DISASTER RECOVERY

Topic Objective:

At the end of this topic student would be able to:

- Understand Classification of Disasters
- Evaluate Security Holes
- Analyze General steps to follow while creating BCP/DRP
- Explain Control majors in recovery plan
- Elaborate Storage area network
- Highlight Network types
- Know about Storage sharing
- Learn about SAN infrastructure
- Highlight Network Edition Disaster Recovery
- Point out Disaster Recovery Planning
- Examine Power considerations
- Identify Change control and documentation
- Define Redundancy
- Describe Capacity
- Explain Security
- Evaluate Plan exercising

Definition/Overview:

Disaster recovery is the process, policies and procedures related to preparing for recovery or continuation of technology infrastructure critical to an organization after a natural or human-induced disaster. Disaster recovery planning is a subset of a larger process known as business continuity planning and should include planning for resumption of applications, data, hardware, communications (such as networking) and other IT infrastructure. A business continuity plan (BCP) includes planning for non-IT related aspects such as key personnel, facilities, crisis communication and reputation protection, and should refer to the disaster recovery plan (DRP) for IT related infrastructure recovery / continuity. This article focuses on disaster recovery planning as related to IT infrastructure.

Key Points:

1. Classification of Disasters

Disaster can be classified in two broad categories. Viz, 1) Natural disasters- Preventing natural disaster is very difficult, but precaution to avoid losses can be taken. These disasters include flood, fire, earthquake, hurricane, smog and chemical gases, etc 2) Man made disasters- These disasters are major reasons for failure. Human error and intervention may be intentional or unintentional which can cause massive failure such as failure in communication and utility. These disasters include walkout, sabotage, burglary, virus, intrusion, etc.

2. Security Holes

Security holes are the vulnerabilities in computing hardware or software. It provides indirect invitation to malicious brains to work on it and exploit it. It is achieved through flaws in network software which allows unintended control within the network. Components of network such as PCs and router hold these holes through their operating systems. Technical details of any systems should not be made public abundantly unless required. Once such holes are discovered, information about it should be immediately passed to security professional responsible for it. On the other hand such information is also passed quickly to hacker who might want to intercept into the network. Security professional should always work to heal such holes to eliminate possible attack.

3. General steps to follow while creating BCP/DRP

- Identify the scope and boundaries of business continuity plan.
- Fist step enables us to define scope of BCP. It provides and idea for limitations and boundaries of plan. It also includes audit and risk analysis reports for institutions assets.
- Create the business impact assessment.
- Business impact analysis is study and assessment of financial losses to institution resulting from destructive event as unavailability of important business services.
- Sell the concept of BCP to upper management and obtain organizational and financial commitment.

- Convincing senior management to approve BCP/DRP is key task. It is very important for security professional to get approval for plan from upper management to bring it to effect.
- Each department will need to understand its role in plan and support to maintain it.
- In case of disaster, each department has to be prepared for the action. To recover and to protect the critical systems each department has to understand the plan follows it accordingly. It is also important to maintain and help in creation of plan for each individual department.
- The BCP project team must implement the plan.
- After approval from upper management plan should be maintained and implemented. Implementation team should follow the guidelines procedures in plan.
- NIST tool set can be used for doing BCP.
- National Institute of standards and Technologies has published good tools which can help in creating BCP.

4. Control majors in recovery plan

Control majors are steps or mechanism that can reduce or eliminate computer security threats. Different types of majors can be included in BCP/DRP

- **Types of majors**
  - **Preventive majors**: These controls can avoid or prevent an event from occurring.
  - **Detective maj**ors: These controls make us capable to detect or discover unwanted event.
  - **Corrective majors**: These controls help to correct or recover the system after disaster or event.

These controls should be always documented and tested regularly.

- **Strategies**

Prior to selecting a disaster recovery strategy, a disaster recovery planner should refer to their organization's business continuity plan which should indicate the key metrics of recovery point objective (RPO) and recovery time objective (RTO) for various business processes (such as the process to run payroll, generate an order, etc). The metrics specified for the business processes must then be mapped to the underlying IT systems and infrastructure that support those processes. Once the RTO and RPO metrics have been mapped to IT infrastructure, the DR planner can determine the most suitable recovery strategy for each system. An important note here however is that the business ultimately sets the IT budget and therefore the RTO and RPO metrics need to fit with the available budget. While most business unit heads would like zero data loss and zero time loss, the cost associated with that level of protection may make the desired high availability solutions impractical. The following is a list of the most common strategies for data protection.

- Backups made to tape and sent off-site at regular intervals (preferably daily)
- Backups made to disk on-site and automatically copied to off-site disk, or made directly to off-site disk
- Replication of data to an off-site location, which overcomes the need to restore the data (only the systems then need to be restored or synced). This generally makes use of storage area network (SAN) technology
- High availability systems which keep both the data and system replicated off-site, enabling continuous access to systems and data

In many cases, an organization may elect to use an outsourced disaster recovery provider to provide a stand-by site and systems rather than using their own remote facilities. In addition to preparing for the need to recover systems, organizations must also implement precautionary measures with an objective of preventing a disaster situation in the first place. These may include some of the following:

- Local mirrors of systems and/or data and use of disk protection technology such as RAID
- Surge protectors to minimize the effect of power surges on delicate electronic equipment
- Uninterruptible power supply (UPS) and/or backup generator to keep systems going in the event of a power failure
- Fire preventions alarms, fire extinguishers
- Anti-virus software and other security measures

5. Storage area network

A storage area network (SAN) is an architecture to attach remote computer storage devices (such as disk arrays, tape libraries, and optical jukeboxes) to servers in such a way that the devices appear as locally attached to the operating system. Although the cost and complexity of SANs are dropping, they are still uncommon outside larger enterprises.

Network attached storage (NAS), in contrast to SAN, uses file-based protocols such as NFS or SMB/CIFS where it is clear that the storage is remote, and computers request a portion of an abstract file rather than a disk block.

6. Network types

Most storage networks use the SCSI protocol for communication between servers and disk drive devices. They do not use SCSI low-level physical interface (e.g. cables), however, as its bus topology is unsuitable for networking. A mapping layer to other low-level protocols is used to form a network:

- ATA over Ethernet (AoE), mapping of ATA over Ethernet,
- Fibre Channel Protocol (FCP), the most prominent one, is mapping of SCSI over Fibre Channel (FC),
- Fibre Channel over Ethernet (FCoE),
- mapping of FICON over FC, used by mainframe computers,
- HyperSCSI, mapping of SCSI over Ethernet,
- iFCP or SANoIP mapping of FCP over IP.
- iSCSI Extensions for RDMA (iSER), mapping of iSCSI over InfiniBand (IB),
- iSCSI, mapping of SCSI over TCP/IP

7. Storage sharing

Historically, data centers first created "islands" of SCSI disk arrays. Each island was dedicated to an application, and visible as a number of "virtual hard drives" (i.e. LUNs). Essentially, a SAN connects storage islands together using a high-speed network, thus allowing all applications to access all disks. Operating systems still view a SAN as a collection of LUNs, and usually maintain their own file systems on them. These local file systems, which cannot be shared among multiple operating systems/hosts, are the most reliable and most widely used. If two independent local file systems resided on a shared LUN, they would be unaware of this fact, would have no means of cache synchronization and eventually would corrupt each other. Thus, sharing data between computers through a SAN requires advanced solutions, such as SAN file systems or clustered computing. Despite such issues, SANs help to increase storage capacity utilization, since multiple servers share the storage space on the disk arrays. The common application of a SAN is for the use of transactionally accessed data that require high-speed block-level access to the hard drives such as email servers, databases, and high usage file servers. In contrast, NAS allows many computers to access the same file system over the network and synchronizes their accesses. Lately, the introduction of NAS heads allowed easy conversion of SAN storage to NAS.

- **Benefits**

Sharing storage usually simplifies storage administration and adds flexibility since cables and storage devices do not have to be physically moved to shift storage from one server to another. Other benefits include the ability to allow servers to boot from the SAN itself. This allows for a quick and easy replacement of faulty servers since the SAN can be reconfigured so that a replacement server can use the LUN of the faulty server. This process can take as little as half an hour and is a relatively new idea being pioneered in newer data centers. There are a number of emerging products designed to facilitate and speed this up still further. Brocade, for example, offers an Application Resource Manager product which automatically provisions servers to boot off a SAN, with typical-case load times measured in minutes. While this area of technology is still new many view it as being the future of the enterprise datacenter. SANs also tend to enable more effective disaster recovery processes. A SAN could span a distant location containing a secondary storage array. This enables storage replication either implemented by disk array controllers, by server software, or by specialized SAN devices. Since IP WANs are often the least costly method of long-distance transport, the Fibre Channel over IP (FCIP) and iSCSI protocols have been developed to allow SAN extension over IP networks. The traditional physical SCSI layer could only support a few meters of distance - not nearly enough to ensure business continuance in a disaster. Demand for this SAN application has increased dramatically after the September 11th attacks in the United States, and increased regulatory requirements associated with Sarbanes-Oxley and similar legislation. The economic consolidation of disk arrays has accelerated the advancement of several features including I/O caching, snapshotting, and volume cloning (Business Continuance Volumes or BCVs).

8. SAN infrastructure

SAN-switch Qlogic with optical Fibre Channel connectors installed. SANs often utilize a Fibre Channel fabric topology - an infrastructure specially designed to handle storage communications. It provides faster and more reliable access than higher-level protocols used in NAS. A fabric is similar in concept to a network segment in a local area network. A typical Fibre Channel SAN fabric is made up of a number of Fibre Channel switches. Today, all major SAN equipment vendors also offer some form of Fibre Channel routing solution, and these bring substantial scalability benefits to the SAN architecture by allowing data to cross between different fabrics without merging them. These offerings use proprietary protocol elements, and the top-level architectures being promoted are radically different. They often enable mapping Fibre Channel traffic over IP or over SONET/SDH.

- **Compatibility**

  One of the early problems with Fibre Channel SANs was that the switches and other hardware from different manufacturers were not entirely compatible. Although the basic storage protocols FCP were always quite standard, some of the higher-level functions did not interoperate well. Similarly, many host operating systems would react badly to other operating systems sharing the same fabric. Many solutions were pushed to the market before standards were finalized and vendors innovated around the standards. Note: fibre channel and SAN are not synonymous

- **SANs at home**

  SANs are primarily used in large scale, high performance enterprise storage operations. It would be unusual to find a single disk drive connected directly to a SAN. Instead, SANs are normally networks of large disk arrays. SAN equipment is relatively expensive, therefore, Fibre Channel host bus adapters are rare in desktop computers. The iSCSI SAN technology is expected to eventually produce cheap SANs, but it is unlikely that this technology will be used outside the enterprise data center environment. Desktop clients are expected to continue using NAS protocols such as SMB and NFS. The exception to this may be remote storage replication.

- **SANs in the media and entertainment**

  Video editing workgroups require very high data transfer rates. Outside of the enterprise market, this is one area that greatly benefits from SANs. Per-node bandwidth usage control, sometimes referred to as quality-of-service (QoS), is especially important in video workgroups as it ensures fair and prioritized bandwidth usage across the network if there is insufficient open bandwidth available. Avid Unity, Apple's Xsan and Tiger Technology MetaSAN are specifically designed for video networks and offer this functionality.

- **Storage virtualization and SANs**

  Storage virtualization refers to the process of completely abstracting logical storage from physical storage. The physical storage resources are aggregated into storage pools, from which the logical storage is created. It presents to the user a logical space for data storage and transparently handles the process of mapping it to the actual physical location. This is currently implemented inside each modern disk array, using vendor's proprietary solution. However, the goal is to virtualize multiple disk arrays, made by different vendors, scattered over the network, into a single monolithic storage device, which can be managed uniformly.

9. Network Edition Disaster Recovery

The new server hardware must meet the requirements described in the Installation Prerequisites section of the ZCS Single Server Installation Guide. Install the new operating system, making any necessary OS configuration modifications as described in the installation guide. Before you begin, make sure that the new server is correctly configured with the IP address and hostname and that ZCS is installed and configured with the same domain, hostname, passwords, etc. as the previous server. See the Single-Server Installation Guide for more information about preparing the server. Before you begin to install ZCS, note the information you need from the old server including: admin account name and password, spam training and non-spam training user account names, exact domain name, and the global document account name.

10. Disaster Recovery Planning

All across the globe businesses small, medium and large are all becoming more and more reliant on their IP networks for survival. This coupled with the growing trend for the convergence of the voice, data, and video over a single IP network make an organizations network infrastructure one of the most critical elements in its overall operation. No longer can organizations afford not to include a thorough and comprehensive plan for their continued availability as part of their business continuity and disaster recovery planning efforts. Company networks now must provide voice, video, and data services that are increasingly integrated with applications. So if the company network fails all forms of communication with customers, suppliers and employees can also fail dramatically. Worse yet access to critical information can be lost or potentially compromised. Yet with these calamitous results real possibilities for companies, I find it surprising that many organizations continue to be inadequately prepared to deal with adverse events relating to their business and network operations.

In 2007 a computer crash in the Customs office of Los Angeles InternationalAirport (LAX) caused hours of delays for more than 17,000 airline passengers. US Customs officials found that a malfunctioning network card caused Customs to lose access to their national systems and databases and their local area network. This connectivity failure created a domino effect leading to a total system failure that caused massive wait times at the airport, stranding some passengers. It took technicians over ten hours to diagnose the problem, halting screening operations until it could be resolved. In another case a major Medical Centerin Bostonrelied heavily upon its networked advanced clinical computing system. With this system, clinicians throughout the medical center and other affiliated hospitals could gain access to laboratory results, radiographs, and electrocardiograms electronically, using a secure Intranet. Patients also had secure access to their test results over the Internet. The outage lasted almost a week during which time the hospital staff had to scramble to utilize hand-carried patient records, laboratory-test results, and countless other documents around the hospital in order to maintain clinical operations. Clearly in both cases disaster recovery plans which address the network and its infrastructure were needed.

It should be clear to all that are responsible for IT within corporations and government agencies that the network and the network infrastructure (comprised of DNS, DHCP, and etc.) are getting more complex and thus harder to manage, yet are a most important part of their organizations overall operational success. Networks provide voice, video, and data services that are increasingly integrated with business critical applications. Applications such as e-mail, CRM, and ERM rely on the network for proper operation. As such the network should be considered of great importance in any business continuity and/or disaster recovery plan. It should be understood that any disaster recovery planning effort should address all of the elements of an organizations network. Most corporate and government networks are comprised of three main elements LAN, WAN, and network infrastructure services. The LAN provides for interconnectivity around a single organizational location or locale, the WAN provides interconnectivity between these locations (interconnecting geographically specific sites) other business partners and access to public networks such as the public switched telephone network in the case of voice traffic and the Internet for data traffic. The network infrastructure services element provides the services that allow control of the network and flow of data such as DNS, DHCP, WINS, FTP, and contain access to the network in the case of Active Directory, RADIUS, and TACACS.

11. Power considerations

According to the many producers of business continuity and disaster recovery surveys and statistics the single largest reason for network and systems failures can be directly attributed to power failures. So planning for power failures is essential in any DR plan. This means that all critical network components at either the primary data center, call center or failover site must be connected to a power source that has a very high availability percentage. In the case of a data center the percentage of availability should be in the area of 99.999 percent. If the LAN provides critical services, as would be the case in a hospital or a bank, then each component of the distribution and access portions of the LAN such as each floor closet should be equipped with uninterruptible power supplies (UPS) which are connected to emergency power sources to maintain internal communication. The WAN routers, switches, firewalls and the like need the same form of protection to provide continuous communication and interconnection to the external sites and other public networks.

Many large data centers or critical operations, such as call centers, rely upon multiple electric power companies to provide utility power to their locations. The power is brought in to the critical site from different geographical locations. In that way if power is interrupted by something like a car-pole accident which severs the electric lines at a particular location the other utility can continue to provide uninterrupted power. Many critical sites operate with emergency power generators, where possible, instead of alternate utilities as described above. These generators, together with UPS equipment can provide a continuous stream of electrical power for days if necessary while utility power is being restored. However, regular maintenance of these generators and their fuel source is critical to ensure their availability when needed.

12. Change control and documentation

We all know that all organization change continuously that is a good thing because without growth many organizations cease to exist. However, change can be challenging for those who need to protect the network and its infrastructure. Every network should be properly inventoried, have network diagrams which show the exact state of the network at a given point in time. Each critical element of the LAN, WAN and Infrastructure services must be known and identified, and properly classified within the business impact analyses which are done periodically. As changes propagate, the network documentation should be continually updated to show the exact configuration of the network topology. Even more importantly, if an alternate recovery site exists it should be subject to the same changes, patches, and configurations as the primary site. People are asked to perform quickly and under extremely difficult conditions when a disaster occurs. The difference between success and failure of a disaster recovery plan may be reliant upon the accuracy of the documentation or the currency of the changes to the DR sites network. Many disaster recovery failures occur because a change to a network element like a switch was never completed on the disaster recovery site switch and a fail-over connection could not be made. Processes especially need to be in place to insure that patches, address changes, access control lists, and new network equipment are incorporated into the disaster recovery network as changes are made.

13. Redundancy

Critical network components are identified, impact analyses have been completed, and the necessary recovery point objective has been established. The level of resiliency of the network should be known. Based on that knowledge the planner can determine the levels of redundancy needed in the networks (primary and backup network). There should be considerations for redundancy of components of network elements (e.g., switches, routers, and etc.). There should also be consideration given to redundant components such as power supplies, CPUs, and circuit cards for those network switches and routers. There should also be considerations given to the redundancy and diversity of WAN circuits. Redundancy can be achieved by providing multiple circuits and multiple types of circuits between critical sites and applications. For example if the WAN network utilized MPLS or ATM it might be prudent to provide different circuits such as frame-relay so that if a carriers entire service goes down (which has happened in the past) the organization can have a backup strategy. Satellite or microwave links can also be a strategies considered between some critical sites. Diversity of circuits can be accomplished either by link diversity insuring that if two links are used they travel different routes to your locations. That way if a link was compromised the alternate link, traveling a different path would be potentially unaffected. There is also carrier diversity. This protects against a carriers service failure by utilizing a second carrier to provide a similar service. Many times multiple carriers are used to provide Internet access diversity and redundancy to a company, especially if it relies heavily on Internet connectivity for ecommerce.

14. Capacity

When developing a fail-over scenario the planner must take into consideration several capacity factors. One is the peak capacity of the site where the traffic will be rerouted to. The second is the peak capacity coming from the site which has failed. The size of the WAN circuits should be such so as to allow for both peak capacities plus an additional 25 to 40 percent. The reason for the additional 25 to 40 percent is to accommodate new peak traffic volumes from added VoIP and/or data traffic caused by customers, suppliers, and employees needing to find out about how the problem affects them. The planner should have a clear communication plan for notifying customers, suppliers, and employees in case of an activation of the business continuity plan. Also, the business continuity plan should prepare for additional personnel to be made available to handle additional call volumes at the fail over site otherwise degraded service will be the norm. On a more simplistic note if the disaster recovery site is designed to employ smaller class routers or switches or reduced capacity circuits than the primary site then the disaster recovery plan is already set up for failure. As soon as the traffic is re-routed to the alternate disaster recovery site the normal volumes of traffic will quickly overload these smaller sized network components. These smaller devices will simply not handle the necessary volumes. Sizing of network components based on expected traffic loads is critical to success. Remember experience indicates that when a disaster occurs excesses in traffic volumes will be experienced.

15. Security

Hackers and crackers prey on weak networks. They are like a feeding frenzy of sharks in bloody waters. As soon as they realize a business has suffered problems they look to see if they can breach the information and network security of that organization. Many times they are successful because the same level of security as is found at the organizations primary site is not found at its disaster recovery sites. Firewalls, intrusion detection, virus protection, access controls and the like MUST be at the same level of protection or there will be security breaches. Count on it.

16. Plan exercising

Plans need to be tested to ensure that they will work when they are absolutely needed. As stated before, during times of crisis people are asked to perform under very difficult circumstances. The human thinking process is often times obfuscated mainly due to the stress of the moment. It is important to make sure that each and every participant of the disaster recovery plan knows what is expected of them and they have had an opportunity to perform their duties under better than disastrous conditions. Further, regular planning allows an organization to see if the disaster recovery plan remains fit-for-purpose, that changes to patches and network addresses have been incorporated, and that nothing has changed since the last exercise.

Topic : Preparing To Develop The Disaster Recovery Plan

Topic Objective:

At the end of this topic student would be able to:

- Overview of Network Disaster Recovery
- Explain Why Is Disaster Recovery Important?
- Define Disaster Recovery Planning
- Describe Disaster Recovery Techniques
- Highlight the Idea of Recovery
- Evaluate Information security
- Analyze Risk management

Definition/Overview:

The Objective of IT Continuity Management is to safeguard the performance of service in any eventuality based on planning and implementation of preventive measures. Enterprises depend to a significant on the availability and functionality of the information technology in use. Therefore, preparation for an eventuality, combined with business continuity management, assumes ever greater importance, with the specific goal of safeguarding the availability of services, taking preventive measures to reduce the probability of failures and, if a catastrophic event should occur, restoring services in the required time.

Key Points:

1. Overview of Network Disaster Recovery

IT professionals have recognized the importance of disaster recovery for decades. Both the terrorist attacks of 11 September 2001 and recent IT technology trends have led to more widespread awareness of disaster recovery and other business continuity issues. Organizations face some tough choices, though, in planning for the future. In IT, disaster recovery involves a series of actions to be taken in the event of major unplanned outages to minimize their adverse effects. Disasters can result from events such as:

- Hacker attacks
- Computer viruses
- Electric power failures
- Underground cable cuts or failures
- Fire, flood, earthquake, and other natural disasters at a facility
- Mistakes in system administration

The related concept of business continuity involves insuring that an organization's critical business processes, including those utilizing IT systems, can be maintained in the event of a disaster.

2. Why Is Disaster Recovery Important?

When executed well, disaster recovery procedures save large sums of money. Disaster recovery can also improve the quality of human life, and it may even save lives. The terrorist attacks of 11 September, for example, caused large-scale network outages. Among the affected systems were some of the fiber optic telecommunications services provided by Verizon. Besides the financial impact to Wall Street firms from lost data connectivity, the loss of voice contact with friends and family greatly affected many individuals on that day.

3. Disaster Recovery Planning

The best approach to disaster recovery focuses primarily on planning and prevention. While the damage resulting from the events of 11 September could not have been anticipated, many other more typical disaster scenarios can be analyzed in detail. For those events that can't be prevented, an IT disaster recovery plan takes into account the need to

- Detect the outages or other disaster effects as quickly as possible
- Notify any affected parties so that they can take action
- Isolate the affected systems so that damage cannot spread
- Repair the critical affected systems so that operations can be resumed

4. Disaster Recovery Techniques

All good IT disaster recovery plans consider the three main components of operations:

- Data
- Systems
- People

From the technical perspective, most organizations rely on some form of redundancy to make possible the recovery of data and systems. Redundancy allows secondary data or system resources to be pressed into service on short notice should primary resources fail or otherwise become unavailable. Traditional backup strategies, for example, archive copies of critical data at a given point in time so that they can be restored later if needed. Organizations may also choose to replicate servers and other critical hardware at multiple locations to guard against any single point of failure. More advanced network technologies, like SONET, and some forms of clustering, incorporate built-in failover capabilities that attempt to automatically recover from some failures. While these and similar approaches have been a part of IT practice for many years, more sophisticated disaster recovery techniques have grown in popularity due to the events of 11 September 2001.

Periodic data backups, for example, have limited value if the "snapshots" are not taken frequently enough. Some organizations now generate so much data that even daily backups are too infrequent. A more sophisticated approach like disk mirroring ensures that data remain available from multiple sources in near real-time. However, traditional mirroring only works over limited distances. Storage area network (SAN) and other competing technologies can alleviate this problem, albeit at a higher cost. Another recent trend in IT disaster recovery planning, third-party relocation services, gives organizations access to fully-equipped operations space at temporary facilities in remote locations. These facilities can be a wonderful option in times of crisis... if trained personnel are available to staff them.

5. Idea of Recovery

Data recovery is not a new idea. In recent years, data has become a vitally important corporate asset essential to business continuity. The ability to recover crucial data quickly after a disaster is a fundamental requirement of economic viability. Have the terrorist attacks in New York City and Washington, D.C. changed the way we will do business in the future? On this page, we explore the concept and techniques of business continuity planning as well as examine how present systems functioned on Sept. 11, 2001. We also try to analyze how this particular disaster may change the way the way we do business. The recovery time objective (RTO) is the maximum tolerable length of time that a computer, system, network, or application can be down after a failure or disaster occurs. The RTO is a function of the extent to which the interruption disrupts normal operations and the amount of revenue lost per unit time as a result of the disaster. These factors in turn depend on the affected equipment and application.

6. Information security

Information security means protecting information and information systems from unauthorized access, use, disclosure, disruption, modification or destruction. The terms information security, computer security and information assurance are frequently incorrectly used interchangeably. These fields are interrelated often and share the common goals of protecting the confidentiality, integrity and availability of information; however, there are some subtle differences between them. These differences lie primarily in the approach to the subject, the methodologies used, and the areas of concentration. Information security is concerned with the confidentiality, integrity and availability of data regardless of the form the data may take: electronic, print, or other forms. Computer security can focus on ensuring the availability and correct operation of a computer system without concern for the information stored or processed by the computer. Governments, military, corporates, financial institutions, hospitals, and private businesses amass a great deal of confidential information about their employees, customers, products, research, and financial status. Most of this information is now collected, processed and stored on electronic computers and transmitted across networks to other computers.

Should confidential information about a businesss customers or finances or new product line fall into the hands of a competitor, such a breach of security could lead to lost business, law suits or even bankruptcy of the business. Protecting confidential information is a business requirement, and in many cases also an ethical and legal requirement. For the individual, information security has a significant effect on privacy, which is viewed very differently in different cultures. The field of information security has grown and evolved significantly in recent years. As a career choice there are many ways of gaining entry into the field. It offers many areas for specialization including, securing network(s) and allied infrastructure, securing applications and databases, security testing, information systems auditing, business continuity planning and digital forensics science, to name a few.

7. Risk management

A comprehensive treatment of the topic of risk management is beyond the scope of this article. We will however, provide a useful definition of risk management, outline a commonly used process for risk management, and define some basic terminology. The CISA Review Manual 2006 provides the following definition of risk management: "Risk management is the process of identifying vulnerabilities and threats to the information resources used by an organization in achieving business objectives, and deciding what countermeasures, if any, to take in reducing risk to an acceptable level, based on the value of the information resource to the organization." There are two things in this definition that may need some clarification. First, the process of risk management is an ongoing iterative process. It must be repeated indefinitely. The business environment is constantly changing and new threats and vulnerability emerge every day. Second, the choice of countermeasures (controls) used to manage risks must strike a balance between productivity, cost, effectiveness of the countermeasure, and the value of the informational asset being protected.

Risk is the likelihood that something bad will happen that causes harm to an informational asset (or the loss of the asset). Vulnerability is a weakness that could be used to endanger or cause harm to an informational asset. A threat is anything (man made or act of nature) that has the potential to cause harm. The likelihood that a threat will use a vulnerability to cause harm creates a risk. When a threat does use a vulnerability to inflict harm, it has an impact. In the context of information security, the impact is a loss of availability, integrity, and confidentiality, and possibly other losses (lost income, loss of life, loss of real property). It should be pointed out that it is not possible to identify all risks, nor is it possible to eliminate all risk. The remaining risk is called residual risk. A risk assessment is carried out by a team of people who have knowledge of specific areas of the business. Membership of the team may vary over time as different parts of the business are assessed. The assessment may use a subjective qualitative analysis based on informed opinion, or where reliable dollar figures and historical information is available, the analysis may use quantitative analysis.

In Section 2 of this course you will cover these topics:

- Assessing Impact And Risks In The Enterprise
- Prioritizing Systems And Functions For Recovery

Topic : Assessing Impact And Risks In The Enterprise

Topic Objective:

At the end of this topic student would be able to:

- Explain IBM Global Mirror
- Highlight Key concepts
- Describe Backup site
- Elaborate Network Visibility: The Key to Risk Management

Definition/Overview:

Business continuity planning (BCP) is the creation and validation of a practiced logistical plan for how an organization will recover and restore partially or completely interrupted critical (urgent) functions within a predetermined time after a disaster or extended disruption.

The logistical plan is called a business continuity plan. In plain language, BCP is working out how to stay in business in the event of disaster. Incidents include local incidents like building fires, regional incidents like earthquakes, or national incidents like pandemic illnesses. BCP may be a part of an organizational learning effort that helps reduce operational risk associated with lax information management controls. This process may be integrated with improving information security and corporate reputation risk management practices.

In December 2006, the British Standards Institution (BSI) released a new independent standard for BCP BS 25999-1. Prior to the introduction of BS 25999, BCP professionals relied on BSI information security standard BS 7799, which only peripherally addressed BCP to improve an organization's information security compliance. BS 25999's applicability extends to organizations of all types, sizes, and missions whether governmental or private, profit or non-profit, large or small, or industry sector. In 2007, the BSI published the second part, BS 25999-2 "Specification for Business Continuity Management", that specifies requirements for implementing, operating and improving a documented Business Continuity Management System (BCMS). In 2004, the United Kingdom enacted the Civil Contingencies Act 2004, a statute that instructs all emergency services and local authorities to actively prepare and plan for emergencies. Local authorities also have the legal obligation under this act to actively lead promotion of business continuity practices amongst its geographical area.

Key Points:

1. IBM Global Mirror

Global Mirror is an IBM technology that provides data replication over extended distances between two sites for business continuity and disaster recovery. If adequate bandwidth exists, Global Mirror provides an recovery point objective (RPO) of as low as 3-5 seconds between the two sites at extended distances with no performance impact on the application at the primary site. It replicates the data asynchronously and also forms a consistency group at a regular interval allowing a clean recovery of the application. The two sites can be on separate continents or simply on different utlility grids. IBM also provides a synchronous data replication called Metro Mirror, which is designed to support replication at "Metropolitan" distances of (normally) less than 300 km.

Global Mirror is based on existing IBM Copy Services functions: Global Copy and FlashCopy. Global Mirror periodically invokes a point-in-time copy at the primary site, at regular intervals, without impacting the I/O to the source volumes. Then it transfers the copy to the recovery site. By grouping many volumes into one Global Mirror session multiple volumes may be copied to the recovery site simultaneously while maintaining point-in-time consistency across those volumes. Global Mirror can be combined with a wide area clustering product like Geographically Dispersed Parallel Sysplex (GDPS), HACMP/XD, or CAW to provide for automated failover between sites. This combined solution provides lower recovery time objective (RTO), because it allows most applications to automatically resume productive operation in 30-600 seconds. The Global Mirror function is available on IBM's enterprise storage devices including the DS8100, the DS8300, the DS6800, the Enterprise Storage Server Models 800 and 750, midrange storage servers DS4000 family and the IBM SAN Volume Controller.

2. Key concepts

For over twenty years information security has held that confidentiality, integrity and availability (known as the CIA Triad) as the core principles of information security.

- **Confidentiality**

Confidentiality is the property of preventing disclosure of information to unauthorized individuals or systems. For example, a credit card transaction on the Internet requires the credit card number to be transmitted from the buyer to the merchant and from the merchant to a transaction processing network. The system attempts to enforce confidentiality by encrypting the card number during transmission, by limiting the places where it might appear (in databases, log files, backups, printed receipts, and so on), and by restricting access to the places where it is stored. If an unauthorized party obtains the card number in any way, a breach of confidentiality has occurred. Breaches of confidentiality take many forms. Permitting someone to look over your shoulder at your computer screen while you have confidential data displayed on it could be a breach of confidentiality. If a laptop computer containing sensitive information about a company's employees is stolen or sold, it could result in a breach of confidentiality. Giving out confidential information over the telephone is a breach of confidentiality if the caller is not authorized to have the information.

Confidentiality is necessary for maintaining the privacy of the people whose personal information a system holds.

- **Integrity**

In information security, integrity means that data cannot be modified without authorization. This is not the same thing as referential integrity in databases. Integrity is violated when an employee accidentally or with malicious intent deletes important data files, when a computer virus infects a computer, when an employee is able to modify his own salary in a payroll database, when an unauthorized user vandalizes a web site, when someone is able to cast a very large number of votes in an online poll, and so on. There are many ways in which integrity could be violated without malicious intent. In the simplest case, a user on a system could mis-type someone's address. On a larger scale, if an automated process is not written and tested correctly, bulk updates to a database could alter data in an incorrect way, leaving the integrity of the data compromised. Information security professionals are tasked with finding ways to implement controls that prevent errors of integrity.

- **Availability**

For any information system to serve its purpose, the information must be available when it is needed. This means that the computing systems used to store and process the information, the security controls used to protect it, and the communication channels used to access it must be functioning correctly. High availability systems aim to remain available at all times, preventing service disruptions due to power outages, hardware failures, and system upgrades. Ensuring availability also involves preventing denial-of-service attacks. In 2002, Donn Parker proposed an alternative model for the classic CIA triad that he called the six atomic elements of information. The elements are confidentiality, possession, integrity, authenticity, availability, and utility. The merits of the Parkerian hexad are a subject of debate amongst security professionals.

- **Authenticity**

In computing, e-Business and information security it is necessary to ensure that the data, transactions, communications or documents (electronic or physical) are genuine. It is also important for authenticity to validate that both parties involved are who they claim they are.

- **Non-repudiation**

In law, non-repudiation implies one's intention to fulfill their obligations to a contract. It also implies that one party of a transaction cannot deny having received a transaction nor can the other party deny having sent a transaction. Electronic commerce uses technology such as digital signatures and encryption to establish authenticity and non-repudiation.

3. Backup site

A backup site is a location where an organization can easily relocate following a disaster, such as fire, flood, terrorist threat or other disruptive event. This is an integral part of the disaster recovery plan and wider business continuity planning of an organisation. A backup site can be another location operated by the organisation, or contracted via a company that specializes in disaster recovery services. In some cases, an organisation will have an agreement with a second organisation to operate a joint backup site. There are three types of backup sites, including cold sites, warm sites, and hot sites. The differences between the types are determined by the costs and effort required to implement each. Another term used to describe a backup site is a work area recovery site.

- **Cold Sites**

A cold site is the most inexpensive type of backup site for an organization to operate. It does not include backed up copies of data and information from the original location of the organization, nor does it include hardware already set up. The lack of hardware contributes to the minimal startup costs of the cold site, but requires additional time following the disaster to have the operation running at a capacity close to that prior to the disaster.

- **Hot Sites**

A hot site is a duplicate of the original site of the organization, with full computer systems as well as near-complete backups of user data. Real time synchronization between the two sites may be used to completely mirror the data environment of the original site using wide area network links and specialized software. Following a disruption to the original site, the hot site exists so that the organization can relocate with minimal losses to normal operations. Ideally, a hot site will be up and running within a matter of hours or even less. Personnel may still have to be moved to the hot site so it is possible that the hot site may be operational from a data processing perspective before staff has relocated. The capacity of the hot site may or may not match the capacity of the original site depending on the organizations requirements. This type of backup site is the most expensive to operate. Hot sites are popular with organizations that operate real time processes such as financial institutions, government agencies and ecommerce providers.

- **Warm Sites**

A warm site is, quite logically, a compromise between hot and cold. These sites will have hardware and connectivity already established, though on a smaller scale than the original production site or even a hot site. Warm sites will have backups on hand, but they may not be complete and may be between several days and a week old. An example would be backup tapes sent to the warm site by courier.

- **Choosing**

Choosing the type is mainly decided by an organisations cost vs. benefit strategy. Hot sites are traditionally more expensive than cold sites since much of the equipment the company needs has already been purchased and thus the operational costs are higher. However if the same organisation loses a substantial amount of revenue for each day they are inactive then it may be worth the cost. Another advantage of a hot site is that it can be used for operations prior to a disaster happening. The advantages of a cold site are simple--cost. It requires much fewer resources to operate a cold site because no equipment has been bought prior to the disaster. The downside with a cold site is the potential cost that must be incurred in order to make the cold site effective. The costs of purchasing equipment on very short notice may be higher and the disaster may make the equipment difficult to obtain. When contracting services from a commercial provider of backup site capability organisations should take note of contractual usage provision and invocation procedures, providers may sign up more than one organisation for a given site or facility, often depending on various service levels. This is a reasonable proposition as it is unlikely that all organisations using the service are likely to need it at the same time and it allows the provider to offer the service at an affordable cost. However, in a large scale incident that affects a wide area it is likely that these facilities will become over subscribed.

4. Network Visibility: The Key to Risk Management

Although a subset of network-based problems (including computer worms and viruses, intrusion attempts, and denial-of-service attacks on business Web sites) often garners most of the attention from IT security groups, they are only part of the risk picture. IT traditionally has responded to these threats in piecemeal fashion, often by buying a "black box" solution, but these exposures need to be evaluated in the context of the total risk picture. Responses need to be planned according to the business value of each risk and the available resources for mitigation. Further, they must be planned with the realization that risk can never be eliminated and that part of risk mitigation is planning for the inevitable events that will occur no matter how well IT protects itself. IT faces three major classes of exposures:

- Dioecious - having unisexual reproductive units with male and female plants. (flowers, conifer cones, or functionally equivalent structures) occurring on different individuals; from Greek for "two households". Individual plants are not called dioecious: they are either gynoecious (female plants) or androecious (male plants).
- **Technology risks**: These are the traditional concerns of IT security such as viruses and hardware failures. The best mitigation strategies for these issues are usually technical tools backed by strong policies and specific tactics for dealing with problems. Today many of these problems are network-related, making strong network management a central tool for risk mitigation.
- **Legal and personnel risks**: These are varied and can include hostile workplace suits, document preservation to meet legal discovery requirements, and sabotage or business espionage by employees. For the most part, mitigation hinges on good management practices.

Managers should be trained in the skills of management, just as engineers are trained in the knowledge and skills of their jobs.

- **Natural and man-made disasters**: These include fire, flood, earthquake, storm, and, in this post 9/11 world, terrorist action. Mitigation includes setting up data centers in low risk areas and disaster recovery.

Chances are there is something from each one of these categories that keeps most IT managers up at night. We can't eliminate risk, but we can prepare for it. We offer three steps IT managers should consider when establishing a "reality-based" mitigation plan.

These risks are not always obvious. For instance, one of the most important under-identified IT risks is the chronic slowdown of traffic to vital IT assets due to network overload. This is particularly true as organizations replace traditional analog telephone systems with internal VoIP and as data traffic consists of larger media files (such as MP3s, JPEGs, and digital video), all of which increase network demand. If a delay of one minute in traffic supporting a key IT asset can cause a business loss, then what is the business value of chronic delays of fractions of a minute for every transaction sustained over much of the business day? These issues make strong network management (including prioritization and traffic policy automation) vital. Identifying the business value of IT assets gives network managers a huge head start in traffic prioritization because these figures can provide a good basis for developing the prioritization plan. The overall risk tolerance of your organization must also be assessed. Risk tolerance varies widely. For example, financial institutions tend to be risk averse; highly entrepreneurial companies may be risk tolerant. Any exposure that falls inside your organization's risk tolerance can be given lower priority for mediation.

Topic : Prioritizing Systems And Functions For Recovery

Topic Objective:

At the end of this topic student would be able to:

- Understand Network communication
- Explain the Considerations of network issues
- Evaluate Location of People and Resources
- Elaborate Production Network Specifications
- Highlight Applications
- Identify about Security
- Learn about Recovery Options
- Know about Frame Relay Network Recovery
- Define Private Line Network Recovery
- Describe Satellite (VSAT)
- Analyze Virtual Private Network (VPN)

Definition/Overview:

A network backup/recovery system is provided which can perform backup operation and/or recovery operation in a computer system, so as to backup data and system configurations and/or recover from a system crash, which can also minimize time consumption and cost. A network backup/recovery system backs up data and immediately returns a storage device to a preceding state. The network backup/recovery system comprises a client management part for controlling at least one client and a server management part for controlling a server and monitoring each using condition of the clients. The server management part sends a predetermined message to the client management part according to each using condition of the clients; and the client management part involves implementing a predetermined backup/recovery procedure according to the received predetermined message.

Key Points:

1. Network communication

Network communication is arguably the most critical component in the complex information fabric of the modern business world. Although systems and applications store and manipulate data, network communication allows data to travel worldwide nearly instantaneously. Network data flows are responsible for carrying everything from media and entertainment content to financial transactions and medical information. Missioncritical systems and applications are useless without being tied into the resources and users that allow them to fulfill business functions. Network communication is responsible for providing this connection. Disaster Recovery and Business Continuity (DR/BC) programs must account for potential disruptions to network communications. Network communications recovery must also be considered if business processing is resumed at an alternate location. This white paper discusses some of the considerations that need to be accounted for when designing wide-area network (WAN) recovery solutions, as well as providing several strategies for bringing network resources back online following an unplanned outage.

2. Considerations

Numerous considerations must be factored in when planning for an effective network recovery strategy. These consist of many of the same considerations made when designing a production network. The fundamental difference is that unplanned outages come in many shapes and sizes, and a sound network recovery strategy must adapt to as many disruptive events as possible. The following section discusses some of the metrics that must be captured to understand network recovery requirements.

3. Location of People and Resources

It is imperative to know where the critical people and resources reside during recovery. Key personnel responsible for critical business functions require access to the resources they need to perform their job functions. Resources can include systems, applications, or information. The locations of these people and resources determine where connectivity must be implemented and what the available options are. Locations can vary depending on the nature of the outage. Some disruptive events might require full recovery at locations ranging from hot sites to mobile facilities, whereas others might require only a partial recovery where few employees travel to another office or even their homes. Regardless of the nature of the disaster, the network recovery strategy should be flexible enough to adapt to support conditions. Critical IT staff should also be considered as they will most likely be responsible for executing the recovery. It is imperative that recovery plans provision the appropriate staff and place them in a position to bring network resources back online.

4. Production Network Specifications

The production network specifications play a large role in determining what recovery options are available. Network technologies vary greatly in type, complexity, and interoperability. In common network implementations, numerous technologies and network carriers might be deployed individually or together. Each technology has unique characteristics and requirements that must be considered from a network recovery perspective. In many cases, recovery can be accounted for during the implementation of the production network through the deployment of redundant connections. Carriers usually offer circuit redundancy that can seamlessly fail over traffic flows if a circuit outage occurs. If the production network does not have appropriate levels of redundancy, steps must be taken to incorporate network recovery into the BC/DR program. Networks rely on technologies including T-1/E-1, DS-x, frame relay, ATM, SONET, and more recently, Multi Protocol Label Switching (MPLS) and Metro Ethernet. When formulating a network recovery strategy, the technology providing the production communication must be considered. In addition to the technology, the carrier also plays a large role in determining what options are available for recovery. Many carriers offer services that are geared for DR. In the recovery options section, examples of network recovery designs are discussed.

5. Applications

Applications requirements for bandwidth, protocols, and security vary greatly based on many factors. For example, an Internet facing web application might have vastly different network requirements than an internal application. These requirements should be documented upfront to ensure that recovery connectivity supports the required applications. Bandwidth is a key consideration with all applications. Bandwidth requirements might be less than production if there are fewer users on the system. Bandwidth utilized by normal user activity and other activities such as application integration points should be planned for. Although most applications utilize TCP/IP for communication, there can be legacy applications requiring older protocols, such as SNA or IPX/SPX. These apps require special attention during recovery planning to ensure compatibility. There are typically interdependencies between applications that need to be accounted for. These are often overlooked during recovery efforts rendering the application useless during a recovery effort. Internal and external (third-party) feeds between systems must be provisioned for in the network recovery plan.

6. Security

During a recovery effort, businesses often overlook security in their haste to bring resources back online. This can be a costly mistake that can incur more damage than the original outage. Security might even be more paramount if the nature of the outage is related to a security breech or malicious attack. Attackers can specifically target businesses recovering from an unplanned outage in hopes that production controls were bypassed in the interest of a rapid recovery. Security protocol and procedures must be adhered to at every step of the recovery process just as in normal production operations. This can be accomplished by ensuring that the security team has input in the formulation of recovery strategies and their associated plans. From a WAN perspective, resources should be secured against unauthorized access with the appropriate mechanisms. Prior to the resumption of business operations, the environment needs to be validated by qualified individuals to ensure security exists at the appropriate levels. If controls are bypassed because of the circumstances of a disaster, they should be analyzed against the associated risk upfront and approved by the appropriate personnel.

7. Recovery Options

After the considerations in the previous section have been made, you can determine the options for network recovery. This section discusses the characteristics of common recovery strategies and their advantages and disadvantages. In a recovery scenario, lower layer connectivity must first be restored before higher layer application communication can occur. Restoring lower layer connectivity requires that physical links be in place. The following section provides recovery solutions based on several common network topologies found in the field. These include Frame Relay and private line networks. Two additional recovery options that will be discussed are Satellite and virtual private network (VPN).

8. Frame Relay Network Recovery

Cloud-based technologies such as Frame Relay implement connections through virtual circuits rather than physical ones. This characteristic facilitates recovery of these types of connections. Through the use of backup circuits, connections can often be quickly recovered to an alternate location. Most of the major carriers offer services that include options to redirect circuits to alternate locations. A requirement for this strategy is that the target location for recovery has an existing connection to the respective carrier's network and the appropriate termination equipment for the circuit. If DR has been outsourced, the recovery provider should have shared access to your carrier's service, which can be configured to recover your network on demand. This can lower cost by avoiding procurement of additional circuits that are used only in a disaster or DR test. To initiate the failover, the subscriber simply places a call to the carriers and instructs them on which circuits need to be redirected. The carrier executes a predetermined script to transfer the circuits to the backup location. This strategy works well in topologies where several remote locations require communication with the primary data center at a main corporate office. The following diagram illustrates this type of recovery.

After the circuits have been redirected to the recovery site, configuration steps might be necessary to bring up the terminating equipment and configure protocols and routing. Procedure for this activity should be integrated into the recovery plan.

9. Private Line Network Recovery

Private lines include T-1, E-1, and DS-3, and other point-to-point private connections. Recently, Metro-Ethernet services have become popular and might also fall into this category. Networks utilizing private lines require that recovery be provisioned into the initial WAN design. If the WAN is sufficiently meshed, you might not need to take additional recovery measures. In many instances, this is not the case, and provisions must be made to recover critical connections following an unplanned outage. Because private line circuits implement end-to-end physical connections, a secondary connection is typically required for recovery purposes. This connection can be a mirror of the primary circuit that is always active. Circuits might be provisioned with separate carriers for additional diversity. A lower bandwidth standby circuit can also be utilized. For example, a T-1 circuit might have a backup circuit that utilizes DSL or ISDN. When private lines form a hub-and-spoke network topology, recovery of a single circuit can typically be easily accomplished through the mechanisms previously detailed. However, if the primary site is completely lost, connectivity must be recovered to the target recovery site. This is typically accomplished through an additional circuit to the target recovery site. The following diagram depicts this scenario.

Recovery of hub-and-spoke networks utilizing private lines can be costly. An alternative to deploying additional connections to the recovery site would be to bring critical users to the recovery location. Additional alternatives are discussed in the following section.

10. Satellite (VSAT)

Very Small Aperture Terminal (VSAT) provides for the transmission voice, video, and data via satellite. VSAT technologies have advanced significantly over recent years and provide an excellent mechanism for network recovery. VSAT can be particularly useful in a mobile recovery scenario because it provides true flexibility in recovery location. Also, many traditional circuits do not survive a large disaster, leaving VSAT as one of the few options to provide connectivity to disaster-affected areas. When VSAT is utilized to recover a remote location, it ultimately ties into a terrestrial circuit that completes the recovery connection. This terrestrial circuit might tie into a private WAN, or the public Internet. The connectivity might provide access to normal production resources or possibly to resources that have been recovered to an alternate location. The following diagram depicts a sample VSAT recovery environment.

In addition to the recovery of a remote office, VSAT can also be combined with mobile units for the recovery of data center resources. Due to the extended transmission distance, VSAT introduces significant latency into the communication process. Prior to the deployment of a VSAT recovery solution, you should first insure that your applications can tolerate connections with high latency.

11. Virtual Private Network (VPN)

Virtual private networks are an excellent technology for use in disaster recovery scenarios. If deployed properly, they are capable of quickly providing secure links between resources in virtually any location. The primary requirements are a connection to the Internet with the appropriate bandwidth and a VPN-capable device at the target locations. VPNs also support the implementation of multi-point connections that can be utilized to recover multiple locations during a disaster. The following diagram depicts a sample VPN recovery environment.

VPN recovery can also be combined with other technologies such as VSAT to provide a more flexible and secure recovery network. In recent years, there has been a substantial push toward Multi Protocol Label Switching (MPLS). Although MPLS still requires technologies such as Frame Relay on the network edge, it enables much greater flexibility in redirecting connections because the virtual circuits are formed in software on routers. If your network utilizes MPLS, your carrier can provide a network recovery solution that integrates cleanly into your existing service.

In Section 3 of this course you will cover these topics:

- Identifying Data Storage And Recovery Sites
- Developing Plans And Procedures, And Relationships

Topic : Identifying Data Storage And Recovery Sites

Topic Objective:

At the end of this topic student would be able to:

- Explain Encrypting File System

- Elaborate Bootable business card
- Highlight Computer data storage
- Examine Hierarchy of storage
- Identify Information processing
- Define Signal processing
- Analyze Index card

Definition/Overview:

Network communications are an integral component of modern business operations and, therefore, must be provisioned for in DR/BC planning. Many organizations overlook the network component when formulating their plans resulting in an unrecoverable environment. Depending on environment characteristics, there are multiple options that can be leveraged to provide a successful network recovery strategy. Due to the vast differences in networks deployed today, these characteristics need to be carefully considered before a specific strategy is selected. Network technologies are constantly evolving, and carriers are now offering DR services within their standard services.

The network recovery strategy must be directly incorporated into the overall DR/BC plans and maintained on a regular basis. If carefully planned, implemented and maintained, your network recovery strategy will insure that critical communications components are available to the business following a disruptive event.

Key Points:

1. Encrypting File System

The Encrypting File System (EFS) is a file system driver that provides filesystem-level encryption in Microsoft Windows (2000 and later) operating systems, except Windows XP Home Edition, Windows Vista Basic, and Windows Vista Home Premium. The technology enables files to be transparently encrypted on NTFS file systems to protect confidential data from attackers with physical access to the computer. User authentication and access control lists can protect files from unauthorized access while the operating system is running, but are easily circumvented if an attacker gains physical access to the computer. One solution is to store the files encrypted on the disks of the computer. EFS does this using public key cryptography, and aims to ensure that decrypting the files is extremely difficult without the correct key. However, EFS is in practice susceptible to brute-force attacks against the user account passwords. In other words, encryption of files is only as strong as the password to unlock the decryption key.

EFS works by encrypting a file with a bulk symmetric key, also known as the File Encryption Key, or FEK. It uses a symmetric encryption algorithm because it takes a relatively smaller amount of time to encrypt and decrypt large amounts of data than if an asymmetric key cipher is used. The symmetric encryption algorithm used will vary depending on the version and configuration of the operating system; see #Algorithms Used by Operating System Version below. The FEK (the symmetric key that is used to encrypt the file) is then encrypted with a public key that is associated with the user who encrypted the file, and this encrypted FEK is stored in the $EFS alternate data stream of the encrypted file. To decrypt the file, the EFS component driver uses the private key that matches the EFS digital certificate (used to encrypt the file) to decrypt the symmetric key that is stored in the $EFS stream. The EFS component driver then uses the symmetric key to decrypt the file. Because the encryption& decryption operations are performed at a layer below NTFS, it is transparent to the user and all their applications.

Folders whose contents are to be encrypted by the file system are marked with an encryption attribute. The EFS component driver treats this encryption attribute in a way that is analogous to the inheritance of file permissions in NTFS: if a folder is marked for encryption, then by default all files and subfolders that are created under the folder are also encrypted. When encrypted files are moved within an NTFS volume, the files remain encrypted. However, there are a number of occasions in which the file could be decrypted without the user explicitly asking Windows to do so. Files and folders are decrypted before being copied to a volume formatted with another file system, like FAT32. Finally, when encrypted files are copied over the network using the SMB/CIFS protocol, the files are decrypted before they are sent over the network. The most significant way of preventing the decryption-on-copy is using backup applications that are aware of the "Raw" APIs. Backup applications that have implemented these Raw APIs will simply copy the encrypted file stream and the $EFS alternate data stream as a single file. In other words, the files are "copied" (e.g. into the backup file) in encrypted form, and are not decrypted during backup. Starting with Windows Vista, a user's private key can be stored on a smart card; Data Recovery Agent (DRA) keys can also be stored on a smart card.

2. Bootable business card

A bootable business card (BBC) is a CD-ROM that has been cut, pressed, or molded to the size and shape of a business card (designed to fit in a wallet or pocket). Alternative names for this form factor include "credit card", "hockey rink" and "wallet-size". The cards are designed to hold about 50 MB. The CD-ROM business cards are generally used for commercial product demos, are mailed to prospective customers and are given away at trade shows. Although the term "bootable business card" could be applied to any bootable CD-ROM in the business card form factor, it almost always refers one which contains a compact Linux distribution generally containing a suite of system diagnostic and rescue tools and/or demos of specific packages.

The key of the bootable business card is that it runs completely from the CD and the system's memory (RAM), as several "Live" CD versions of Linux have been doing for years. One simply puts the CD disc into the drive, powers up the computer and ensures that the CD drive is selected for boot before the hard drive. Once booted, the operating system runs from the CD and out of the system's RAM. Because the business card form factor has such a small capacity the Linuxcare developers chose to use a compressed filesystem. This allows the typical BBC to contain about 100 megabytes of software in only about 50 megabytes of disc space. The original BBC and most of its clones and derivatives will scan the system for recognized filesystems, automatically "mounting" these up in read-only mode. This makes filesystems on any local hard disks accessible while minimizing the risk of inadvertent corruption, deletion or other damage to files on local drives. A typical BBC contains a suite of networking, back-up and data recovery utilities, which is why they are valued by Linux system administrators as rescue tools. Many BBCs use the cloop (compressed loopback) driver which provided a compressed read-only filesystem for Linux. Of course they typically have some of the system's memory (RAM) configured as a ramdisk (or perhaps several RAM disks). This typically leaves the CD-ROM drive dedicated for the duration of the system usage. However, some BBCs create a larger ramdisk and copy the entire system off the CD, thus making the drive available for other CDs or DVDs. This is useful because some PCs have only a single CD or DVD drive. Once booted, these systems provide a UNIX/Linux command line prompt (generally as the root user). Some also provide some very compact graphical user interface (GUI) tools. The LNX-BBC includes a small X (X Window System) server and a web browser called BrowseX (among other tools). At their core most BBCs are rescue and diagnostics tools for expert professionals, and normal user-operations are catered for better by Live CD distributions.

3. Computer data storage

Computer data storage, often called storage or memory, refers to computer components, devices, and recording media that retain digital data used for computing for some interval of time. Computer data storage provides one of the core functions of the modern computer, that of information retention. It is one of the fundamental components of all modern computers, and coupled with a central processing unit (CPU, a processor), implements the basic computer model used since the 1940s. In contemporary usage, memory usually refers to a form of semiconductor storage known as random access memory (RAM) and sometimes other forms of fast but temporary storage. Similarly, storage today more commonly refers to mass storage - optical discs, forms of magnetic storage like hard disks, and other types slower than RAM, but of a more permanent nature. Historically, memory and storage were respectively called primary storage and secondary storage. The contemporary distinctions are helpful, because they are also fundamental to the architecture of computers in general. As well, they reflect an important and significant technical difference between memory and mass storage devices, which has been blurred by the historical usage of the term storage. Nevertheless, this article uses the traditional nomenclature.

4. Hierarchy of storage

- **Primary storage**

Primary storage, presently known as memory, is the only one directly accessible to the CPU. The CPU continuously reads instructions stored there and executes them as required. Any data actively operated on is also stored there in uniform manner. Historically, early computers used delay lines, Williams tubes, or rotating magnetic drums as primary storage. By 1954, those unreliable methods were mostly replaced by magnetic core memory, which was still rather cumbersome. Undoubtedly, a revolution was started with the invention of a transistor, that soon enabled then-unbelievable miniaturization of electronic memory via solid-state silicon chip technology. This led to a modern random access memory (RAM). It is small-sized, light, but quite expensive at the same time. (The particular types of RAM used for primary storage are also volatile, i.e. they lose the information when not powered). As shown in the diagram, traditionally there are two more sub-layers of the primary storage, besides main large-capacity RAM:

- Processor registers are located inside the processor. Each register typically holds a word of data (often 32 or 64 bits). CPU instructions instruct the arithmetic and logic unit to perform various calculations or other operations on this data (or with the help of it). Registers are technically among the fastest of all forms of computer data storage.
- Processor cache is an intermediate stage between ultra-fast registers and much slower main memory. It's introduced solely to increase performance of the computer. Most actively used information in the main memory is just duplicated in the cache memory, which is faster, but of much lesser capacity. On the other hand it is much slower, but much larger than processor registers. Multi-level hierarchical cache setup is also commonly usedprimary cache being smallest, fastest and located inside the processor; secondary cache being somewhat larger and slower.

Main memory is directly or indirectly connected to the CPU via a memory bus, today sometimes referred to as a front side bus. It is actually comprised of two buses (not on the diagram): an address bus and a data bus. The CPU firstly sends a number through an address bus, a number called memory address, that indicates the desired location of data. Then it reads or writes the data itself using the data bus. Additionally, a memory management unit (MMU) is a small device between CPU and RAM recalculating the actual memory address, for example to provide an abstraction of virtual memory or other tasks. As the RAM types used for primary storage are volatile (cleared at start up), a computer containing only such storage would not have a source to read instructions from, in order to start the computer. Hence, non-volatile primary storage containing a small startup program (BIOS) is used to bootstrap the computer, that is, to read a larger program from non-volatile secondary storage to RAM and start to execute it. A non-volatile technology used for this purpose is called ROM, for read-only memory (the terminology may be somewhat confusing as most ROM types are also capable of random access).

Many types of "ROM" are not literally read only, as updates are possible; however it is slow and memory must be erased in large portions before it can be re-written. Some embedded systems run programs directly from ROM (or similar), because such programs are rarely changed. Standard computers do not store non-rudimentary programs in ROM, rather use large capacities of secondary storage, which is non-volatile as well, and not as costly. Recently, primary storage and secondary storage in some uses refer to what was historically called, respectively, secondary storage and tertiary storage.

- **Secondary storage**

A hard disk drive with protective cover removed. Secondary storage, or storage in popular usage, differs from primary storage in that it is not directly accessible by the CPU. The computer usually uses its input/output channels to access secondary storage and transfers the desired data using intermediate area in primary storage. Secondary storage does not lose the data when the device is powered downit is non-volatile. Per unit, it is typically also an order of magnitude less expensive than primary storage. Consequently, modern computer systems typically have an order of magnitude more secondary storage than primary storage and data is kept for a longer time there. In modern computers, hard disks are usually used as secondary storage. The time taken to access a given byte of information stored on a hard disk is typically a few thousandths of a second, or milliseconds. By contrast, the time taken to access a given byte of information stored in random access memory is measured in billionths of a second, or nanoseconds. This illustrates the very significant access-time difference which distinguishes solid-state memory from rotating magnetic storage devices: hard disks are typically about a million times slower than memory. Rotating optical storage devices, such as CD and DVD drives, have even longer access times.

Some other examples of secondary storage technologies are: flash memory (e.g. USB sticks or keys), floppy disks, magnetic tape, paper tape, punch cards, standalone RAM disks, and Zip drives. The secondary storage is often formatted according to a filesystem format, which provides the abstraction necessary to organize data into files and directories, providing also additional information (called metadata) describing the owner of a certain file, the access time, the access permissions, and other information. Most computer operating systems use the concept of virtual memory, allowing utilization of more primary storage capacity than is physically available in the system. As the primary memory fills up, the system moves the least-used chunks (pages) to secondary storage devices (to a swap file or page file), retrieving them later when they are needed. As more of these retrievals from slower secondary storage are necessary, the more the overall system performance is degraded.

- **Tertiary storage**

Tertiary storage or tertiary memory, provides a third level of storage. Typically it involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage device according to the system's demands; this data is often copied to secondary storage before use. It is primarily used for archival of rarely accessed information since it is much slower than secondary storage (e.g. 5-60 seconds vs. 1-10 milliseconds). This is primarily useful for extraordinarily large data stores, accessed without human operators. Typical examples include tape libraries and optical jukeboxes. When a computer needs to read information from the tertiary storage, it will first consult a catalog database to determine which tape or disc contains the information. Next, the computer will instruct a robotic arm to fetch the medium and place it in a drive. When the computer has finished reading the information, the robotic arm will return the medium to its place in the library.

- **Off-line storage**

Off-line storage, also known as disconnected storage, is a computer data storage on a medium or a device that is not under the control of a processing unit. The medium is recorded, usually in a secondary or tertiary storage device, and then physically removed or disconnected. It must be inserted or connected by a human operator before a computer can access it again. Unlike tertiary storage, it cannot be accessed without human interaction. Off-line storage is used to transfer information, since the detached medium can be easily physically transported. Additionally in case a disaster, for example a fire, destroys the original data, a medium in a remote location will be probably unaffected, enabling disaster recovery. Off-line storage increases a general information security, since it is physically inaccessible from a computer, and data confidentiality or integrity cannot be affected by computer-based attack techniques. Also, if the information stored for archival purposes is accessed seldom or never, off-line storage is less expensive than tertiary storage. In modern personal computers, most secondary and tertiary storage media are also used for off-line storage. Optical discs and flash memory devices are most popular, and to much lesser extent removable hard disk drives. In enterprise uses, magnetic tape is predominant. Older examples are floppy disks, Zip disks, or punched cards.

5. Information processing

Information processing is the change (processing) of information in any manner detectable by an observer. As such, it is a process which describes everything which happens (changes) in the universe, from the falling of a rock (a change in position) to the printing of a text file from a digital computer system. In the latter case, an information processor is changing the form of presentation of that text file. Information processing may more specifically be defined in terms used by Claude E. Shannon as the conversion of latent information into manifest information. Latent and manifest information is defined through the terms of equivocation (remaining uncertainty, what value the sender has actually chosen), dissipation (uncertainty of the sender what the receiver has actually received) and transformation.

Within the field of cognitive psychology, information processing is an approach to the goal of understanding human thinking. It arose in the 1940s and 1950s. The essence of the approach is to see cognition as being essentially computational in nature, with mind being the software and the brain being the hardware. The information processing approach in psychology is closely allied to cognitivism in psychology and functionalism in philosophy although the terms are not quite synonymous. Information processing may be sequential or parallel, either of which may be centralized or decentralized (distributed). The parallel distributed processing approach of the mid-1980s became popular under the name connectionism. In the early 1950s Friedrich Hayek was ahead of his time when he posited the idea of spontaneous order in the brain arising out of decentralized networks of simple units (neurons). However, Hayek is rarely cited in the literature of connectionism. In the 1970s, Abraham Moles and Frieder Nake were among the first to establish and analyze links between information processing and aesthetics.

6. Signal processing

Signal processing is the analysis, interpretation, and manipulation of signals. Signals of interest include: sound, images, time-varying measurement values and sensor data, for example biological data such as electrocardiograms, control system signals, telecommunication transmission signals such as radio signals, and many others. Processing of such signals includes: filtering (for example in tone controls, equalizers and image enhancement software), adaptive filtering (for example for echo-cancellation in a conference telephone, or separation of information from noise for aircraft identification by radar), compression (for example, image compression) and feature extraction (for example speech-to-text conversion), spectrum analysis (for example in magnetic resonance imaging), wavetable synthesis (in modems and music synthesizers), storage, digitalization and reconstruction. Signals are electrical representations of time-varying or spatial-varying physical quantities, either analog or digital, and may come from various sources. In the context of signal processing, arbitrary binary data streams and on-off signals are not considered as signals, but only analog and digital signals that are representations of analog physical quantities. In communication systems, signal processing occurs at OSI layer 1, the Physical Layer (modulation, equalization, multiplexing, radio transmission, etc) in the seven layer OSI model, as well as at OSI layer 6, the Presentation Layer (source coding, including analog-to-digital conversion and data compression).

7. Index card

An index card is heavy paper stock cut to a standard size. Index cards are often used for recording individual items of information that can then be easily rearranged and filed (i.e. random access memory). The most common size in the United States and Russiais 3 in by 5 in (76 by 127 mm), hence the common name 3-by-5 card. Other sizes widely available include 4 in by 6 in (102 by 152 mm), 5 in by 8 in (127 by 203 mm) and ISO-size A7 (74 mm by 105 mm). Cards are available in blank, ruled and grid styles in a variety of colors. Special divider cards with protruding tabs and a variety of cases and trays to hold the cards are also sold by stationers.

As the name implies, index cards were widely used in the nineteenth and twentieth centuries to create an index to large collections of documents. A major law firm, for example, might have a room full of metal cabinets with drawers designed to hold index cards. Clerks might fill out several cards for an individual document or legal case, allowing them to be filed alphabetically under a number of terms.

Topic : Developing Plans And Procedures, And Relationships

Topic Objective:

At the end of this topic student would be able to:

- Analyze Risk Analysis
- Explain Establish the Budget
- Elaborate Develop the Plan
- Highlight Procedure
- Evaluate Recovery point objective
- Define Recovery time objective

Definition/Overview:

The key to surviving such an event is a business continuity strategy, a set of policies and procedures for reacting to and recovering from an IT-disabling disaster, and the main component of a business continuity strategy is a disaster recovery plan (DRP). In this article, DevX and Cole Emerson, President of Cole Emerson& Associates, Inc., a business-continuity consulting firm, and chairman of the board of DRI International, administrators of a global certification program for business continuity/disaster recovery planners, walk through the basics of creating an effective DRP.

Key Points:

1. Risk Analysis

The first step in drafting a disaster recovery plan is conducting a thorough risk analysis of your computer systems. List all the possible risks that threaten system uptime and evaluate how imminent they are in your particular IT shop. Anything that can cause a system outage is a threat, from relatively common manmade threats like virus attacks and accidental data deletions to more rare natural threats like floods and fires. Determine which of your threats are the most likely to occur and prioritize them using a simple system: rank each threat in two important categories, probability and impact. In each category, rate the risks as low, medium, or high. For example, a small Internet company (less than 50 employees) located in California could rate an earthquake threat as medium probability and high impact, while the threat of utility failure due to a power outage could rate high probability and high impact. So in this company's risk analysis, a power outage would be a higher risk than an earthquake and would therefore be a higher priority in the disaster recovery plan.

2. Establish the Budget

Once you've figured out your risks, ask 'what can we do to suppress them, and how much will it cost?' Can I detect a threat before it hits? How do I reduce the potential of it occurring? How do I minimize its impact to the business? For example, our small California Internet company could employ an emergency power supply to mitigate its power outage threat and have all its data backed up daily on RAID tapes, which are stored at a remote site in case of an earthquake. The more preventative measures you establish upfront the better. Emerson says, "dollars spent in prevention are worth more than dollars spent in recovery." The results of Step 1 should be a comprehensive list of possible threats, each with its corresponding solution and cost. It is imperative that IT presents all of these threats to the business operations units, so they can make an informed decision regarding the size of the disaster recovery budget (i.e., which risks the company can afford to tolerate and which it must pay to mitigate). Emerson believes IT "falls down" in its failure to communicate the real risks for system downtime to the business operations units of their companies. He says, "It's okay for operations to say no; it's not okay for IT not to let them know the risks."

A good place to begin is by presenting the cost of downtime to the business. How long can your business afford to be without its computer systems should one of your threats occur? Ultimately, the business operations unit decides which threats the business can tolerate. According to Emerson, when developing a DRP, IT departments are "shooting in the dark without those business indications." Both IT and the business units must agree on which data and applications are most critical to the business and need to be recovered most quickly in a disaster. The management of our small Internet company, for example, may decide they can supply the budget only for the emergency generators and the company will have to assume the risk of an earthquake. Disaster recovery budgets vary from company to company but they typically run between 2 and 8 percent of the overall IT budget. Companies for which system availability is crucial usually are on the higher end of the scale, while companies that can function without it are on the lower end. However, these percentages may be too small. For a large IT shop 15 percent is a best practice rule of thumb according to Emerson.

3. Develop the Plan

The feedback from the business units will begin to shape your DRP procedures. If, for example, they determine that the company must be up within 48 hours of an incident to stay viable, then you can calculate the amount of time it would take to execute the recovery plan and have the business back up in that timeframe. Emerson suggests that you have the recovery systems tested, configured, and retested 24 hours prior to launching them. He says the set up takes anywhere from 40 hours to days to complete. The recovery procedure should be written in a detailed plan or "script." Establish a

Recovery Team from among the IT staff and assign specific recovery duties to each member. The manner in which your team conducts its recovery probably will be no different than its regular production procedures: the chain of command likely won't change and neither will the aspects of the network for which each member is responsible. Define how to deal with the loss of various aspects of the network (databases, servers, bridges/routers, communications links, etc.) and specify who arranges for repairs or reconstruction and how the data recovery process occurs. The script will also outline priorities for the recovery: What needs to be recovered first? What is the communication procedure for the initial respondents? To complement the script, create a checklist or test procedure to verify that everything is back to normal once repairs and data recovery have taken place.

4. Procedure

A disaster recovery plan (DRP) - sometimes referred to as a business continuity plan (BCP) or business process contingency plan (BPCP) - describes how an organization is to deal with potential disasters. Just as a disaster is an event that makes the continuation of normal functions impossible, a disaster recovery plan consists of the precautions taken so that the effects of a disaster will be minimized and the organization will be able to either maintain or quickly resume mission-critical functions. Typically, disaster recovery planning involves an analysis of business processes and continuity needs; it may also include a significant focus on disaster prevention. Disaster recovery is becoming an increasingly important aspect of enterprise computing. As devices, systems, and networks become ever more complex, there are simply more things that can go wrong. As a consequence, recovery plans have also become more complex. For example, fifteen or twenty years ago if there was a threat to systems from a fire, a disaster recovery plan might consist of powering down the mainframe and other computers before the sprinkler system came on, disassembling components, and subsequently drying circuit boards in the parking lot with a hair dryer. Current enterprise systems tend to be too large and complicated for such simple and hands-on approaches, however, and interruption of service or loss of data can have serious financial impact, whether directly or through loss of customer confidence.

Appropriate plans vary from one enterprise to another, depending on variables such as the type of business, the processes involved, and the level of security needed. Disaster recovery planning may be developed within an organization or purchased as a software application or a service. It is not unusual for an enterprise to spend 25% of its information technology budget on disaster recovery. Nevertheless, the consensus within the DR industry is that most enterprises are still ill-prepared for a disaster. According to the Disaster Recovery site, "Despite the number of very public disasters since 9/11, still only about 50 percent of companies report having a disaster recovery plan. Of those that do, nearly half have never tested their plan, which is tantamount to not having one at all."

5. Recovery point objective

Recovery Point Objective (RPO) describes the acceptable amount of data loss measured in time. The Recovery Point Objective (RPO) is the point in time to which you must recover data as defined by your organization. This is generally a definition of what an organization determines is an "acceptable loss" in a disaster situation. If the RPO of a company is 2 hours and the time it takes to get the data back into production is 5 hours, the RPO is still 2 hours. Based on this RPO the data must be restored to within 2 hours of the disaster. Example: If there is a complete replication at 10:00am and the system dies at 11:59am without a new replication, the loss of the data written between 10:00am and 11:59am will not be recovered from the replica. This amount of time data has been lost has been deemed acceptable because of the 2 hour RPO. This is the case even if it takes an additional 3 hours to get the site back into production. The production will continue from the point in time of 10:00am. All data in between will have to be manually recovered through other means. The RPO in conjunction with the Recovery Time Objective (RTO) is the basis on which data protection strategy is developed.

6. Recovery time objective

The Recovery Time Objective (RTO) is the duration of time and a service level within which a business process must be restored after a disaster (or disruption) in order to avoid unacceptable consequences associated with a break in business continuity. The business continuity timeline usually runs parallel with an incident management timeline and may start at the same, or different, points. In accepted business continuity planning methodology the RTO is established during the Business Impact Analysis (BIA) by the owner of a process (usually in conjunction with the Business Continuity planner). The RTOs are then presented to senior management for acceptance. It should be noted that the

RTO attaches to the business process and not the resources required to support the process.

The RTO and the results of the BIA in its entirety provide the basis for identifying and analyzing viable strategies for inclusion in the business continuity plan. Viable strategy options would include any which would enable resumption of a business process in a time frame at or near the RTO. This would include alternate or manual workaround procedures and would not necessarily require computer systems to meet the RTOs. It is important to remember that the "O" in RTO stands for objective, not mandate. In reality, strategy is often selected that will not meet the RTO. In this instance the RTO will not be met but should still remain an objective of future strategy revision.

In Section 4 of this course you will cover these topics:

- Developing Procedures For Special Circumstances
- Testing The Disaster Recovery Plan

Topic : Developing Procedures For Special Circumstances

Topic Objective:

At the end of this topic student would be able to:

- Explain Business Continuance
- Highlight Cleanup Taskforce
- Analyze the concept of Test
- Elaborate Remote backup service
- Evaluate Secure virtual office
- Define Seven tiers of disaster recovery
- Describe Virtual tape library

- Examine Emulation
- Identify Address space remapping

Definition/Overview:

Network disaster recovery plans (DRPs) are just emerging as a business issue. The need to make contingency plans in case of a computer disaster, or major disruptions to the network, is a vital part of corporate strategy. Identifies users as responsible directly for the problem of security and for establishing the levels of security; whereas IT management staff must be responsible for the availability of IT. Highlights what measures can be taken to set up network DRPs, although advises against too complex strategies and plans, and emphasizes easy-to-maintain strategies. Concludes that a corporate policy must be developed to accommodate all the changes needed in selecting an appropriate DRP.

Key Points:

1. Business Continuance

Business continuance (sometimes referred to as business continuity) describes the processes and procedures an organization puts in place to ensure that essential functions can continue during and after a disaster. Business continuance planning seeks to prevent interruption of mission-critical services, and to reestablish full functioning as swiftly and smoothly as possible.

2. Cleanup Taskforce

This Disaster Recovery Plan (DRP) can be used as a Disaster Planning template for any size of enterprise. The Disaster Recovery template and supporting material have been updated to be Sarbanes-Oxley and HIPAA compliant. Preparation for Disaster Recovery / Business Continuity in light of SOX has two primary parts. The first is putting systems in place to completely protect all financial and other data required to meet the reporting regulations and to archive the data to meet future requests for clarification of those reports. The second is to clearly and expressly document all these procedures so that in the event of a SOX audit, the auditors clearly see that the DRP exists and will appropriately protect the data.

3. Test

Once your DRP is set, test it frequently. Eventually you'll need to perform a component-level restoration of your largest databases to get a realistic assessment of your recovery procedure, but a periodic walk-through of the procedure with the Recovery Team will assure that everyone knows their roles. Test the systems you're going to use in recovery regularly to validate that all the pieces work. Always record your test results and update the DRP to address any shortcomings. As your business environment changes, so should your DRP. Reexamine the plan every year on a high level: Do you still need every part of the plan? Do you need to add to it? Will the budget need to be adjusted to accommodate changes to the plan? As applications, hardware, and software are added to your network, they must be brought into the plan. New employees must be trained on recovery procedures. New threats to business seem to pop up every week and a sound DRP takes all of them into account.

4. Remote backup service

A remote, online, or managed backup service is a service that provides users with an online system for backing up and storing computer files. Managed backup providers are companies that provide this type of service. Online backup systems are typically built around a client software program that runs on a schedule, typically once a day. This program collects, compresses, encrypts, and transfers the data to the remote backup service provider's servers. Other types of product are also available in the market, such as remote continuous data protection (CDP). Providers of this type of service frequently target specific market segments. High-end LAN-based backup systems may offer services such as near-realtime transaction-level replication or open file backups. Consumer online backup companies frequently have beta software offerings and/or free-trial backup services. Most online/remote backup services came into existence during the heyday of the dotcom boom in the late 1990s with the exception of a few early pioneers like industry originator Rob Cosgrove, CEO of Remote Backup Systems. While the initial years of these service providers were about capturing market share distributed among the top few providers, the large industry players took cognizance of the importance and the role that these online backup providers were playing in the web services arena and M&A activity has became quite predominant in the last few years. Today, most service providers of online backup services position their services using the SaaS (software as a service) strategy and its relevance is predicted to increase exponentially in the years to come as personal and enterprise data storage needs rise. The last few years have also witnessed a healthy rise in the number of online backup providers with them existing independently as also as part of a business unit of a larger industry behemoth.

5. Secure virtual office

A Secure Virtual Office is a software environment which allows people to securely access and run applications on a remote server from an Internet connection, as if the application and data was on their own desktop machine. This allows client machines to be cheap personal computers or a thin client. IT administration costs are significantly lower as just a few large servers need to be maintained rather than the many PCs found in a conventional office. A Secure Virtual Office facility often includes standard office productivity applications like Microsoft Office or OpenOffice.org as well as access to useful software like accountancy or CRM packages. Secure Virtual Offices are usually found located in a data centre in large cities around the world. A modern Secure Virtual Office shares many similarities to older client-server architectures. But the recent re-emergence is a far cry from the dumb terminal green screens of the past. The recent rise in popularity has been brought about by fast, secure internet connections, cheap, reliable servers and software vendors providing applications that are application service provider (ASP) ready.

6. Seven tiers of disaster recovery

The Seven Tiers of Disaster Recovery was originally defined by Share to help identify the various methods of recovering mission-critical computer systems as required to support business continuity. Although the original known published concept dates back to the 1990s, Business Continuity Planning (BCP) and Disaster Recovery Specialists today continue to use the 7-Tiers to illustrate continuity capabilities and costs at a very high level. The definitions for the various Tiers have been updated as technology has evolved in support of today's business requirements and their associated Recovery Time Objective (RTO) and Recovery Point Objective (RPO). The seven tiers of business continuity solutions offer a simple method to define current service levels and associated risks.

- **Tier 0: No off-site data Possibly no recovery**

  Businesses with a Tier 0 business continuity solution have no business continuity plan. There is no saved information, no documentation, no backup hardware, and no contingency plan. The time necessary to recover in this instance is unpredictable. In fact, it may not be possible to recover at all.

- **Tier 1: Data backup with no hot site**

  Businesses that use Tier 1 continuity solutions back up their data and send these backups to an off-site storage facility. The method of transporting these backups is often referred to as "PTAM" - the "Pick-up Truck Access Method." Depending on how often backups are created and shipped, these organizations must be prepared to accept several days to weeks of data loss, but their backups are secure off-site. However, this tier lacks the systems on which to restore data.

- **Tier 2: Data backup with a hot site**

  Businesses using Tier 2 business continuity solutions make regular backups on tape. This is combined with an off-site facility and infrastructure (known as a hot site) in which to restore systems from those tapes in the event of a disaster. This solution will still result in the need to recreate several hours or even days worth of data, but the recovery time is more predictable.

- **Tier 3: Electronic vaulting**

  Tier 3 solutions build on the components of Tier 2. Additionally, some mission critical data is electronically vaulted. This electronically vaulted data is typically more current than that which is shipped via PTAM. As a result there is less data recreation or loss after a disaster occurs. The facilities for providing Electronic Remote Vaulting consists of high-speed communication circuits, some form of channel extension equipment and either physical or virtual Tape devices and an automated tape library at the remote site. IBM's Peer-to-Peer VTS and Sun's VSM Clustering are two examples of this type implementation.

- **Tier 4: Point-in-time copies**

  Tier 4 solutions are used by businesses that require both greater data currency and faster recovery than users of lower tiers. Rather than relying largely on shipping tape, as is common on the lower tiers, Tier 4 solutions begin to incorporate more disk based solutions. Several hours of data loss is still possible, but it is easier to make such point-in-time (PiT) copies with greater frequency than tape backups even when electronically vaulted.

- **Tier 5: Transaction integrity**

  Tier 5 solutions are used by businesses with a requirement for consistency of data between the production and recovery data centers. There is little to no data loss in such solutions, however, the presence of this functionality is entirely dependent on the application in use.

- **Tier 6: Zero or near-Zero data loss**

  Tier 6 business continuity solutions maintain the highest levels of data currency. They are used by businesses with little or no tolerance for data loss and who need to restore data to applications rapidly. These solutions have no dependence on the applications or applications staffs to provide data consistency. Tier 6 solutions often require some form of Disk mirroring. There are various synchronous and asynchronous solutions available from the mainframe storage vendors. Each solution is somewhat different, offering different capabilities and providing different Recovery Point and Recovery Time objectives. Often some form of automated tape solution is also required. However, this can vary somewhat depending on the amount and type of data residing on tape.

- **Tier 7: Highly automated, business integrated solution**

  Tier 7 solutions include all the major components being used for a Tier 6 solution with the additional integration of automation. This allows a Tier 7 solution to ensure consistency of data above that which is granted by Tier 6 solutions. Additionally, recovery of the applications is automated, allowing for restoration of systems and applications much faster and more reliably than would be possible through manual business continuity procedures.

7. Virtual tape library

A virtual tape library (VTL) is a data storage virtualization technology used typically for backup and recovery purposes. A VTL presents a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software. Virtualizing the disk storage as tape allows integration of VTLs with existing backup software and existing backup and recovery processes and policies. The benefits of such virtualization include storage consolidation and faster data restore processes. Most current VTL solutions use PATA or SATA disk arrays as the primary storage component due to their relatively low cost. The use of array enclosures increases the scalability of the solution by allowing the addition of more disk drives and enclosures to increase the storage capacity. The shift to VTL also eliminates streaming problems that often impair efficiency in tape drives as disk technology does not rely on streaming and can write effectively regardless of data transfer speeds.

By backing up data to disks instead of tapes, VTL often increases performance of both backup and recovery operations. Restore processes are found to be faster than backup regardless of implementations. In some cases, the data stored on the VTL's disk array is exported to other media, such as physical tapes, for disaster recovery purposes (scheme called disk-to-disk-to-tape, or D2D2T). Alternatively, most contemporary backup software products introduce also direct usage of the file system storage (especially network-attached storage, accessed through NFS and CIFS protocols over IP networks) not requiring a tape library emulation at all. They also often offer a disk staging feature: moving the data from disk to a physical tape for a long-term storage.

8. Emulation

The word emulation refers to an ambition and effort to equal, excel or surpass another; to compete or rival with some degree of success, especially through imitation. It can also refer to the simulation of equipment or phenomena by artificial means, such as by software modeling. Emulation or Emulator may also refer to:

- **Emulator**, imitation of behavior of a computer or other electronic system with the help of another type of computer/system
- **Console emulator**, a program that allows a computer or modern console to emulate another video game console
- **In-circuit emulator**, a program used to emulate the processor in an embedded system, to aid in debugging
- **Hardware emulation**, the use of special purpose hardware to emulate the behavior of a yet-to-be-built system, with greater speed than pure software emulation
- **Emulation for Logic Validation**, used to emulate hardware in manufacturing automation

9. Address space remapping

Virtualization of storage helps achieve location independence by abstracting the physical location of the data. The virtualization system presents to the user a logical space for data storage and itself handles the process of mapping it to the actual physical location. The actual form of the mapping will depend on the chosen implementation. Some implementations may limit the granularity of the mapping which itself may limit the capabilities of the device. Typical granularities range from a single physical disk down to some small subset (multiples of megabytes or gigabytes) of the physical disk. In a block-based storage environment, a single block of information is addressed using a logical unit identifier (LUN) and an offset within that LUN - known as a Logical Block Address (LBA). The address space mapping is between a logical disk, usually referred to as a virtual disk (vdisk) and a logical unit presented by one or more storage controllers. The LUN itself may be also a product of virtualization in a different layer.

Topic : Testing The Disaster Recovery Plan

Topic Objective:

At the end of this topic student would be able to:

- Explain Network simulation for disaster recovery plan testing

- Analyze Lack of disaster recovery plan testing
- Highlight Network simulation technology
- Elaborate Distance testing for DR sites
- Evaluate Vendor testing
- Identify User testing and validation
- Define Testing options

Definition/Overview:

Disaster recovery is the process, policies and procedures related to preparing for recovery or continuation of technology infrastructure critical to an organization after a natural or human-induced disaster. Disaster recovery planning is a subset of a larger process known as business continuity planning and should include planning for resumption of applications, data, hardware, communications (such as networking) and other IT infrastructure. A business continuity plan (BCP) includes planning for non-IT related aspects such as key personnel, facilities, crisis communication and reputation protection, and should refer to the disaster recovery plan (DRP) for IT related infrastructure recovery / continuity. This article focuses on disaster recovery planning as related to IT infrastructure.

Key Points:

1. Network simulation for disaster recovery plan testing

IT disaster recovery (DR) has long been a primary focus of storage resellers. One of the main selling points for backup storage is that it helps businesses avoid interruption of critical operations in the event of a fire, power outage, hurricane or other disaster. Recovering from disaster, however, is about more than protecting data. To continue functioning, a business may require secondary offsite servers, alternative network connections and the ability to relocate core personnel. Anyone selling storage as part of an overall DR strategy must look at the bigger picture as well. Network simulation technology can help you test clients' disaster recovery plans and ensure that they'll actually work. IT disaster recovery (DR) has long been a primary focus of storage resellers. One of the main selling points for backup storage is that it helps businesses avoid interruption of critical operations in the event of a fire, power outage, hurricane or other disaster. Recovering from disaster, however, is about more than protecting data. To continue functioning, a business may require secondary offsite servers, alternative network connections and the ability to relocate core personnel. Anyone selling storage as part of an overall DR strategy must look at the bigger picture as well.

2. Lack of disaster recovery plan testing

Unfortunately, most clients -- and most resellers -- are never totally sure that their DR plans will work when the time comes. They may consider multiple disaster scenarios. They may invest in what they hope will be adequate DR infrastructure. They may even test that infrastructure by having technicians access failover servers via failover network links. But they never fully validate their disaster recovery plan investments. The need for such validation should not be underestimated. Say, for example, that your client replicates a critical server and locates it a safe distance from the primary data center. A technician tests the performance of this replicated server and it seems fine. So they believe they are safe. But they're not. When end users access a replicated server that is physically farther from them than the primary server, their application sessions can take longer. These longer sessions can translate directly into reduced capacity. As a result the end user may experience unacceptable application performance during an emergency, despite the fact that everything looked OK on paper and with a single user.

3. **Network simulation technology**

This technology allows them to fully replicate the network conditions that would exist in the event of a real-world disaster, including all the characteristics of the network (such as latency and bandwidth constraints) and realistic traffic levels for all the applications needed to support the continuity of the business. They can then ensure that the DR infrastructure they put in place will in fact deliver required levels of service to all necessary end users in all locations. They will also be able to "right size" their DR investments, and demonstrate to auditors and regulators that they have exercised appropriate diligence in formulating and executing their DR plans. Of course, network simulation technology has many other important uses. It's helpful in capacity planning, evaluating new products and helping developers design applications that perform well under real-world network conditions. For storage resellers, its DR applicability makes it particularly appealing. Extreme weather and concerns about terrorism have made business leaders more sensitive than ever about business continuity. You can address those sensitivities by delivering the right DR infrastructure solutions and the right technology for validating your clients' DR plans. As businesses become increasingly dependent on their IT systems and data, a DR plan is essential to their survival and success--both for business and regulatory reasons. Most businesses today rely on key systems, networks, and data for every aspect of their operations. If these go down (or their data becomes unavailable), the company is effectively out of business. So the key question in devising a DR plan is How long can we afford to be out of business, and what kind of systems and processes do we need to implement to achieve this level of availability? On the regulatory side, the requirements of Sarbanes-Oxley, Gramm-Leach-Bliley, HIPAA, FISMA [Federal Information Security Management Act], ISO 17799, and other regulations have raised the bar for data protection, security, and availability--and increased the need for DR capabilities and effective testing.

4. Distance testing for DR sites

The new guidelines are also causing companies to plan, build and test their secondary DR sites at greater distances. Some federal agencies are suggesting a minimum 200-to-300-mile separation of facilities, and many enterprises prefer up to 1,000 km or more. As a result, one area getting more attention is the effect of distance in the DR/backup network connection. As distance increases, the effects of latency increase, along with bit error rates and, depending on the underlying technologies, potential packet loss and packet-sequencing issues. Some users are extremely sophisticated in their testing process, running full emulation testing of their backup/DR procedures. This is accomplished using a simulated link between their primary and secondary sites. In fact, some leverage this testing environment even further and use it to evaluate various hardware and software technologies and products to be used or purchased for DR.

5. Vendor testing

Some users assume that this type of testing is conducted by the equipment manufacturers and carriers. Given all the talk in the past on interoperability testing, it would be reasonable to believe that vendor interoperability labs do it all. For some vendors and carriers, this is definitely the case. Clearly, there are certain storage manufacturers, switch manufacturers (both Fibre Channel and IP) and WAN carriers that have solid plans and tools in place for testing the effects of distance and error injection on their products in various multi-vendor configurations. In fact, I've been impressed with the level of commitment and detail I've heard from some of them. Others will say that they have a good plan, but they won't talk about it. Unfortunately, based on stories from various users, it seems apparent that some vendors have not been as vigilant in their testing. One switch vendor that will go unnamed provides a good model of how to do it right. As a vendor with high quality standards, its quality assurance group is responsible for long-distance testing of the company's Fibre Channel switches. The group's goal is to test over distances ranging from 100 to 200 km (with plans to go up to 1,000 km), handling the associated delays, error rates and potential fiber cuts.

Initially the company's testing was done with a dark fiber spindle in a box. As the distance increased, this became both more cumbersome (dealing with all the necessary equipment) and more expensive (including the cost of the fiber and all the optics needed for repeating and amplifying). Today, the company's lab environment uses a network emulator product from Anue Systems, allowing it to emulate whatever distance and delay it chooses for any test, using a programmable box. Network emulators for various protocols are available from a variety of vendors, including Agilent, Anue Systems, Empirix, PacketStorm Communications, Shunra Software, Simena, Spirent Communications and others. According to the switch vendor's QA manager, "The Anue box was small, stable and easy to use, and the support was excellent. As a result, the space reduction and cost savings were huge." The vendor has now automated much of the testing using a command line interface. Currently it's testing 1G and 2G Fibre Channel, but it will be moving to 4G and 10G Fibre Channel as these products roll out.

6. User testing and validation

While we hope that all equipment manufacturers are doing this type of testing, users need to own the responsibility for testing and validating in their own environment before production deployment (as well as on an ongoing basis). One major insurance provider I spoke with, as part of a new and improved DR plan, is in the process of establishing a full-backup, secondary data center 1,400 miles away. Because it couldn't add test traffic to its 24x7 production network, the company needed a test environment to simulate the distance to test its configuration of servers, storage arrays, tape, LANs, SANs and WANs. The company's testing required Gigabit Ethernet and SONET running OC-48, moving to OC-192 and 10G Ethernet. When asked about his criteria for a distance simulation tool, the design architect responsible for the DR plan and test lab pointed first to ease of use. "We like the Anue box because anybody can set it up and use it. Cost, functionality and ease of use what else is there?" The company has been using it for three months of testing, and he says it has "performed flawlessly." In addition to using Anue for validating its DR plan, the company is currently evaluating various technology options against each other and plans to use its test environment to do specific product evaluations, as well. Another large enterprise user already running multiple data centers is making some major network changes from channel extenders and SRDF to Fibre Channel arrays connected via FCIP over GigE or SONET, or Fibre Channel over ATM.

The company is evaluating various vendors, but it cannot run field trials over its production network. It therefore needs to mimic its production environment, including servers, arrays, tape libraries and Fibre Channel and IP network elements over distance. It uses its production environment data (a 700G-byte Oracle database and latency numbers from its live network) and plugs it into its test environment. Currently evaluating network emulator products from multiple vendors, the company is looking closely at features, ease of use and, of course, price.

7. Testing options

Some network emulator vendors only support certain technologies, while others have multi-protocol products (Fibre Channel, Ethernet, IP and SONET). Functionality and prices cover a broad range. Some products focus on basic required functions, simplicity and ease of use with a low price point, while other high-end products offer sophisticated modeling techniques, output to Visio diagrams, advanced technical features, extensive GUIs and the higher price tag to go with them.

In Section 5 of this course you will cover these topics:

- Continued Assessment Of Needs, Threats, And Solutions

Topic : Continued Assessment Of Needs, Threats, And Solutions

Topic Objective:

At the end of this topic student would be able to:

- Understand Analysis
- Explain Threat analysis
- Elaborate Definition of impact scenarios
- Define Recovery requirement documentation
- Describe Solution design
- Evaluate Implementation
- Learn about Recovering data after physical damage
- Know about Recovery techniques
- Examine Recovering data after logical damage
- Analyze Preventing logical damage
- Highlight Recovery techniques

Definition/Overview:

Business continuity planning: Business continuity planning (BCP) is the creation and validation of a practiced logistical plan for how an organization will recover and restore partially or completely interrupted critical (urgent) functions within a predetermined time after a disaster or extended disruption. The logistical plan is called a business continuity plan. In plain language, BCP is working out how to stay in business in the event of disaster. Incidents include local incidents like building fires, regional incidents like earthquakes, or national incidents like pandemic illnesses. BCP may be a part of an organizational learning effort that helps reduce operational risk associated with lax information management controls. This process may be integrated with improving information security and corporate reputation risk management practices.

In December 2006, the British Standards Institution (BSI) released a new independent standard for BCP BS 25999-1. Prior to the introduction of BS 25999, BCP professionals relied on BSI information security standard BS 7799, which only peripherally addressed BCP to improve an organization's information security compliance. BS 25999's applicability extends to organizations of all types, sizes, and missions whether governmental or private, profit or non-profit, large or small, or industry sector. In 2007, the BSI published the second part, BS 25999-2 "Specification for Business Continuity Management", that specifies requirements for implementing, operating and improving a documented Business Continuity Management System (BCMS). In 2004, the United Kingdom enacted the Civil Contingencies Act 2004, a statute that instructs all emergency services and local authorities to actively prepare and plan for emergencies. Local authorities also have the legal obligation under this act to actively lead promotion of business continuity practices amongst its geographical area.

Key Points:

1. Analysis

The analysis phase in the development of a BCP manual consists of an impact analysis, threat analysis, and impact scenarios with the resulting BCP plan requirement documentation. An impact analysis results in the differentiation between critical (urgent) and non-critical (non-urgent) organization functions/ activities. A function may be considered critical if the implications for stakeholders of damage to the organization resulting are regarded as unacceptable. Perceptions of the acceptability of disruption may be modified by the cost of establishing and maintaining appropriate business or technical recovery solutions. A function may also be considered critical if dictated by law. For each critical (in scope) function, two values are then assigned:

- **Recovery Point Objective** (RPO)- the acceptable latency of data that will be recovered
- **Recovery Time Objective** (RTO) - the acceptable amount of time to restore the function

The Recovery Point Objective must ensure that the Maximum Tolerable Data Loss for each activity is not exceeded. The Recovery Time Objective must ensure that the Maximum Tolerable Period of Disruption (MTPD) for each activity is not exceeded. Next, the impact analysis results in the recovery requirements for each critical function. Recovery requirements consist of the following information:

- The business requirements for recovery of the critical function, and/or
- The technical requirements for recovery of the critical function

2. Threat analysis

After defining recovery requirements, documenting potential threats is recommended to detail a specific disasters unique recovery steps. Some common threats include the following:

- Disease
- Earthquake
- Fire
- Flood
- Cyber attack
- Bribery
- Hurricane
- Utility outage
- Terrorism

All threats in the examples above share a common impact: the potential of damage to organizational infrastructure - except one (disease). The impact of diseases can be regarded as purely human, and may be alleviated with technical and business solutions. However, if the humans behind these recovery plans are also affected by the disease, then the process can fall down. During the 2002-2003 SARS outbreak, some organizations grouped staff into separate teams, and rotated the teams between the primary and secondary work sites, with a rotation frequency equal to the incubation period of the disease. The organizations also banned face-to-face contact between opposing team members during business and non-business hours. With such a split, organizations increased their resiliency against the threat of government-ordered quarantine measures if one person in a team contracted or was exposed to the disease. Damage from flooding also has a unique characteristic. If an office environment is flooded with non-salinated and contamination-free water (e.g., in the event of a pipe burst), equipment can be thoroughly dried and may still be functional.

3. Definition of impact scenarios

After defining potential threats, documenting the impact scenarios that form the basis of the business recovery plan is recommended. In general, planning for the most wide-reaching disaster or disturbance is preferable to planning for a smaller scale problem, as almost all smaller scale problems are partial elements of larger disasters. A typical impact scenario like 'Building Loss' will most likely encompass all critical business functions, and the worst potential outcome from any potential threat. A business continuity plan may also document additional impact scenarios if an organization has more than one building. Other more specific impact scenarios - for example a scenario for the temporary or permanent loss of a specific floor in a building - may also be documented.

4. Recovery requirement documentation

After the completion of the analysis phase, the business and technical plan requirements are documented in order to commence the implementation phase. A good asset management program can be of great assistance here and allow for quick identification of available and re-allocateable resources. For an office-based, IT intensive business, the plan requirements may cover the following elements which may be classed as ICE (In Case of Emergency) Data:

- The numbers and types of desks, whether dedicated or shared, required outside of the primary business location in the secondary location
- The individuals involved in the recovery effort along with their contact and technical details
- The applications and application data required from the secondary location desks for critical business functions
- The manual workaround solutions
- The maximum outage allowed for the applications
- The peripheral requirements like printers, copier, fax machine, calculators, paper, pens etc.

Other business environments, such as production, distribution, warehousing etc will need to cover these elements, but are likely to have additional issues to manage following a disruptive event.

5. Solution design

The goal of the solution design phase is to identify the most cost effective disaster recovery solution that meets two main requirements from the impact analysis stage. For IT applications, this is commonly expressed as:

- The minimum application and application data requirements
- The time frame in which the minimum application and application data must be available

Disaster recovery plans may also be required outside the IT applications domain, for example in preservation of information in hard copy format, or restoration of embedded technology in process plant. This BCP phase overlaps with Disaster recovery planning methodology. The solution phase determines:

- The crisis management command structure
- The location of a secondary work site (where necessary)
- Telecommunication architecture between primary and secondary work sites
- Data replication methodology between primary and secondary work sites
- The application and software required at the secondary work site, and
- The type of physical data requirements at the secondary work site.

6. Implementation

The implementation phase, quite simply, is the execution of the design elements identified in the solution design phase. Work package testing may take place during the implementation of the solution, however; work package testing does not take the place of organizational testing.

- **Testing and organizational acceptance**

The purpose of testing is to achieve organizational acceptance that the business continuity solution satisfies the organization's recovery requirements. Plans may fail to meet expectations due to insufficient or inaccurate recovery requirements, solution design flaws, or solution implementation errors. Testing may include:

- Crisis command team call-out testing
- Technical swing test from primary to secondary work locations
- Technical swing test from secondary to primary work locations
- Application test
- Business process test

At minimum, testing is generally conducted on a biannual or annual schedule. Problems identified in the initial testing phase may be rolled up into the maintenance phase and retested during the next test cycle.

- **Maintenance**

Maintenance of a BCP manual is broken down into three periodic activities. The first activity is the confirmation of information in the manual, roll out to ALL staff for awareness and specific training for individuals whose roles are identified as critical in response and recovery. The second activity is the testing and verification of technical solutions established for recovery operations. The third activity is the testing and verification of documented organization recovery procedures. A biannual or annual maintenance cycle is typical.

- **Information update and testing**

All organizations change over time, therefore a BCP manual must change to stay relevant to the organization. Once data accuracy is verified, normally a call tree test is conducted to evaluate the notification plan's efficiency as well as the accuracy of the contact data. Some types of changes that should be identified and updated in the manual include:

- Staffing changes
- Staffing persona
- Changes to important clients and their contact details
- Changes to important vendors/suppliers and their contact details
- Departmental changes like new, closed or fundamentally changed departments.
- Changes in company investment portfolio and mission statement
- Changes in upstream/downstream supplier routes

- **Testing and verification of technical solutions**

  As a part of ongoing maintenance, any specialized technical deployments must be checked for functionality. Some checks include:

  - Virus definition distribution
  - Application security and service patch distribution
  - Hardware operability check
  - Application operability check
  - Data verification

- **Testing and verification of organization recovery procedures**

  As work processes change over time, the previously documented organizational recovery procedures may no longer be suitable. Some checks include:

  - Are all work processes for critical functions documented?
  - Have the systems used in the execution of critical functions changed?
  - Are the documented work checklists meaningful and accurate for staff?
  - Do the documented work process recovery tasks and supporting disaster recovery infrastructure allow staff to recover within the predetermined recovery time objective.

- **Treatment of test failures**

  As suggested by the diagram included in this article, there is a direct relationship between the test and maintenance phases and the impact phase. When establishing a BCP manual and recovery infrastructure from scratch, issues found during the testing phase often must be reintroduced to the analysis phase.

7. Data recovery

Data recovery is the process of salvaging data from damaged, failed, corrupted, or inaccessible secondary storage media when it cannot be accessed normally. Often the data are being salvaged from storage media formats such as hard disk drives, storage tapes, CDs, DVDs, RAID, and other electronics. Recovery may be required due to physical damage to the storage device or logical damage to the file system that prevents it from being mounted by the host operating system. Although there is some confusion as to the term, data recovery can also be the process of retrieving and securing deleted information from a storage media for forensic purposes or spying.

8. Recovering data after physical damage

A wide variety of failures can cause physical damage to storage media. CD-ROMs can have their metallic substrate or dye layer scratched off; hard disks can suffer any of several mechanical failures, such as head crashes and failed motors; tapes can simply break. Physical damage always causes at least some data loss, and in many cases the logical structures of the file system are damaged as well. This causes logical damage that must be dealt with before any files can be salvaged from the failed media. Most physical damage cannot be repaired by end users. For example, opening a hard disk in a normal environment can allow airborne dust to settle on the platter and become caught between the platter and the read/write head, causing new head crashes that further damage the platter and thus compromise the recovery process. Furthermore, end users generally do not have the hardware or technical expertise required to make these repairs. Consequently, costly data recovery companies are often employed to salvage important data. These firms often use "Class 100" / ISO-5 cleanroom facilities to protect the media while repairs are being made. (Any data recovery firm without a pass certificate of IS0-5 or better will not be accepted by hard drive manufacturers for warranty purposes.)

9. Recovery techniques

Recovering data from physically-damaged hardware can involve multiple techniques. Some damage can be repaired by replacing parts in the hard disk. This alone may make the disk usable, but there may still be logical damage. A specialized disk-imaging procedure is used to recover every readable bit from the surface. Once this image is acquired and saved on a reliable medium, the image can be safely analyzed for logical damage and will possibly allow for much of the original file system to be reconstructed.

- **Hardware repair**

Media that has suffered a catastrophic electronic failure will require data recovery in order to salvage its contents. Examples of physical recovery procedures are: removing a damaged PCB (printed circuit board) and replacing it with a matching PCB from a healthy drive, performing a live PCB swap (in which the System Area of the HDD is damaged on the target drive which is then instead read from the donor drive, the PCB then disconnected while still under power and transferred to the target drive), read/write head assembly with matching parts from a healthy drive, removing the hard disk platters from the original damaged drive and installing them into a healthy drive, and often a combination of all of these procedures. Some data recovery companies have procedures that are highly technical in nature and are not recommended for an untrained individual. Any of them will almost certainly void the manufacturer's warranty.

- **Disk imaging**

Result of a failed data recovery from a Hard disk drive. The extracted raw image can be used to reconstruct usable data after any logical damage has been repaired. Once that is complete, the files may be in usable form although recovery is often incomplete. Open source tools such as DCFLdd v1.3.4-1 or DOS tools such as HDClone can usually recover data from all but the physically-damaged sectors. A 2007 Defense Cyber Crime Institute study shows that the DCFLdd v1.3.4-1 installed on a Linux 2.4 Kernel system produces extra "bad sectors", resulting in the loss of information that is actually available. The study states that when installed on a FreeBSD Kernel system, only the bad sectors are lost. Another tool that can correctly image damaged media is ILook IXImager, a tool available only to government and Law Enforcement.

Typically, Hard Disk Drive data recovery imaging has the following abilities: (1) Communicating with the hard drive by bypassing the BIOS and operating system which are very limited in their abilities to deal with drives that have "bad sectors" or take a long time to read. (2) Reading data from bad sectors rather than skipping them (by using various read commands and ECC to recreate damaged data). (3) Handling issues caused by unstable drives, such as resetting/repowering the drive when it stops responding or skipping sectors that take too long to read (read instability can be caused by minute mechanical wear and other issues). and (4) Pre-configuring drives by disabling certain features, such a SMART and G-List re-mapping, to minimize imaging time and the possibility of further drive degradation.

10. Recovering data after logical damage

Logical damage is primarily caused by power outages that prevent file system structures from being completely written to the storage medium, but problems with hardware (especially RAID controllers) and drivers, as well as system crashes, can have the same effect. The result is that the file system is left in an inconsistent state. This can cause a variety of problems, such as strange behavior (e.g., infinitely recursing directories, drives reporting negative amounts of free space), system crashes, or an actual loss of data. Various programs exist to correct these inconsistencies, and most operating systems come with at least a rudimentary repair tool for their native file systems. Linux, for instance, comes with the fsck utility, Mac OS X has Disk Utility and Microsoft Windows provides chkdsk. Third-party utilities such as The Coroners Toolkit and The Sleuth Kit are also available, and some can produce superior results by recovering data even when the disk cannot be recognized by the operating system's repair utility. Utilities such as TestDisk can be useful for reconstructing corrupted partition tables. Some kinds of logical damage can be mistakenly attributed to physical damage. For instance, when a hard drive's read/write head begins to click, most end-users will associate this with internal physical damage. This is not always the case, however. Another possibility is that the firmware of the drive or its controller needs to be rebuilt in order to make the data accessible again.

11. Preventing logical damage

The increased use of journaling file systems, such as NTFS 5.0, ext3, and XFS, is likely to reduce the incidence of logical damage. These file systems can always be "rolled back" to a consistent state, which means that the only data likely to be lost is what was in the drive's cache at the time of the system failure. However, regular system maintenance should still include the use of a consistency checker. This can protect both against bugs in the file system software and latent incompatibilities in the design of the storage hardware. One such incompatibility is the result of the disk controller reporting that file system structures have been saved to the disk when it has not actually occurred. This can often occur if the drive stores data in its write cache, then claims it has been written to the disk. If power is lost, and this data contains file system structures, the file system may be left in an inconsistent state such that the journal itself is damaged or incomplete. One solution to this problem is to use hardware that does not report data as written until it actually is written. Another is using disk controllers equipped with a battery backup so that the waiting data can be written when power is restored. Finally, the entire system can be equipped with a battery backup that may make it possible to keep the system on in such situations, or at least to give enough time to shut down properly.

12. Recovery techniques

Two common techniques used to recover data from logical damage are consistency checking and data carving. While most logical damage can be either repaired or worked around using these two techniques, data recovery software can never guarantee that no data loss will occur. For instance, in the FAT file system, when two files claim to share the same allocation unit ("cross-linked"), data loss for one of the files is essentially guaranteed.

12.1 Consistency checking

The first, consistency checking, involves scanning the logical structure of the disk and checking to make sure that it is consistent with its specification. For instance, in most file systems, a directory must have at least two entries: a dot (.) entry that points to itself, and a dot-dot (..) entry that points to its parent. A file system repair program can read each directory and make sure that these entries exist and point to the correct directories. If they do not, an error message can be printed and the problem corrected. Both chkdsk and fsck work in this fashion. This strategy suffers from two major problems. First, if the file system is sufficiently damaged, the consistency check can fail completely. In this case, the repair program may crash trying to deal with the mangled input, or it may not recognize the drive as having a valid file system at all. The second issue that arises is the disregard for data files. If chkdsk finds a data file to be out of place or unexplainable, it may delete the file without asking. This is done so that the operating system may run smoother, but the files deleted are often important user files which cannot be replaced. Similar issues arise when using system restore disks (often provided with proprietary systems like Dell and Compaq), which restore the operating system by removing the previous installation. This problem can often be avoided by installing the operating system on a separate partition from your user data.

12. Data carving

Data Carving is a data recovery technique that allows for data with no file system allocation information to be extracted by identifying sectors and clusters belonging to the file. Data carving usually searches through raw sectors looking for specific desired file signatures. The fact that there is no allocation information means that the investigator must specify a block size of data to carve out upon finding a matching file signature. This presents the challenge that the beginning of the file is still present and that there is (depending on how common the file signature is) a risk of many false hits. Also, data carving requires that the files recovered be located in sequential sectors (rather than fragmented) as there is no allocation information to point to fragmented file portions. This method can be time and resource intensive.