5.4 Disaster Recovery

Objectives:

- Develop effective disaster recovery plans
- Use IT contingency planning to eliminate single points of failure and provide fault tolerance and redundancy

Disaster Recovery Planning

Business continuity means planning systems and operations so that they are as little affected by incidents as possible and so that the resources are available to recover from them. Disaster recovery plans describe the procedures to follow in the event of severe incidents. When planning a network, the administrator must consider the impact of data loss and service unavailability on the organization. If a local plumber loses his email connection for a day, then the chances are he or she will lose little business because of it and he will still be able to function in his primary role of plumber (conversely, losing a cell phone could be catastrophic). Therefore little consideration would be needed to plan for disasters. Conversely, if a major bank lost its trading floor link to its partners, even for an hour, since the organization's primary function (trading) would be impossible, huge losses may result. Significant thought is required to ensure that a disaster does not cause this kind of loss to what are termed mission critical systems.

A disaster could be anything from a fairly trivial loss of power, or failure of a minor component, through to man-made or natural disasters, such as fires, earthquakes, or acts of terrorism.

Creating Disaster Recovery Plans

An organization sensitive to these risks will develop an effective, documented Disaster Recovery Plan (DRP). This should accomplish the following:

- Identify scenarios for natural and non-natural disaster and options for protecting systems. Plans need to take account of risk (a combination of the likelihood the disaster will occur and the possible impact on the organization) and cost. There is no point implementing disaster recovery plans that financially cripple the organization. The business case is made by comparing the cost of recovery measures against the cost of downtime. Downtime cost is calculated from lost revenues and ongoing costs (principally salary). The recovery plan should not generally exceed the downtime cost. Of course, downtime will include indefinable costs, such as loss of customer goodwill restitution for not meeting service contracts, and so on. For each disaster scenario. A Business Impact Analysis is made identifying risks and vulnerabilities and advising on appropriate countermeasures.

- Identify tasks, resources, and responsibilities for responding to a disaster.
    - Who is responsible for doing what? How can they be contacted? What happens if they are not available?
    - Which functions are most critical? Where should effort be concentrated first?
    - What resources are available? Should they be pre-purchased and held as stock? Will the disaster affect availability of supplies?
    - What are the timescales for resuming normal operations?

- Train Staff in the disaster planning procedures and how to react well to change.

Succession Planning

As well as risks to systems, a BIA has to take on the macabre issue of human capital resilience. Put bluntly, this means "is someone else available to fulfil the same role if an employee is incapacitated?"

Succession planning targets the specific issue of leadership and senior management. Most disaster recovery plans are heavily dependent on a few key people to take charge during the disaster and ensure that the plan is put into effect. Succession planning ensures that these sorts of competencies are widely available to an organization.

Testing Disaster Recovery Plans

It is necessary to test disaster recovery procedures. There are two means of doing this:

- Table top exercise - staff "ghost" the same procedures as they would in a disaster, without actually creating disaster conditions or applying or changing anything. These are simple to set up but do not provide any sort of practical evidence of things that could go wrong, time to complete and so on.
- Disaster recovery exercises designed to simulate different scenarios. These are extremely complex to set up, especially if run in the production environment rather than a test environment.

Also identify timescales for disaster plans to be , to take account of changing circumstances and business needs. Following an incident, it is vital to hold a review meeting to analyse why the incident occurred, what could have been done to prevent it, and how effective was the response.

As well as restoring systems, the disaster recovery plan should identify stakeholders who need to be informed about any security incidents. There may be a legal requirement to inform the police or fire service or buildings inspectors about any safety-related or criminal incidents. If third-party or personal data is lost or stolen, the data subjects may need to be informed. If the disaster affects services, customers need to be informed about the time-to-fix and any alternative arrangements that can be made.

Secure Recovery

When a disaster has occurred, then the recovery plan will swing into action to get the failed part of the network operational as soon as possible. If a disk has failed, swap it out. If a node in a cluster has failed, remove and replace or repair the node to provide for high reliability as soon as possible. If data becomes corrupted or lost, utilize your restore plan to recover the data.

With the pressure to get systems running again though, it is important not to overlook the process of re-securing the system against intrusion. The system needs to be restored to its baseline secure configuration.

Another issue is that creating a duplicate of anything doubles the complexity of securing that resource properly. The same security procedures must apply to redundant sites, spare systems, and backup data as applies to the main copy.

IT Contingency Planning

Computer systems require protection from hardware failure, software failure, and system failure (failure of network connectivity devices for instance).

When implementing a network, the goal will always be to minimize the single points of failure and to allow ongoing service provision despite a disaster. To perform IT Contingency Planning (ITCP), think of all the things that could fail, determine whether the result would be critical loss of service, and whether this is unacceptable. Then work out how to make the service fault tolerant.

Fault Tolerance

A system that can experience failures and continue to provide the same (or nearly the same) level of service is said to be fault tolerant. Fault tolerance is often achieved by provisioning redundant components. A redundant component is one that is not essential to the normal function of a system but that allows the system to recover from the failure of another component.

Examples of devices and solutions that provide fault tolerance and high availability include the following:

- Redundant components (power supplies, network cards, drives (RAID), and cooling fans) provide protection against hardware failures. Hot swappable components allow for easy replacement (without having to shut down the server).
- Uninterruptible Power Supplies (UPS) and Standby Power Supplies.
- Backup strategies - provide protection for data.
- Cluster services are an (expensive) way of ensuring that the total failure of a server does not disrupt services generally.

While these computer systems are important, thought also needs to be given to how to make the business "fault tolerant" in terms of staffing, utilities (heat, power, communications, and transport), customers, and suppliers.

The following components are often provided with redundancy on server systems.

- Load Balancing Network Links - Without a network connection, a server is not of much use! As network cards are cheap, it is commonplace for a server to have multiple cards (adapter fault tolerance). Multiple adapters can be configured to work together (adapter teaming). This provides fault tolerance (if one adapter fails, the network connection will not be lost) and can also provide load balancing (connections can be spread between the cards). Network cabling should be designed to allow for multiple paths between the various servers, so that during a failure of one part of the network, the rest remains operational (redundant connections). Routers are great fault tolerant devices, because they can communicate system failures and IP packets can be routed via an alternate device.

- Power Supply – Power supplies are one of the components more likely to break down than others. They are cheap, but they take up a lot of space. A second power supply can be configured to take over from the first if it fails. A redundant power supply unit does not protect against problems with the supply of power itself. For this you need a Universal Power Supply and/or backup generator. A UPS is always required to provide against interruption to computer services. A backup generator cannot be brought online fast enough to respond to a power failure. The issue with a

UPS is that it is battery based and therefore only able to supply power for a set amount of time which ranges from a few minutes to a few hours depending on the size of the battery bank and the service load. The UPS only provides power to critical server components.
A backup generator can provide power to the whole building and uses fuel such as diesel which can keep it going for several days. The main issue is the safe storage of the fuel.

- Cooling Fans – Servers contain many heat generating components (notably processors, hard drives and power supplies), efficient cooling systems are particularly important. Fans have are more likely to fail over other components and therefore you should plan for redundancy.

- Hardware and Spare Parts – Within reason an inventory of spare parts should be kept. However as parts become obsolete quite quickly care should be taken not to overstock on parts.
- Hard Disks – These store the most valuable components of a computer system. Data can be protected by making back-ups, but only very expensive backup solutions can make up to the minute copies. Multiple disks can be configured as a RAID solution.
- Drive Arrays (RAID) – RAID (Redundant Array of Independent Disks), many disks can act as backups for each other to increase reliability and fault tolerance. If one disk fails, the data is not lost and the server can keep functioning. RAID levels are numbered 1 to 6 where each level corresponds to a specific type of fault tolerance.

RAID Level	Fault Tolerance
Level 0	Striping without Parity (no Fault tolerance) Data is written in blocks across multiple disks simultaneously. This can improve performance, but if one disk fails, so does the whole volume and data on the disk will be corrupted
Level 1	Mirroring – data is written to 2 disks simultaneously, provides redundancy (if one fails there is a copy on the other disk) However, storage efficiency is only 50%
Level 5	Striping with Parity – data is written across three or more disks but additionally parity is calculated. This allows the volume to continue if one disk is lost. This has better storage efficiency than RAID 1
Nested (0+1, 1+0 or 5+0)	Nesting RAID sets generally improves performance or redundancy as nested solutions can support the failure of more than one disk.

Clusters and Sites

As well as providing redundancy at device level, whole servers and even sites can be provisioned:

Clusters – This is a group of servers each of which is called a node. A cluster provides fault tolerance for critical applications. They can do this by taking over the position of a failed node should a problem occur. There are 2 types of clustering, Active/Active and Active/Passive.

- Active/Active – consists of a number of nodes that process concurrently. This allows for the maximum capacity from the available hardware while all nodes are functional. In the event of a failover the workload of the failed node is distributed to the remaining nodes. The problem is that during a failover the performance of the remaining nodes is affected.

- Active/Passive – These use a redundant node to failover. So if a cluster had 8 nodes, the 8th node would not do anything until a failover occurs. It then takes over responsibility for the failed node's services. The advantage of this is that performance is not affected during failover. However, the hardware and operating system costs are higher due to unused capacity.

Some applications and services will not function in a clustered environment. This should e taken account of if you decide to operate a cluster service.

Uses for Clustering

Generally provides a level of fault tolerance for back-end applications. So if you had an online purchasing system you might want to implement a cluster service to support the SQL database.

At the front-end you might to deploy a network load balancer for web servers which in fact is more reliable than clustering. This is because the web servers do not contain data that changes, they only handle client requests for data. The SQL servers will contain data that changes.

Implementation of clusters is based on linking nodes together using a private network, which does not support client connections. A message is sent periodically between the nodes to demonstrate that the node is healthy and available. Absence of this message from a node would cause a failover.

To allow for a quick shift of services from one node to another, most vendors support shared disk technology. As well as their own private storage, nodes share access to a disk array (SAN) which contains the data that will be switched during a failover.

Hot, Warm and Cold Spares

In the absence of configuring clusters, many admins have standby servers available as spares but that are not connected to the network. Standby servers are classified as Cold, warm or hot.

- Cold Server – is configured with an OS but is turned off
- Warm Server – the server is updated periodically with backup data
- Hot Server – the server is updated regularly and is ready to take over in the event of a failover.

Hot, Warm and Cold Sites

Enterprise networks often provide spare sites. A site is another location that can provide the same or similar levels of service. Operations are swapped over to the new site until the existing site can be brought back on line.

Sites are also referred to as cold, warm or hot. A hot site can be brought in to operation almost immediately. It usually means that the site is owned by the organisation and is ready

to deploy in less than 24 hours. A cold site takes longer typically in the region of one week, it could typically be an empty building that is leased but would not contain any IT equipment.

Providing redundancy on this scale will be expensive. Sites are often leased from service providers.

Replication

This is the process of duplicating data between different servers or sites. RAID mirroring and server clustering are examples of disk to disk and server to server replication.

Replication can be synchronous or asynchronous. Synchronous replication means that the data must be written at both sites before it can be considered committed. Asynchronous replication means that the data is mirrored from a primary site to a secondary site.

Disk to disk and server to server replication are relatively simple to accomplish as they can use direct access RAID or local network technologies. Site-to-site replication is considerably harder and more expensive as it relies on Wide area Network Technologies. Distance between sites has to be considered. If they are too close, both could be affected by the disaster. If further apart, the more costly it will be to replicate over a longer communications pathway and latency will be greater. This can be mitigated by using fiber optic links.