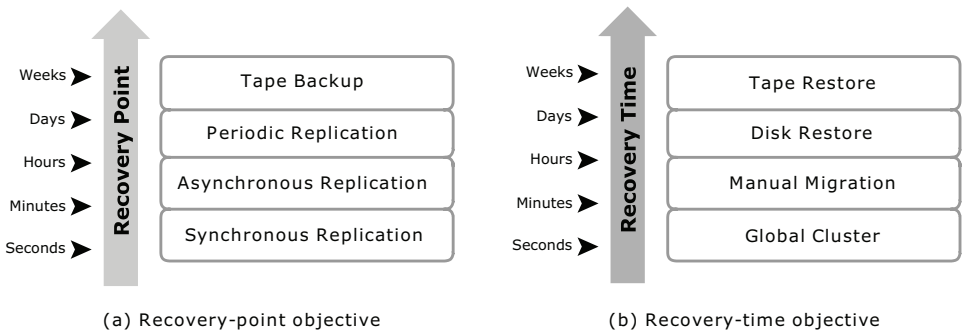

Figure 11-2: Strategies to meet RPO and RTO targets

- **Recovery-Time Objective (RTO):** The time within which systems, applications, or functions must be recovered after an outage. It defines the amount of downtime that a business can endure and survive. Businesses can optimize disaster recovery plans after defining the RTO for a given data center or network. For example, if the RTO is two hours, then use a disk backup because it enables a faster restore than a tape backup. However, for an RTO of one week, tape backup will likely meet requirements. Some examples of RTOs and the recovery strategies to ensure data availability are listed below (refer to Figure 11-2):
  - **RTO of 72 hours:** Restore from backup tapes at a cold site.
  - **RTO of 12 hours:** Restore from tapes at a hot site.
  - **RTO of 4 hours:** Use a data vault to a hot site.

- **RTO of 1 hour:** Cluster production servers with controller-based disk mirroring.
- **RTO of a few seconds:** Cluster production servers with bidirectional mirroring, enabling the applications to run at both sites simultaneously.

- **Data vault:** A repository at a remote site where data can be periodically or continuously copied (either to tape drives or disks), so that there is always a copy at another site.
- **Hot site:** A site where an enterprise's operations can be moved in the event of disaster. It is a site with the required hardware, operating system, application, and network support to perform business operations, where the equipment is available and running at all times.
- **Cold site:** A site where an enterprise's operations can be moved in the event of disaster, with minimum IT infrastructure and environmental facilities in place, but not activated.
- **Cluster:** A group of servers and other necessary resources, coupled to operate as a single system. Clusters can ensure high availability and load balancing. Typically, in failover clusters, one server runs an application and updates the data, and another server is kept redundant to take over completely, as required. In more sophisticated clusters, multiple servers may access data, and typically one server performs coordination.

11.3 BC Planning Lifecycle

BC planning must follow a disciplined approach like any other planning process. Organizations today dedicate specialized resources to develop and maintain BC plans. From the conceptualization to the realization of the BC plan, a lifecycle of activities can be defined for the BC process. The BC planning lifecycle includes five stages (see Figure 11-3):

1. Establishing objectives
2. Analyzing
3. Designing and developing
4. Implementing
5. Training, testing, assessing, and maintaining

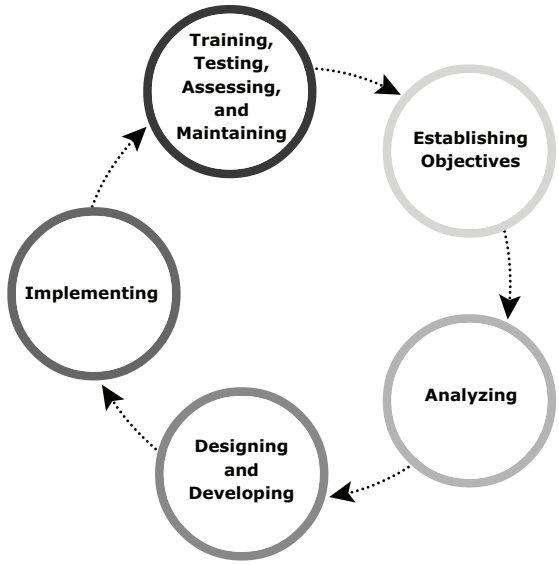

Figure 11-3: BC planning lifecycle

Several activities are performed at each stage of the BC planning lifecycle, including the following key activities:

1. Establishing objectives
   - Determine BC requirements.
   - Estimate the scope and budget to achieve requirements.
   - Select a BC team by considering subject matter experts from all areas of the business, whether internal or external.
   - Create BC policies.
2. Analyzing
   - Collect information on data profiles, business processes, infrastructure support, dependencies, and frequency of using business infrastructure.
   - Identify critical business needs and assign recovery priorities.
   - Create a risk analysis for critical areas and mitigation strategies.
   - Conduct a Business Impact Analysis (BIA).
   - Create a cost and benefit analysis based on the consequences of data unavailability.
   - Evaluate options.

3. Designing and developing
   - Define the team structure and assign individual roles and responsibilities. For example, different teams are formed for activities such as emergency response, damage assessment, and infrastructure and application recovery.
   - Design data protection strategies and develop infrastructure.
   - Develop contingency scenarios.
   - Develop emergency response procedures.
   - Detail recovery and restart procedures.
4. Implementing
   - Implement risk management and mitigation procedures that include backup, replication, and management of resources.
   - Prepare the disaster recovery sites that can be utilized if a disaster affects the primary data center.
   - Implement redundancy for every resource in a data center to avoid single points of failure.
5. Training, testing, assessing, and maintaining
   - Train the employees who are responsible for backup and replication of business-critical data on a regular basis or whenever there is a modification in the BC plan.
   - Train employees on emergency response procedures when disasters are declared.
   - Train the recovery team on recovery procedures based on contingency scenarios.
   - Perform damage assessment processes and review recovery plans.
   - Test the BC plan regularly to evaluate its performance and identify its limitations.
   - Assess the performance reports and identify limitations.
   - Update the BC plans and recovery/restart procedures to reflect regular changes within the data center.

11.4 Failure Analysis

Failure analysis involves analyzing the data center to identify systems that are susceptible to a single point of failure and implementing fault-tolerance mechanisms such as redundancy.

11.4.1 Single Point of Failure

A *single point of failure* refers to the failure of a component that can terminate the availability of the entire system or IT service. Figure 11-4 illustrates the possibility of a single point of failure in a system with various components: server, network, switch, and storage array. The figure depicts a system setup in which an application running on the server provides an interface to the client and performs I/O operations. The client is connected to the server through an IP network, the server is connected to the storage array through a FC connection, an HBA installed at the server sends or receives data to and from a storage array, and an FC switch connects the HBA to the storage port.

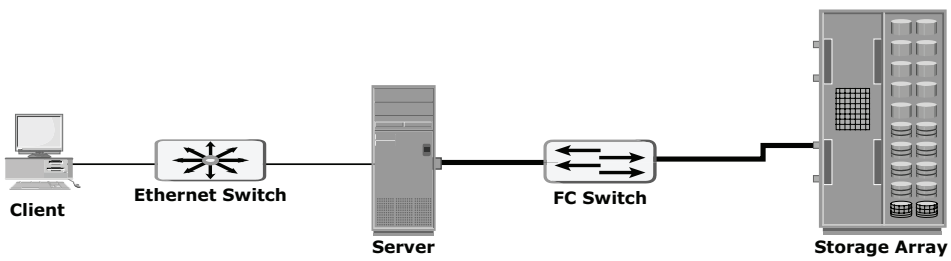

Figure 11-4: Single point of failure

In a setup where each component must function as required to ensure data availability, the failure of a single component causes the failure of the entire data center or an application, resulting in disruption of business operations. In this example, several single points of failure can be identified. The single HBA on the server, the server itself, the IP network, the FC switch, the storage array ports, or even the storage array could become potential single points of failure. To avoid single points of failure, it is essential to implement a fault-tolerant mechanism.

11.4.2 Fault Tolerance

To mitigate a single point of failure, systems are designed with redundancy, such that the system will fail only if all the components in the redundancy group fail. This ensures that the failure of a single component does not affect

data availability. Figure 11-5 illustrates the fault-tolerant implementation of the system just described (and shown in Figure 11-4).

Data centers follow stringent guidelines to implement fault tolerance. Careful analysis is performed to eliminate every single point of failure. In the example shown in Figure 11-5, all enhancements in the infrastructure to mitigate single points of failures are emphasized:

- Configuration of multiple HBAs to mitigate single HBA failure.
- Configuration of multiple fabrics to account for a switch failure.
- Configuration of multiple storage array ports to enhance the storage array's availability.
- RAID configuration to ensure continuous operation in the event of disk failure.
- Implementing a storage array at a remote site to mitigate local site failure.
- Implementing server (host) clustering, a fault-tolerance mechanism whereby two or more servers in a cluster access the same set of volumes. Clustered servers exchange *heartbeats* to inform each other about their health. If one of the servers fails, the other server takes up the complete workload.

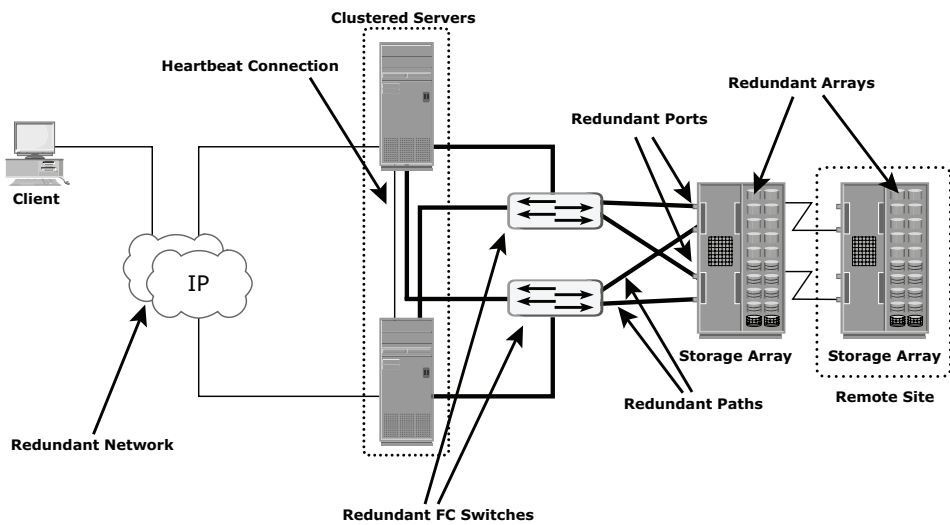

Figure 11-5: Implementation of fault tolerance

11.4.3 Multipathing Software

Configuration of multiple paths increases the data availability through path failover. If servers are configured with one I/O path to the data there will be no access to the data if that path fails. Redundant paths eliminate the path to become single points of failure. Multiple paths to data also improve I/O performance through load sharing and maximize server, storage, and data path utilization.

In practice, merely configuring multiple paths does not serve the purpose. Even with multiple paths, if one path fails, I/O will not reroute unless the system recognizes that it has an alternate path. Multipathing software provides the functionality to recognize and utilize alternate I/O path to data. Multipathing software also manages the load balancing by distributing I/Os to all available, active paths.

11.5 Business Impact Analysis

A *business impact analysis (BIA)* identifies and evaluates financial, operational, and service impacts of a disruption to essential business processes. Selected functional areas are evaluated to determine resilience of the infrastructure to support information availability. The BIA process leads to a report detailing the incidents and their impact over business functions. The impact may be specified in terms of money or in terms of time. Based on the potential impacts associated with downtime, businesses can prioritize and implement countermeasures to mitigate the likelihood of such disruptions. These are detailed in the BC plan. A BIA includes the following set of tasks:

- Identify the key business processes critical to its operation.
- Determine the attributes of the business process in terms of applications, databases, and hardware and software requirements.
- Estimate the costs of failure for each business process.
- Calculate the maximum tolerable outage and define RTO and RPO for each business process.
- Establish the minimum resources required for the operation of business processes.
- Determine recovery strategies and the cost for implementing them.

- Optimize the backup and business recovery strategy based on business priorities.
- Analyze the current state of BC readiness and optimize future BC planning.

11.6 BC Technology Solutions

After analyzing the business impact of an outage, designing appropriate solutions to recover from a failure is the next important activity. One or more copies of the original data are maintained using any of the following strategies, so that data can be recovered and business operations can be restarted using an alternate copy:

- **Backup and recovery:** Backup to tape is the predominant method of ensuring data availability. These days, low-cost, high-capacity disks are used for backup, which considerably speeds up the backup and recovery process. The frequency of backup is determined based on RPO, RTO, and the frequency of data changes.
- **Storage array-based replication (local):** Data can be replicated to a separate location within the same storage array. The replica is used independently for BC operations. Replicas can also be used for restoring operations if data corruption occurs.
- **Storage array-based replication (remote):** Data in a storage array can be replicated to another storage array located at a remote site. If the storage array is lost due to a disaster, BC operations start from the remote storage array.
- **Host-based replication:** The application software or the LVM ensures that a copy of the data managed by them is maintained either locally or at a remote site for recovery purposes.

11.7 Concept in Practice: EMC PowerPath

PowerPath is a host-based multipathing software that provides path failover and load balancing functionality. PowerPath operates between operating systems and device drivers and supports SCSI, iSCSI, and Fibre Channel environment. It prioritizes I/O bandwidth utilization by using sophisticated load balancing algorithms to ensure optimal application performance. Refer to `http://education.EMC.com/ismbook` for the latest information.

11.7.1 PowerPath Features

PowerPath provides the following features:

- **Online path configuration and management:** PowerPath provides the flexibility to define some paths to a device as "active" and some as "standby." The standby paths are used when all active paths to a logical device have failed. Paths can be dynamically added and removed by setting them in standby or active mode.
- **Dynamic load balancing across multiple paths:** PowerPath distributes the I/O requests across all available paths to the logical device. This reduces bottlenecks and improves application performance.
- **Automatic path failover:** In the event of a path failure, PowerPath fails over seamlessly to an alternative path without disrupting application operations. PowerPath redistributes I/O to the best available path to achieve optimal host performance.
- **Proactive path testing:** PowerPath uses the autoprobe and autorestore functions to proactively test the dead and restored paths, respectively. The PowerPath *autoprobe* function periodically probes inactive paths to identify failed paths before sending the application I/O. This process enables PowerPath to proactively close paths before an application experiences a timeout when sending I/O over failed paths. The PowerPath *autorestore* function runs every five minutes and tests every failed or closed path to determine whether it has been repaired.
- **Cluster support:** The deployment of PowerPath in a server cluster eliminates application downtime due to a path failure. PowerPath detects the path failure and uses an alternate path so the cluster software does not have to reconfigure the cluster to keep the applications running.
- **Interoperability:** PowerPath is supported on many operating systems, storage arrays, and storage interconnected devices, including iSCSI devices.

11.7.2 Dynamic Load Balancing

For every I/O, the PowerPath filter driver selects the path based on the load-balancing policy and failover setting for the logical device. The driver identifies all available paths that read and write to a device and builds a routing table called a volume path set for the devices. PowerPath follows any one of the following user specified load-balancing policies:

- **Round-Robin policy:** I/O requests are assigned to each available path in rotation.
- **Least I/Os policy:** I/O requests are routed to the path with the fewest queued I/O requests, regardless of the total number of I/O blocks.

- **Least Blocks policy:** I/O requests are routed to the path with the fewest queued I/O blocks, regardless of the number of requests involved.
- **Priority-Based policy:** I/O requests are balanced across multiple paths based on the composition of reads, writes, user-assigned devices, or application priorities.

I/O Operation without PowerPath

Figure 11-6 illustrates I/O operations in a storage system environment in the absence of PowerPath. The applications running on a host have four paths to the storage array. However, the applications can use only one of the paths because the LVM that is native to the host operating system allows only one path for application I/O operations.

This example illustrates how I/O throughput is unbalanced without PowerPath. Two applications are generating high I/O traffic, which overloads both paths, but the other two paths are less loaded. In this scenario, some paths may be idle or unused while other paths have multiple I/O operations queued. As a result, the applications cannot achieve optimal performance.

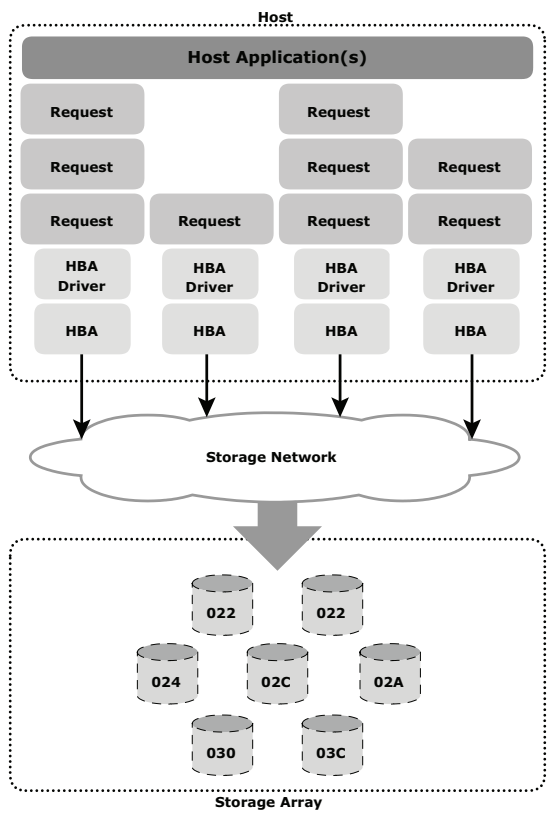

Figure 11-6: I/O without PowerPath

I/O Operation with PowerPath

Figure 11-7 shows I/O operations in a storage system environment that has PowerPath. PowerPath ensures that I/O requests are balanced across the four paths to storage, based on the load-balancing algorithm chosen. As a result, the applications can effectively utilize their resources, thereby improving their performance.

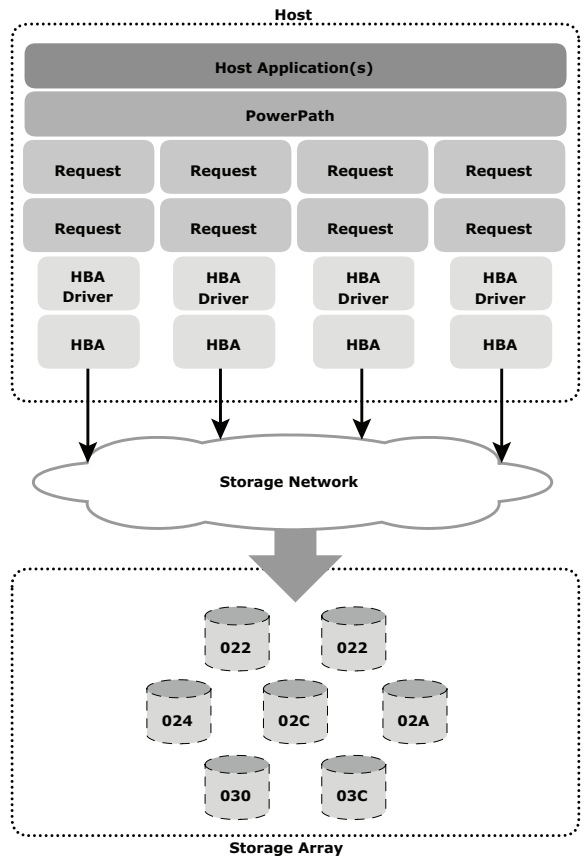

Figure 11-7: I/O with PowerPath

11.7.3 Automatic Path Failover

The next two examples demonstrate how PowerPath performs path failover operations in the event of a path failure for active-active and active-passive array configurations.

Path Failure without PowerPath

Figure 11-8 shows a scenario in which applications use only one of the four paths defined by the operating system. Without PowerPath, the loss of paths (the path failure is marked by a cross "X") due to single points of failure, such as the loss of an HBA, storage array front-end connectivity, switch port, or a failed cable, can result in an outage for one or more applications.

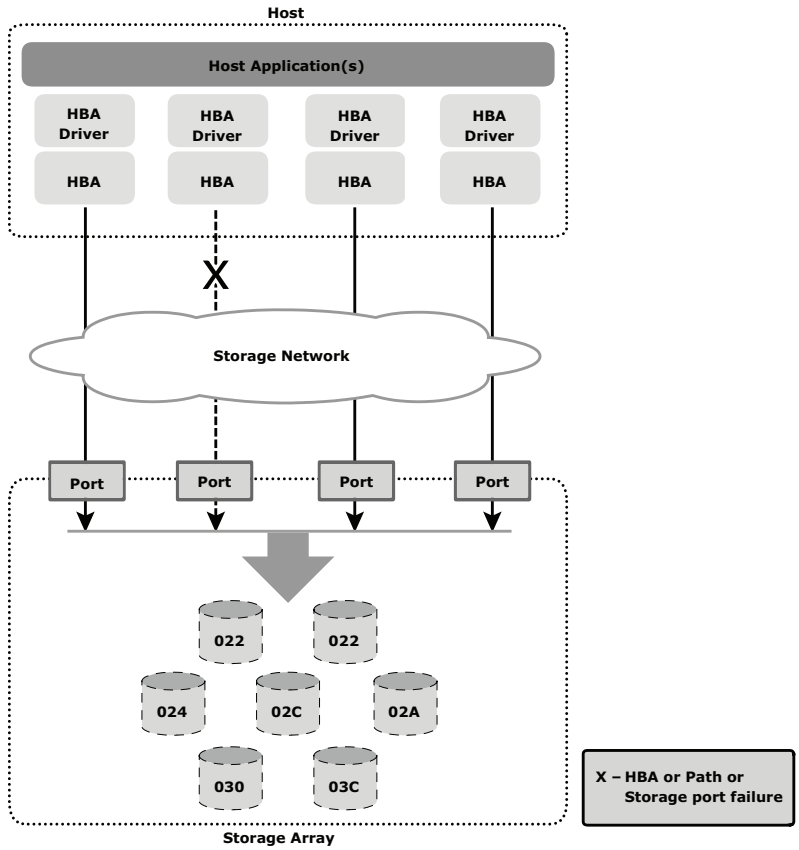

Figure 11-8: Path failure without PowerPath

Path Failover with PowerPath: Active-Active Array

Figure 11-9 shows a storage system environment in which an application uses PowerPath with an active-active array configuration to perform I/O operations. PowerPath redirects the application I/Os through an alternate active path.

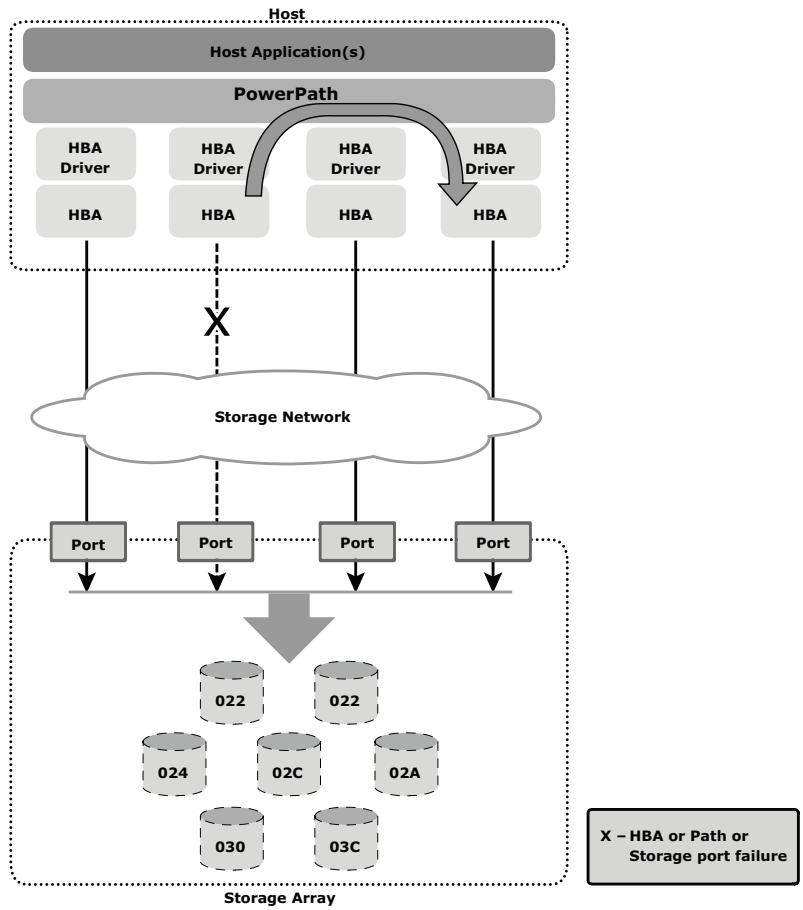

Figure 11-9: Path failover with PowerPath for an active-active array

In the event of a path failure, PowerPath performs the following operations:

1. If an HBA, cable, or storage front-end port fails, the device driver returns a timeout to PowerPath.
2. PowerPath responds by setting the path offline and redirecting the I/O through an alternate path.
3. Subsequent I/Os use alternate active path(s).

Path Failover with PowerPath: Active-Passive Array

Figure 11-10 shows a scenario in which a logical device is assigned to a storage processor B (SP B).

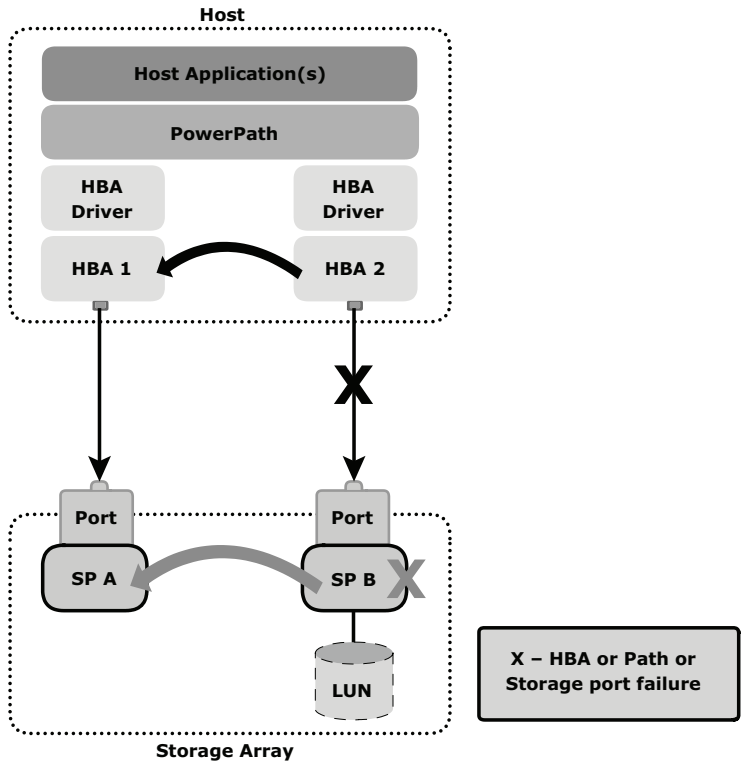

Figure 11-10: Path failover with PowerPath for an active-passive array

Path failure can occur due to a failure of the link, HBA, or storage processor (SP). In the event of a path failure, PowerPath with an active-passive configuration performs the path failover operation in the following way:

- If an active I/O path to SP B through HBA 2 fails, PowerPath uses a passive path to SP B through HBA 1.
- If HBA 2 fails, the application uses HBA 1 to access the logical device.
- If SP B fails, PowerPath stops all I/O to SP B and *trespasses* the device over to SP A. All I/O will be sent down the paths to SP A, this process is referred as *LUN trespassing*. When SP B is brought back online, PowerPath recognizes that it is available and resumes sending I/O down to SP B.

Summary

Technology innovations have led to a rich set of options in terms of storage devices and solutions to meet the needs of businesses for high availability and business continuity. The goal of any business continuity (BC) plan is to identify

and implement the most appropriate risk management and mitigation procedures to protect against possible failures. The process of analyzing the hardware and software configuration to identify any single points of failure and their impact on business operations is critical. A business impact analysis (BIA) helps a company develop an appropriate BC plan to ensure that the storage infrastructure and services are designed to meet business requirements. BC provides the framework for organizations to implement effective and cost-efficient disaster recovery and restart procedures. In a constantly changing business environment, BC can become a demanding endeavor.

The next three chapters discuss specific BC technology solutions, backup and recovery, local replication, and remote replication.

EXERCISES

1. A network router has a failure rate of 0.02 percent per 1,000 hours. What is the MTBF of that component?
2. The IT department of a bank promises customer access to the currency conversion rate table between 9:00 AM and 4:00 PM from Monday to Friday. It updates the table every day at 8:00 AM with a feed from the mainframe system. The update process takes 35 minutes to complete. On Thursday, due to a database corruption, the rate table could not be updated. At 9:05 AM, it was established that the table had errors. A rerun of the update was done and the table was recreated at 9:45 AM Verification was run for 15 minutes and the rate table became available to the bank branches. What was the availability of the rate table for the week in which this incident took place, assuming there were no other issues?
3. "Availability is expressed in terms of 9s." Explain the relevance of the use of 9s for availability, using examples.
4. Provide examples of planned and unplanned downtime in the context of data center operations.
5. How does clustering help to minimize RTO?
6. How is the choice of a recovery site strategy (cold and hot) determined in relation to RTO and RPO?
7. Assume the storage configuration design shown in the following figure:

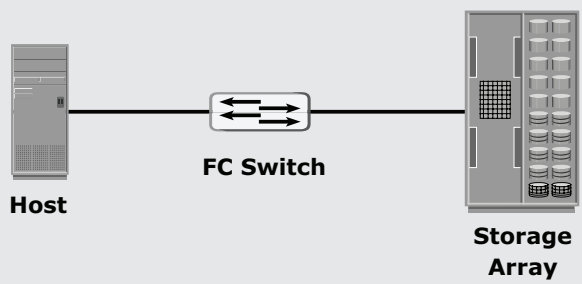

Perform the single point of failure analysis for this configuration and provide an alternate configuration that eliminates all single points of failure.