BSc in Computer Engineering
CMP4103
Computer Systems and Network Security

Lecture 12

Operations Security

Eng Diarmuid O'Briain, CEng, CISSP

Department of Electrical and Computer Engineering,
College of Engineering, Design, Art and Technology,
Makerere University

Table of Contents

This page intentionally left blank

1. Operations

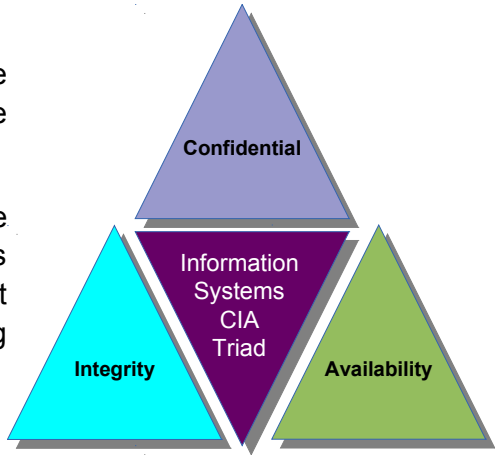

Information Systems Operations involves the confidentiality of integrity of information and the availability of systems.

The CIA Triad while essentially the core principle of information security it also applies to show that Operations Security is about protecting information assets by reducing threats and vulnerabilities.

Operations Security is about:

- identifying the resources to be protected
- defining the privileges that must be restricted
- determining the available control mechanisms
- appreciating the potential for abuse of access
- ensuring the appropriate use of controls
- implementing good security practice

The focus of information systems operations is about people accessing equipment and the backing up of that equipment.

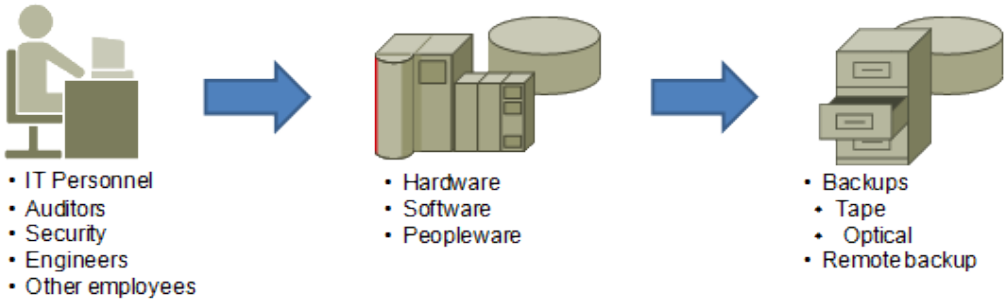

2. Access Control Categories

Access Controls protect confidentiality, integrity and availability of objects.

2.1 Preventive

Stop unwanted or unauthorised activity from occurring. These include biometrics, fences and locks, data classification, job rotation and separation of duties plus auditing, cryptography and monitoring.

2.2 Deterrent

Discourages the violation of security policies, often filling the gap left by preventive controls. They include gates, keyed access and security guards, badges, cameras and intrusion alarms or awareness training, separation of duties and security clearances.

2.3 Detective

Discovers unwanted or unauthorised activity but often take effect after an incident has occurred as opposed to before or during its occurrence. Security patrols, Security badges, guard dogs, security cameras, motion detectors and sound alarms or incident investigations, supervisory review, audits, and violation or exception reports.

2.4 Corrective

These restore systems to a known-good state following a security-related breach or incident. Access termination, service restarting or system rebooting, the implementation of intrusion detection systems, antivirus programs and malware scanners or business continuity planning, disaster recovery planning and security policies are typical corrective access controls.

2.5 Recovery

These are used to repair and restore critically damaged capabilities, functions and resources following a security violation. These are more complex in scale and scope than corrective controls and include backups, rollbacks and restorations, fault-tolerance, redundancy and clustering, or antivirus scanners, database shadowing and data replication.

2.6 Compensation

This provides aid to various other existing controls in the enforcement of system-wide security policies. This includes security policies, operational requirements or utilisation criteria or personnel supervision, monitoring and work procedures.

2.7 Directive

This confines and controls the actions of subjects to enforce and encourage strict security policy compliance. Security guards, guard dogs and security cameras, policy requirements, security criteria and posted notifications, or escape routes, employee supervision and awareness training all come under the umbrella of directive access control.

2.8 Administrative

This is the defined policies and procedures of the organisation that governs overall access, focusing on personnel and business practices. Workplace policies, procedures and hiring practices, background checks, data classification and security training or work reviews, employee supervision and personnel controls are examples.

2.9 Logical or technical

These are hardware and software mechanisms that manage access to and provide protection for shared computer and network resources. Encryption, passwords and smartcards, access control lists, biometrics and constrained interfaces or cryptographic protocols, firewall appliances and network routers.

2.10 Physical

These are structural barriers em-placed to prevent direct access to components of a facility, network or system. Security guards, guard dogs and fences, alarm systems, motion detectors and security windows or security lights, security locks and video cameras.

3. Resource protection

Resource protection involves the protective controls for the personnel, facility and equipment within. It involves a combination of measures including security, fire prevention, incident detection, process control, fire protection systems and incident response.

3.1 Physical protection of equipment

Access to the company sensitive equipment must be highly controlled. Physical access by the wrong people can spell disaster for the organisation. Backups of all data must be regular and sent off-site to a secure storage facility on a regular basis as part of the Business Continuity Plan (BCP).

3.1.1 Media Management

Storage

Storing critical information in electronic format provides businesses with a cost effective means to back up their information systems. The media devices are backed up and stored in an environment that will ensure the longevity and integrity of your vital information. The organisation should be confident that in the event of a disaster the electronic media can be easily retrieved to get operations back up and running.

The following should be considered when storing media:

- Temperature and Humidity Controlled Environment
- Static Free Surroundings
- Fire Suppressant Systems
- Fire Protection

Encryption

Sensitive data on media needs protection. A programme of encryption should be considered in general but particularly so if media is to be stored off-site.

Retrieval

Transporting media to backup locations must be organised and timely. During such transfers it is possible that media will be lost or stolen and a mechanism of retrieval of the data must be put in place until it is confirmed to be in backup.

Disposal

The proper disposal of media must match the highest classification of data which is contained on that device. For example a USB Stick with classified sensitive and restricted data must be disposed of in the manner required for the disposal of restricted data.

For storage devices to be recycled a process confirming the multiple pass secure overwrite should be completed. This is an overwrite of all addressable locations with 2 different characters. A more complete overwrite is the overwrite of all addressable locations with a character, its complement, then a random character and verify.

For media with highly sensitive data or media no longer to be used by the organisation they should be securely destroyed. Processes include disintegrate, incinerate, pulverise, shred, or melt the media.

Marking

It is recommended that data is marked to indicate the level of protection the data requires. A Protective Marking Scheme categorises the level of confidentiality. Military data classification scale of Unclassified, Sensitive But Unclassified (SBU), Restricted, Confidential, Secret and Top Secret or commercial data classification scale of Public, Sensitive, Private and Confidential or similar should be used.

At scheduled reviews data classification should be reviewed to consider re-classification. This should be part of a process and carried out by a employee of sufficient clearance to do so.

3.1.2 Records management

This is the practice of maintaining the records of an organisation from the time they are created up to their eventual disposal. This may include classifying, storing, securing, and destruction or archival preservation of records.

A record can be either a tangible object or digital information, for example, birth certificates, medical x-rays, office documents, databases, application data, and e-mail. Records management is primarily concerned with the evidence of an organisation's activities, and is usually applied according to the value of the records rather than their physical format.

3.1.3 Fire

Fire prevention, detection and suppression systems are essential for the control of fire to ensure:

- Safety of the lives of employees and visitors
- Continuity Of Operations (COOP)
- Property protection

3.1.4 Electrical Power

Ensure the provision of clean steady electrical mains. Uninterruptable Power Supplies (UPS) and generators are essential to ensure that in the event of a mains failure the data centre equipment continues to operate. Service Level Agreements (SLA) with the power companies should be negotiated to ensure timely response to any failures that occur.

3.1.5 Heating, Ventilating, and Air Conditioning (HVAC)

HVAC systems are essential to maintain the correct levels of humidity and temperature in data centres to optimise the operation of the sensitive equipment.

3.1.6 Water

Protection against water and humidity damage in data centres must be considered. Humidity and water sensors to alarm in the event of either and proper planning of the building are essential. It is for this reason that data centres are usually located on a centre floor so leaks in the roof will not impact the data centre and flooding on the ground floor is also unlikely to impact the centre. Another consideration is the location of toilets on the data centre floor and on floors above.

3.1.7 Communications

Communications links in and out of the data centre may be critical to operations. SLAs should be drawn up with providers and redundant links from different providers should be considered to ensure continuity of operations.

4. Administrative Control

Operations team enforces company policies on behalf of management. This enforcement operate and manage access control to systems and data, detect attacks using Intrusion Detection Systems (IDS) and block unauthorised access while permitting authorised communications with the use of firewalls.

The Operations team must be trustworthy as they are the guardians of the network, nevertheless privileges and rights are required that ensure the activities of the operations personnel is logged and that the data they access is on a strict 'need to know' basis and for specific tasks.

4.1 Separation of Duties

It is important that operations personnel do not have too much access to the system. Separation of duties to different personnel will help to prevent one person accessing the chain of resources that will allow them sufficient access to compromise the company's security or commit fraudulent activities. With separation of duties it would require "collusion" with other operations personnel to bypass this security device. Examples are the auditor of access permissions should not be the administrator applying the access permissions or the programmer should not be permitted to be the formal tester of the code that is written. The following jobs should never be performed by the same individual.

- System Administrator
- Network Administrator
- E-mail Administrator
- Security Administrator
- Database Administrator

4.2 Job Rotation

The regular movement of personnel around different functions in the operations team is healthy. All job functions should have more than one person trained to do them.

4.3 Mandatory Vacations

Random mandatory vacations without network access are a good way to reduce the opportunity for fraud.

4.4 Security Violations

Security Violations should be documented fully and root cause should be conducted to determine if changes in processes or systems are needed to prevent the violation occurring again.

4.5 Disciplinary process / Termination

Violations of the security policy should be treated very harshly and should be generally considered subject to employee termination.

5. CIS Critical Security Controls (CSC)

5.1 Center for Internet Security (CIS)

In 2008, the Center for Internet Security's (*https://www.cisecurity.org/*) Critical Security Controls ("CIS Controls") were created as a collaboration between representatives from the U.S. government and private sector security research organisations. A set of practical defences specifically targeted toward stopping cyber attacks, these proposed defences were technical in nature and intended to define specific, practical steps an organisation could take to stop the most common cyber threats from compromising their information systems. The CIS Controls were crafted to answer the frequent question: "Where should I start when I want to improve my cyber defences?"

The first five CIS CSCs are often referred to as providing cybersecurity ***hygiene***, as a number of studies show that implementation of these provide an effective defence against the most common cyber-attacks (~80% of attacks).

The five critical tenets of an effective cyber defence system as reflected in the CIS CSCs are:

Offence informs defence

- Use knowledge of actual attacks that have compromised systems to provide the foundation to continually learn from these events to build effective, practical defences. Include only those controls that can be shown to stop known real-world attacks.

Prioritisation

- Invest first in Controls that will provide the greatest risk reduction and protection against the most dangerous threat actors and that can be feasibly implemented in your computing environment.

Metrics

- Establish common metrics to provide a shared language for executives, IT specialists, auditors, and security officials to measure the effectiveness of security measures within an organisation so that required adjustments can be identified and implemented quickly.

Continuous diagnostics and mitigation

- Carry out continuous measurement to test and validate the effectiveness of current security measures and to help drive the priority of next steps.

Automation

- Automate defences so that organisations can achieve reliable, scalable, and continuous measurements of their adherence to the Controls and related metrics.

5.2 CSC versions

The Controls were developed based on specific knowledge of the threat environment as well as the current technologies in the marketplace upon which communications and data rely. One of the key benefits of the Controls is that they are not static; they are updated regularly and are tailored to address the security issues of the day. These notes are based on version 6.1 of CIS CSC.

5.3 The first 5 CSCs

Controls CSC 1 through CSC 5 are essential to success and should be considered among the very first things to be done. These as considered **Foundational Cyber Hygiene** – the basic things that an organisation must do to create a strong foundation for its defence.

5.3.1 CSC 1 - Inventory of Authorised and Unauthorised Devices

CSC1 helps organisations to define a baseline of what must be defended. Without an understanding of what devices and data are connected, they cannot be defended. Scanning the network to create an inventory is a good start, i.e. Passive Asset Detection System (PADS) for example.

The next step is to prevent unauthorised devices from joining a network. The initial goal is not to prevent attackers from joining the network, as much as it is to understand what is on the network so it can be defended.

5.3.2 CSC 2 - Inventory of Authorised and Unauthorised Software

This control ensures that only authorised software is allowed to execute on an organisation's information systems. While an inventory of software is important, the most crucial control an organisation can implement here is application whitelisting, which limits the ability to run applications to only those which are explicitly approved. CSC 2 is often considered one of the most effective at preventing and detecting cyberattacks, although it is often not easily implemented.

5.3.3 CSC 3 - Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers

Most technology systems are installed with a focus on ease-of-use and not necessarily security. Most organisations already have the technology systems necessary to securely configure their systems at scale. For Unix/Linux systems Ansible, Puppet or Chef are available and for Microsoft there is the Active Directory Group Policy Objects available for organisations. By utilising such configuration standards or benchmarks defined by the Center for Internet Security, or found in the NIST National Checklist Program Repository (*https://www.nist.gov/programs-projects/national-checklist-program*), this Control is achievable by most organisations.

5.3.4 CSC 4 - Continuous Vulnerability Assessment and Remediation

The CSC 4 control is to understand the technical software weaknesses that exist in an organisation's information systems and to remove or remediate those weaknesses. Successful organisations implement patch management systems that cover both Operating System (OS) and third-party application vulnerabilities. This allows for the automatic, ongoing, and proactive installation of updates to address software vulnerabilities. In addition to patch management systems, organisations must implement a vulnerability management system to give themselves the ability to detect where exploitable software weaknesses currently exist so they can be remediated.

5.3.5 CSC 5 - Controlled Use of Administrative Privileges

The purpose of this Control is to ensure that workforce members have only the system rights, privileges and permissions that they need in order to do their job - no more and no less than necessary. Unfortunately, for the sake of speed and convenience, many organisations allow staff to have local system or even domain administrator rights which are too generous and open the door for abuse, accidental or otherwise.

5.4 The other CSCs

The following additional controls should be tackled once the first five have been fully implemented. These controls are listed in order of importance as determined by the CIS.

- CSC 6: Maintenance, Monitoring, and Analysis of Audit Logs.
- CSC 7: Email and Web Browser Protections.
- CSC 8: Malware defences.
- CSC 9: Limitation and Control of Network Ports, Protocols, and Services.
- CSC 10: Data Recovery Capability.
- CSC 11: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches.
- CSC 12: Boundary Defence.
- CSC 13: Data Protection.
- CSC 14: Controlled Access Based on the Need to Know.
- CSC 15: Wireless Access Control.
- CSC 16: Account Monitoring and Control.
- CSC 17: Security Skills Assessment and Appropriate Training to Fill Gaps.
- CSC 18: Application Software Security.
- CSC 19: Incident Response and Management.
- CSC 20: Penetration Tests and Red Team Exercises.

6. Informations Systems Operations functions

6.1 Threat awareness

The operations team should have threat awareness and specific countermeasures for

- Media Libraries
- Errors and Omissions
- Fraud and Theft
- Employee Sabotage
- Loss of Physical Support
- Industrial Espionage
- Loss of infrastructure support
- Hackers
- Malicious code
    - Worms
    - Viruses
    - Trojan horses

6.2 Protection of Information

- Backup of Critical Information regularly
- Perform offsite backups
- Redundancy
    - High Availability (HA)
    - RAID
- System trusted recovery

6.3 Fault Tolerant Systems

Fault tolerance is the feature of a system to continue operating in the event of the failure of some of its components.

The basic characteristics of fault tolerance require:

- No single point of failure
- No single point of repair
- Fault isolation to the failing component
- Fault containment to prevent propagation of the failure
- Availability of reversion modes

A fault tolerant system can be implemented using one of the following three fault tolerant configurations:

- **Hot Standby**
  - The primary and backup systems run simultaneously. The data is mirrored to the secondary server in real time so that both systems contain identical information.
- **Warm Standby**
  - The backup system runs in the background of the primary system. Data is mirrored to the secondary server at regular intervals, which means that there are times when both servers do not contain the exact same data.
- **Cold Standby**
  - The backup system is only called upon when the primary system fails. The system on cold standby receives scheduled data backups, but less frequently than a warm standby. Cold standby systems are used for non-critical applications or in cases where data is changed infrequently.

Additional components is a good method of addressing the standby requirement. This can be addressed in three ways:

- **Replication**
  - Providing multiple identical instances of the same system or subsystem, directing tasks or requests to all of them in parallel, and choosing the correct result on the basis of a quorum;
- **Redundancy**
  - Providing multiple identical instances of the same system and switching to one of the remaining instances in case of a failure (failover);
- **Diversity**
  - Providing multiple different implementations of the same specification, and using them like replicated systems to cope with errors in a specific implementation.

6.3.1 High-availability (HA) clusters

High-availability clusters are implemented primarily for improving the availability of services, which the cluster provides. They operate by having redundant nodes, which are then used to provide service when system components fail. The most common size for an HA cluster is two nodes, which is the minimum requirement to provide redundancy. HA cluster implementations attempt to use redundancy of cluster components to eliminate single points of failure.

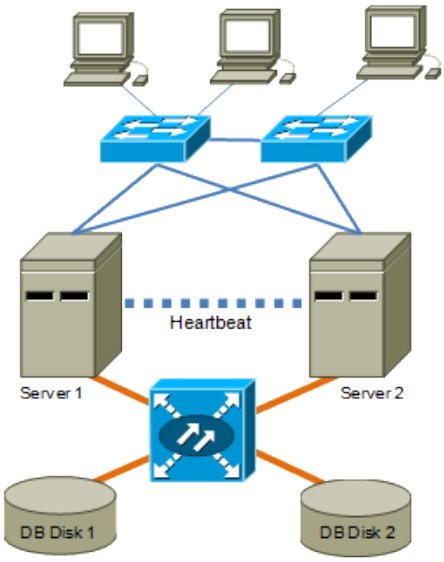

There are many commercial implementations of High-Availability clusters for many operating systems.

6.3.2 Redundant Array of Inexpensive/Independent Disks (RAID)

RAID is a technology that allowed computer users to achieve high levels of storage reliability from low-cost and less reliable PC-class disk-drive components, via the technique of arranging the devices into arrays for redundancy.

Marketers representing industry RAID manufacturers later reinvented the term to describe a redundant array of independent disks as a means of dissociating a "low cost" expectation from RAID technology.

"RAID" is now used as an umbrella term for computer data storage schemes that can divide and replicate data among multiple hard disk drives. The different schemes/architectures are named by the word RAID followed by a number, as in RAID 0, RAID 1, etc. RAID's various designs involve two key design goals: increase data reliability and/or increase input/output performance. When multiple physical disks are set up to use RAID technology, they are said to be in a RAID array. This array distributes data across multiple disks, but the array is seen by the computer user and operating system as one single disk. RAID can be set up to serve several different purposes.

There are various combinations giving different trade-offs of protection against data loss, capacity, and speed. RAID levels 0, 1, and 5 are the most commonly found, and cover most requirements.

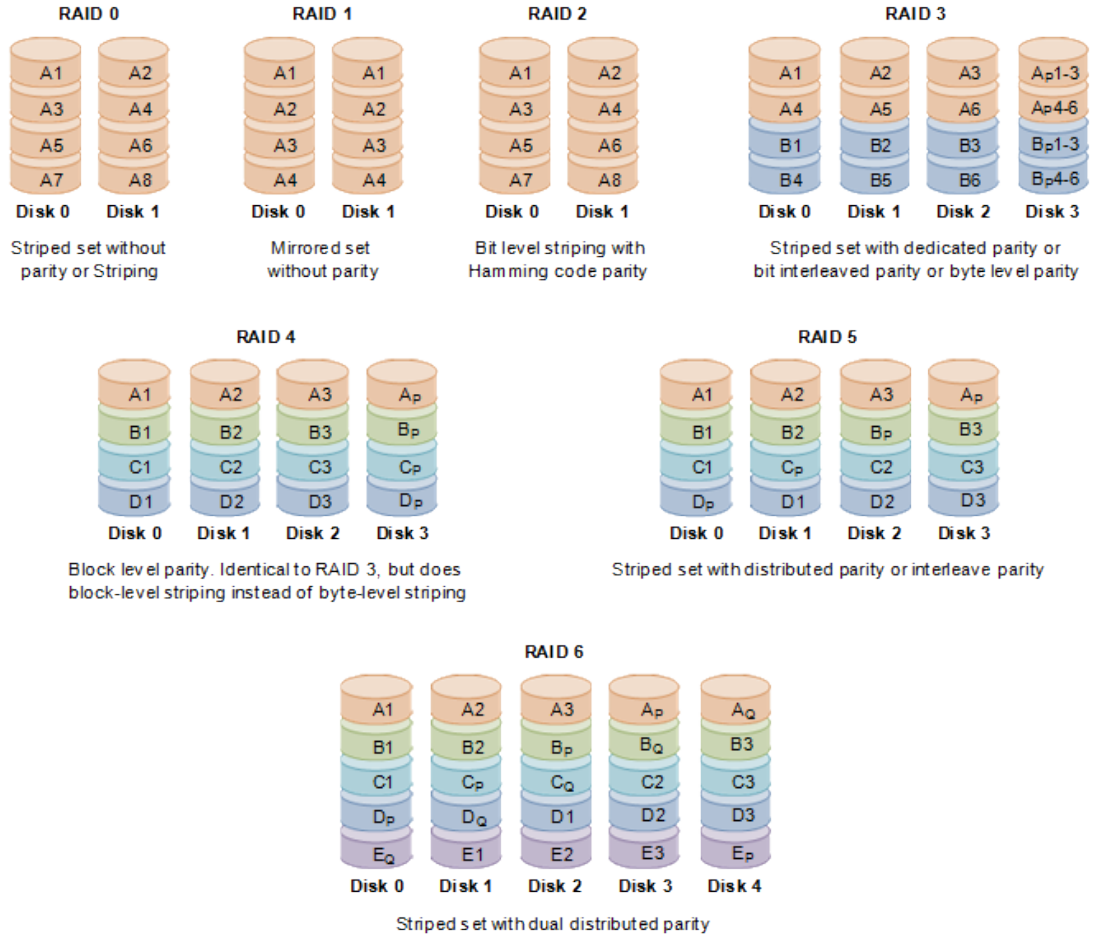

6.3.2.1 RAID 0

(striped disks) distributes data across multiple disks in a way that gives improved speed at any given instant. If one disk fails, however, all of the data on the array will be lost, as there is neither parity nor mirroring. In this regard, RAID 0 is somewhat of a misnomer, in that RAID 0 is non-redundant. A RAID 0 array requires a minimum of two drives. A RAID 0 configuration can be applied to a single drive provided that the RAID controller is hardware and not software (i.e. OS-based arrays) and allows for such configuration. This allows a single drive to be added to a controller already containing another RAID configuration when the user does not wish to add the additional drive to the existing array. In this case, the controller would be set up as RAID only (as opposed to SCSI only (no RAID)), which requires that each individual drive be a part of some sort of RAID array.

6.3.2.2 RAID 1

Mirrors the contents of the disks, making a form of 1:1 ratio real-time backup. The contents of each disk in the array are identical to that of every other disk in the array. A RAID 1 array requires a minimum of two drives. RAID 1 mirrors, though during the writing process copy the data identically to both drives, would not be suitable as a permanent backup solution, as RAID technology by design allows for certain failures to take place.

6.3.2.3 RAID 3 or 4

(striped disks with dedicated parity) combines three or more disks in a way that protects data against loss of any one disk. The storage capacity of the array is reduced by one disk. A RAID 3 or 4 array requires a minimum of three drives: two to hold striped data, and a third drive to hold parity data.

6.3.2.4 RAID 5

(striped disks with distributed parity) combines three or more disks in a way that protects data against the loss of any one disk. The storage capacity of the array is a function of the number of drives minus the space needed to store parity. The maximum number of drives that can fail in any RAID 5 configuration without losing data is only one. Losing two drives in a RAID 5 array is referred to as a "double fault" and results in data loss.

6.3.2.5 RAID 6

(striped disks with dual parity) combines four or more disks in a way that protects data against loss of any two disks.

6.3.2.6 RAID 1+0

(or 10) is a mirrored data set (RAID 1) which is then striped (RAID 0), hence the "1+0" name. A RAID 10 array requires a minimum of two drives, but is more commonly implemented with 4 drives to take advantage of speed benefits.

6.3.2.7 RAID 0+1

(or 01) is a striped data set (RAID 0) which is then mirrored (RAID 1). A RAID 0+1 array requires a minimum of four drives: two to hold the striped data, plus another two to mirror the first pair.

6.3.3 System Recovery

System failure has the possibility of both the risk of loss of data as well as the possibility of security risks. Trusted recovery is the process of the system administrator bringing the system operational before allowing users access to the system while not yet in a fully operational state. Types of trusted recovery include:

- **System Cold Start**
  - A normal system start is not possible due to unexpected failures. The Administrator will need to intervene to bring the system to a normal state. See example.
- **Emergency System Restart**
  - This is typical of when the system fails and brings itself to a maintenance mode to perform file recovery and restarts with none of the user process that existed at the time of the failure restored.
- **System Reboot**
  - This is carried out after the administrator has noticed a failure and shutdown the system in a controlled manner.

Example of a trusted start:

- Reboot system to single user mode
- Recover all active file systems at the time of failure
- Restore missing or damaged files from backups
- Recover security labels to missing files
- Check security critical files
- Allow users access to the system

7. Change and Configuration Management

Change Management is the process of managing change. A Process must exist for users to submit a change request. Once such a request is received, how is it tracked, prioritised and implemented?

Configuration management is the management of security features and assurances through control of changes made to hardware, software, firmware, documentation, test, test fixtures, and test documentation throughout the life cycle of an information system.

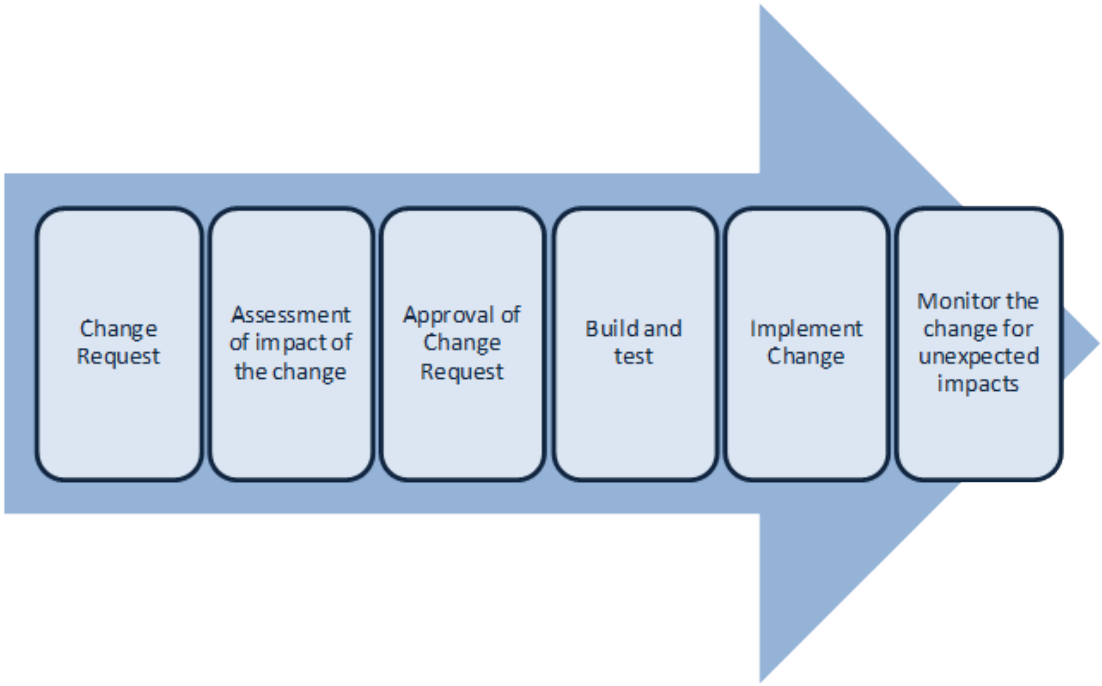

7.1 Change Control Procedures

7.1.1 Record Change Request

The client initiates change request as a formal request for something to be changed. The change control team then records and categorises that request. This categorisation would include estimates of importance, impact, and complexity.

7.1.2 Assessment of the impact of the change

The impact assessor or assessors then make their risk analysis typically by answering a set of questions concerning risk, both to the business and to the process, and follow this by making a judgement on who should carry out the change. If the change requires more than one type of assessment, the head of the change control team will consolidate these.

7.1.3 Approval of the Change Request

Everyone with a stake in the change then must meet to determine whether there is a business or technical justification for the change. If approved the change is then sent to the delivery team for planning and build.

7.1.4 Build and test

Management will assign the change to a specific delivery team, usually one with the specific role of carrying out this particular type of change. The team's first job is to plan the change in detail as well as construct a regression plan in case the change needs to be backed out. If all stakeholders agree with the plan, the delivery team will build the solution, which will then be tested. They will then seek approval and request a time and date to carry out the implementation phase.

7.1.5 Implement Change

All stakeholders must agree to a time, date and cost of implementation. Following implementation, it is usual to carry out a post-implementation review which would take place at another stakeholder meeting.

7.1.6 Monitor

After implementation, the system requires monitoring to ensure the change does not lead to unexpected adverse effects.

8. Exercise

Consider a company who has built an online presence consisting of a redundant website, with on-line shopping facilities. They have dual site redundancy.

Carry-out in groups a plan for the system, consider:

- Location of server(s)
- Contracts with site owner(s).
- Access Agreements.
- Access Control.
- High Availability.
- Security.
- Employee roles.
- Operations policies.
- Change Control Mechanisms.

9. Bibliography

Lee J (2005). Scalable Continuous Media Streaming Systems: Architecture, Design, Analysis. John Wiley & Sons.

(ISC)² (2012) Official (ISC)² Guide to the CISSP Common Body of Knowledge. Third Edition.

This page is intentionally blank