Cybersecurity for Industrial Networks

Topic 1
Operations, Business Continuity and Disaster Recovery

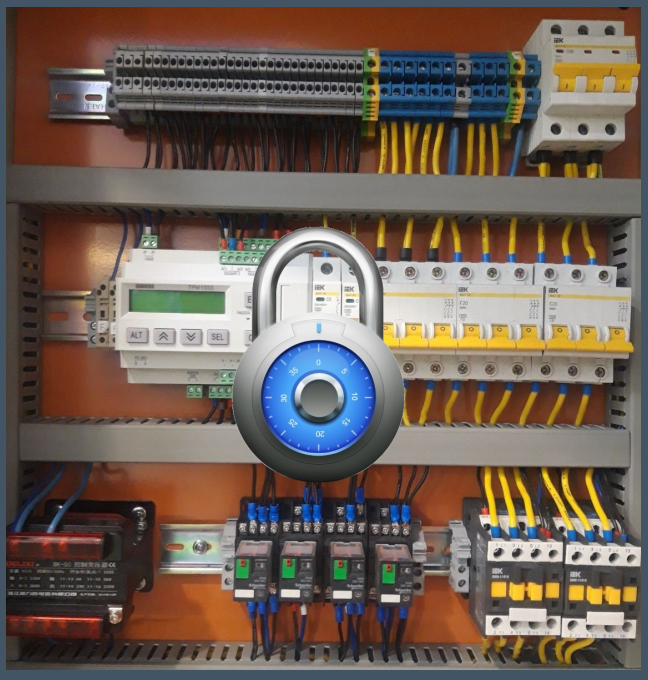

Dr Diarmuid Ó Briain

Version: 1.0

Ollscoil Teicneolaíochta an Oirdheiscirt
South East Technological University

Dr Diarmuid Ó Briain

Table of Contents

Illustration Index

1 Objectives

By the end of this topic, you will be able to:

- Understand the principles of managing and protecting information systems operations.
- Apply effective access control strategies to protect information systems assets.
- Implement comprehensive resource protection measures to safeguard information systems.
- Develop and execute an effective Business Continuity Plan (BCP) to mitigate disruptions.

2 Operations

As illustrated in Figure 1, Information Systems Operations involves the Confidentiality of the Integrity of information and the Availability of systems; however, for Operational Technology (OT) the focus switches to Human Safety, the Availability of systems, the integrity of those systems and their data, followed by Confidentiality of data.

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

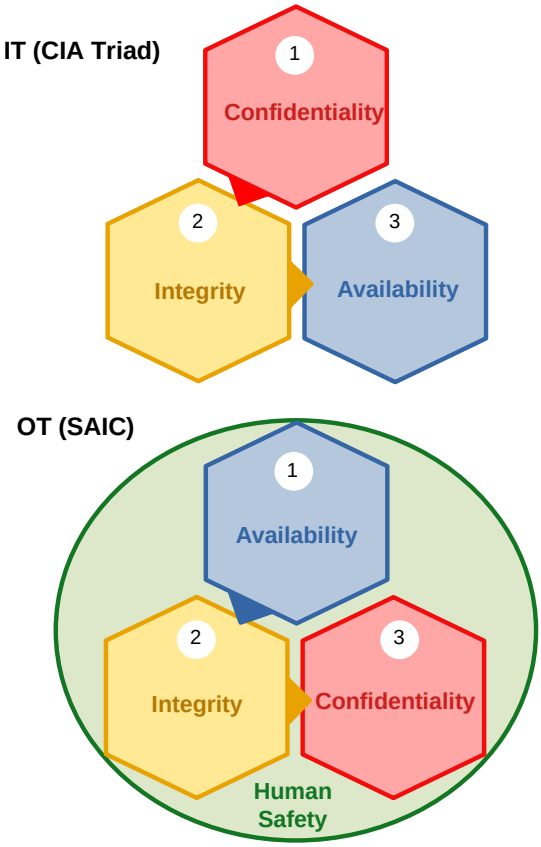

Figure 1: CIA / SAIC

The focus of Information and Operational Technology Systems Operations, shown in Figure 2, is about people accessing equipment and the backing up of that equipment.

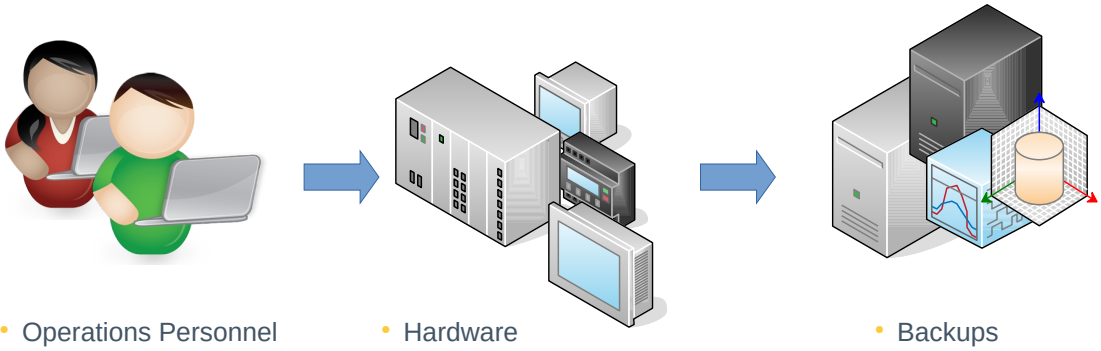

- Operations Personnel
- IT Personnel
- Auditors
- Security
- Engineers
- Other Employees

- Hardware
- Software
- Peopleware

- Backups
  - Tape
  - Optical
- Remote Backups

Figure 2: Information and Operational Technology Systems Operations

3 Access Control Categories

Access Controls protect Confidentiality, Integrity and Availability of objects.

3.1 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.

3.2 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.

3.3 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.

3.4 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.

3.5 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.

3.6 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.

3.7 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.

3.8 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.

3.9 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.

3.10 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.

4 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.

4.1 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 BCP.

4.1.1 Media Management

Storage

Storing critical information in electronic format provides businesses with a cost effective means to back up data from their operational and informational technology 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.

4.1.2 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.

4.1.3 Fire

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

1. Safety of the lives of employees and visitors
2. Continuity Of OPerations (COOP)
3. Property protection

4.1.4 Electrical Power

Ensure the provision of clean steady electrical mains. Uninterruptible 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.

4.1.5 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.

4.1.6 Water

Protection against water and humidity damage in manufacturing floors and 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.

4.1.7 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.

5 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.

5.1 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
- Operations Engineer
- Database Administrator
- Security Administrator

5.2 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.

5.3 Mandatory Vacations

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

5.4 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.

5.5 Disciplinary process / Termination

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

6 CIS Critical Security Controls

6.1 Center for Internet Security (CIS)

In 2008, the Center for Internet Security's (*https://www.cisecurity.org*) Critical Security Controls (CSC) were created as a collaboration between representatives from the U.S. government and private sector security research organisations [1]. 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 CIS Controls are a prioritised set of safeguards to mitigate the most prevalent cyber-attacks against systems and networks. They are considered the gold standard for cybersecurity best practices and are widely used by organisations of all sizes to improve their security posture.

Version 8 of the controls reflect the evolving cybersecurity landscape and introduces several key changes, including:

- **Focus on activity-based controls:** CIS Controls v8 groups controls by activities rather than by who manages the devices. This change reflects the growing importance of cloud-based computing and the increasing interconnectedness of IT systems.
- **Reduction in the number of controls:** The number of CIS Controls has been reduced from 20 to 18 to streamline the implementation process and make it easier for organisations to focus on the most critical controls.
- **Emphasis on hybrid and cloud environments:** CIS Controls v8 provides specific guidance for organisations that are moving to a hybrid or fully cloud environment.

6.1.1 Design Principles

The CIS Design Principles include:

Offence Informs Defence

- CIS Controls are selected, dropped, and prioritised based on data, and on specific knowledge of attacker behaviour and how to stop it.

Focus

- Help defenders identify the most critical things they need to do to stop the most important attacks.
- Avoid being tempted to solve every security problem—avoid adding “good things to do” or “things you could do”.

Feasible

- All individual recommendations (Safeguards) must be specific and practical to implement.

Measurable

- All CIS Controls, especially for Implementation Group 1, must be measurable.
- Simplify or remove ambiguous language to avoid inconsistent interpretation.
- Some Safeguards may have a threshold.

Align

- Create and demonstrate “peaceful co-existence” with other governance, regulatory, process management schemes, framework, and structures.
- Cooperate with and point to existing, independent standards and security recommendations where they exist, e.g., National Institute of Standards and Technology (NIST), Cloud Security Alliance (CSA), Software Assurance Forum for Excellence in Code (SAFECode), MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK), Open Web Application Security Project (OWASP).

6.2 Implementation Groups

CIS Controls Implementation Groups (IG) were added in version 8 as recommended guidance to prioritise implementation. Each IG identifies a subset of the CIS Controls that the community has broadly assessed to be applicable for an enterprise with a similar risk profile and resources to strive to implement. These IGs represent a horizontal look across the CIS Controls tailored to different types of enterprises. Specifically, IG1 is defined as “*essential cyber hygiene*,” the foundational set of cyber defence Safeguards that every enterprise should apply to guard against the most common attacks.

Each subsequent IG then builds upon the previous one: IG2 includes IG1, and IG3 includes all CIS Safeguards in IG1 and IG2.

- **IG1** - Safeguards selected for IG1 should be implementable with limited cybersecurity expertise and aimed to thwart general, non-targeted attacks. These Safeguards will also typically be designed to work in conjunction with small or home office Commercial Off-The-Shelf (COTS) hardware and software.
- **IG2** (Includes IG1) - An IG2 enterprise employs individuals responsible for managing and protecting IT infrastructure. These enterprises support multiple departments with differing risk profiles based on job function and mission. Small enterprise units may have regulatory compliance burdens. IG2 enterprises often store and process sensitive client or enterprise information and can withstand short interruptions of service. A major concern is loss of public confidence if a breach occurs. Safeguards selected for IG2 help security teams cope with increased operational complexity. Some Safeguards will depend on enterprise-grade technology and specialised expertise to properly install and configure.
- **IG3** (Includes IG1 and IG2) - An IG3 enterprise employs security experts that specialise in the different facets of cybersecurity (e.g., risk management, penetration testing, application security). IG3 assets and data contain sensitive information or functions that are subject to regulatory and compliance oversight. An IG3 enterprise must address availability of services and the confidentiality and integrity of sensitive data. Successful attacks can cause significant harm to the public welfare. Safeguards selected for IG3 must abate targeted attacks from a sophisticated adversary and reduce the impact of zero-day attacks.

6.3 Critical Security Controls

There are 18 CSCs in all. In earlier versions of the CIS Controls the first five were considered **Foundational Cyber Hygiene** – the basic things that an organisation must do to create a strong foundation for its defence and essential to success. However, since the introduction of IGs there are safeguards within each CSC associated with each IG that must be implemented.

6.4 The 18 CSCs

6.4.1 *CSC 1* - Inventory and Control of Enterprise Assets

Actively manage (inventory, track, and correct) all enterprise assets (end-user devices, including portable and mobile; network devices; non-computing/Internet of Things (IoT) devices; and servers) connected to the infrastructure physically, virtually, remotely, and those within cloud environments, to accurately know the totality of assets that need to be monitored and protected within the enterprise. This will also support identifying unauthorised and unmanaged assets to remove or remediate.

6.4.2 *CSC 2* - Inventory and Control of Software Assets

Actively manage (inventory, track, and correct) all software (operating systems and applications) on the network so that only authorised software is installed and can execute, and that unauthorised and unmanaged software is found and prevented from installation or execution.

6.4.3 *CSC 3* – Data Protection

Develop processes and technical controls to identify, classify, securely handle, retain, and dispose of data.

6.4.4 *CSC 4* – Secure Configuration of Enterprise Assets and Software

Establish and maintain the secure configuration of enterprise assets (end-user devices, including portable and mobile; network devices; non-computing/IoT devices; and servers) and software (operating systems and applications).

6.4.5 *CSC 5* – Account Management

Use processes and tools to assign and manage authorisation to credentials for user accounts, including administrator accounts, as well as service accounts, to enterprise assets and software.

6.4.6 *CSC 6* – Access Control Management

Use processes and tools to create, assign, manage, and revoke access credentials and privileges for user, administrator, and service accounts for enterprise assets and software.

6.4.7 *CSC 7* – Continuous Vulnerability Management

Develop a plan to continuously assess and track vulnerabilities on all enterprise assets within the enterprise's infrastructure, in order to remediate, and minimise, the window of opportunity for attackers. Monitor public and private industry sources for new threat and vulnerability information.

6.4.8 *CSC 8* – Audit Log Management

Collect, alert, review, and retain audit logs of events that could help detect, understand, or recover from an attack.

6.4.9 *CSC 9* – Email and Web Browser Protections

Improve protections and detections of threats from email and web vectors, as these are opportunities for attackers to manipulate human behaviour through direct engagement.

6.4.10 *CSC 10* – Malware Defences

Prevent or control the installation, spread, and execution of malicious applications, code, or scripts on enterprise assets.

6.4.11 CSC 11 – Data Recovery

Establish and maintain data recovery practices sufficient to restore in-scope enterprise assets to a pre-incident and trusted state.

6.4.12 CSC 12 – Network Infrastructure Management

Establish, implement, and actively manage (track, report, correct) network devices, in order to prevent attackers from exploiting vulnerable network services and access points.

6.4.13 CSC 13 – Network Monitoring and Defence

Operate processes and tooling to establish and maintain comprehensive network monitoring and defence against security threats across the enterprise's network infrastructure and user base.

6.4.14 CSC 14 – Security Awareness and Skills Training

Establish and maintain a security awareness program to influence behaviour among the workforce to be security conscious and properly skilled to reduce cybersecurity risks to the enterprise.

6.4.15 CSC 15 – Service Provider Management

Develop a process to evaluate service providers who hold sensitive data, or are responsible for an enterprise's critical IT platforms or processes, to ensure these providers are protecting those platforms and data appropriately.

6.4.16 CSC 16 – Application Software Security

Manage the security life cycle of in-house developed, hosted, or acquired software to prevent, detect, and remediate security weaknesses before they can impact the enterprise.

6.4.17 CSC 17 – Incident Response Management

Establish a programme to develop and maintain an incident response capability (e.g., policies, plans, procedures, defined roles, training, and communications) to prepare, detect, and quickly respond to an attack.

6.4.18 CSC 18 – Penetration Testing

Test the effectiveness and resiliency of enterprise assets through identifying and exploiting weaknesses in controls (people, processes, and technology), and simulating the objectives and actions of an attacker.

6.5 Safeguards

Each CSC has a number of safeguards associated with it. As an example consider "CSC 1 - **Inventory and Control of Enterprise Assets**". There are five safeguards in total, in IG1 there are two from the five, IG2 there are for of the five and all five apply to IG3. All five controls in this CSC apply to devices.

6.5.1 Safeguard 1.1 - Establish and Maintain Detailed Enterprise Asset Inventory

This safeguard has a security function to Identify and it applies across all three IGs.

6.5.2 Safeguard 1.2 - Address Unauthorised Assets

This safeguard has a security function to Respond and it also applies across all three IGs.

6.5.3 Safeguard 1.3 - Utilise an Active Discovery Tool

This safeguard has a security function to Detect and it only applies to IGs 2 and 3.

6.5.4 Safeguard 1.4 - Use DHCP Logging to Update Enterprise Asset Inventory

Like 1.1, this safeguard has a security function to Identify; however it only applies to IGs 2 and 3.

6.5.5 Safeguard 1.5 - Use a Passive Asset Discovery Tool

This safeguard has a security function to Detect and it only applies to IGs 3.

In a similar way each of the other CSCs have a set of safeguards defined.

7 Informations Systems Operations functions

7.1 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

7.2 Protection of Information

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

7.3 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.

7.3.1 High-availability clusters

High-availability clusters, as illustrated in Figure 3 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.

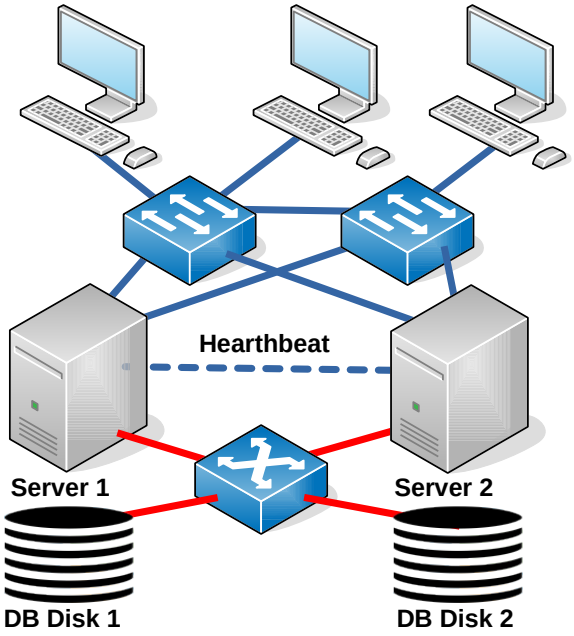

Figure 3: HA Clusters

7.3.2 Redundant Array of Independent Disks

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.

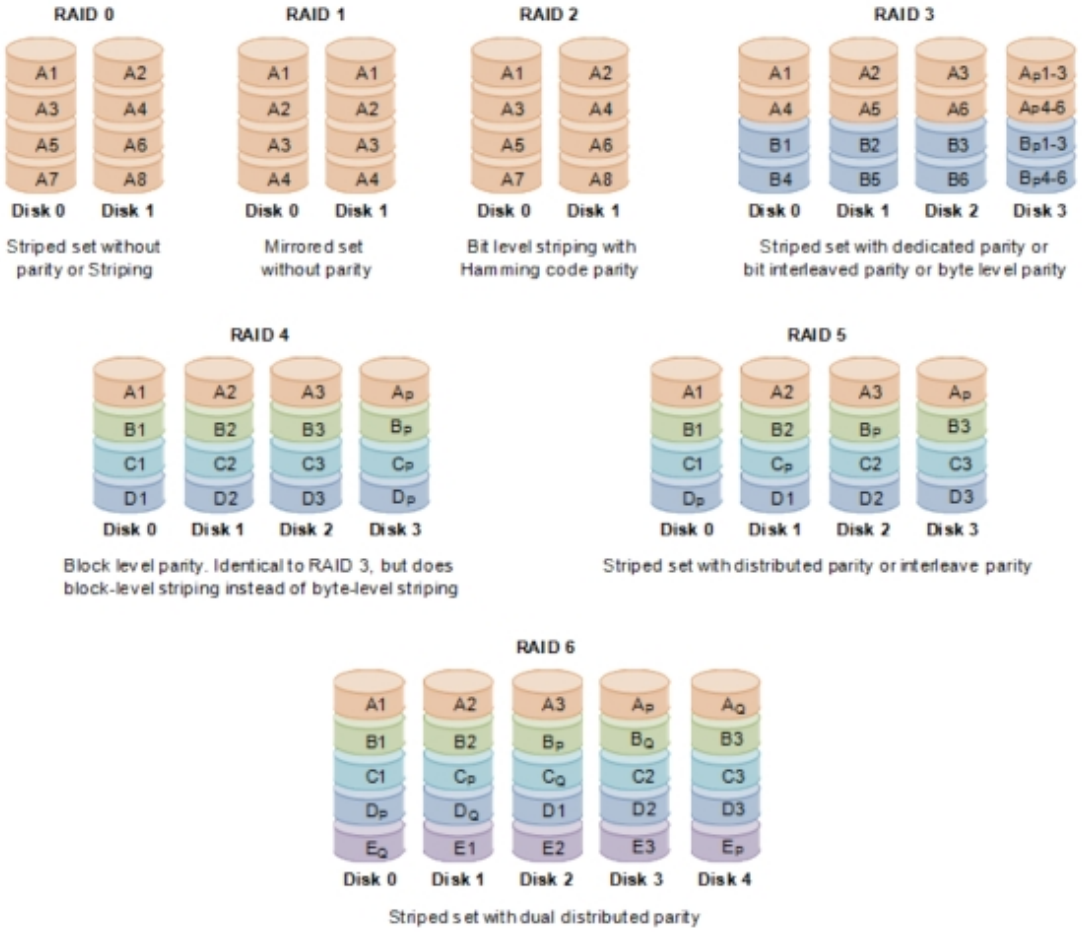

Figure 4: RAID

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.

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.

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.

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.

RAID 6

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

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.

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.

7.3.3 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

8 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.

8.1 Change Control Procedures

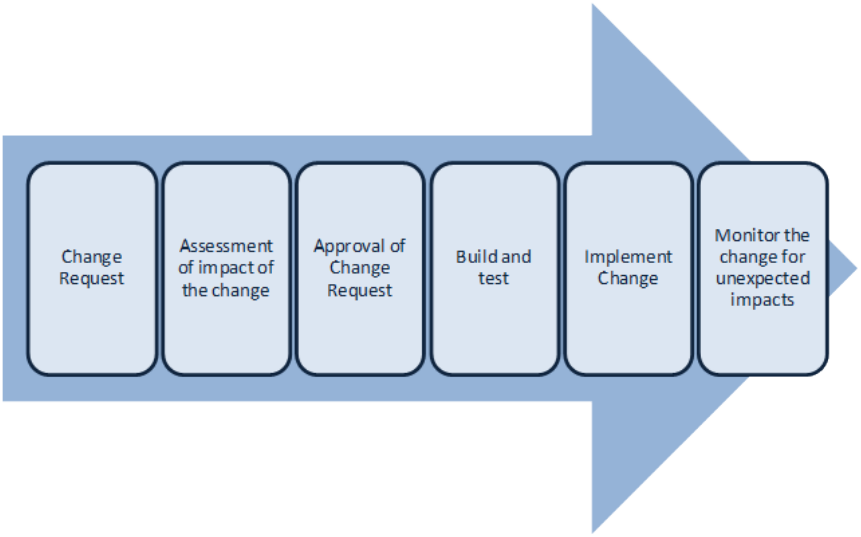

Figure 5: Change Control Procedures

8.1.1 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.

8.1.2 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.

8.1.3 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.

8.1.4 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.

8.1.5 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.

8.1.6 Monitor

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

9 Exercise #1: Operations

Consider a company who has built an online presence, for Just In Time (JIT) manufacturing, consisting of a redundant website, with on-line ordering 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.

10 Business Continuity Planning

BCP is the creation and validation of a practised logistical plan for how an organisation will recover and restore partially or completely interrupted critical functions within a predetermined time after a disaster or extended disruption.

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. There are two ISO standards that are relevant to BCP as well as NIST Special Publication 800-82r3.

10.1 ISO 22301:2019 - Business Continuity Management Systems

This standard specifies the requirements for an organisation's Business Continuity Management System (BCMS), which is a systematic approach to ensuring that the organisation can continue to operate critical business functions in the event of a disruptive event. The standard covers all aspects of BCP, from identifying and assessing risks to planning and implementing recovery procedures [2].

10.2 ISO 27031:2011 - ICT readiness for Business Continuity

This standard, formally titled "*Information technology — Security techniques — Guidelines for information and communication technology readiness for business continuity*" provides guidance on how to implement an Information and Communications Technology (ICT) readiness capability for business continuity. It covers all aspects of ICT readiness, from identifying and assessing ICT risks to planning and implementing ICT recovery procedures [3].

10.3 NIST Special Publication 800-82r3

BCP plays a crucial role in ensuring the resilience of OT systems against cyberattacks and other disruptive events. NIST Special Publication (SP) 800-82r3, "*Guide to Operational Technology (OT) Security*," provides comprehensive guidance on developing and implementing a robust BCP for OT environments [4].

Along with objectives, such as Protection of critical OT assets, Minimising downtime, Preventing data loss, the SP has a specific objective to Preserve operational continuity. This objective focuses on the development procedures to ensure the continuity of mission-critical OT processes during and after disruptions.

SP 800-82r3 3.3.9. outlines the steps to develop recovery and restoration capability to recover from cybersecurity incidents and to restore the assets and services that were impaired by the cybersecurity incident to a pre-cyber- incident state.

This capability typically includes the following tasks:

- Define recovery objectives when recovering from disruptions. For example, the recovery capability shall prioritise human safety and environmental safety prior to restarting the OT operation that was impaired by the cybersecurity event.
- Develop a site Disaster Recovery Plan (DRP) and BCP to prepare the OT organisation to respond appropriately to significant disruptions in their operations due to the cybersecurity incident. Establish backup systems and processes to back up the relevant OT systems' state, data, configuration files, and programs at regular intervals to support recovery to a stable state.
- Establish processes for restoring relevant OT systems' state, data, configuration files, and programs from backups in a timely manner.
- Establish recovery processes and procedures that will be executed to restore the OT assets and services affected by cybersecurity incidents.
- Establish communication plans to coordinate restoration activities with internal and external stakeholders and the executive management team.
- Establish communication plans to manage public relations.
- Establish a task for lessons learned as part of the recovery process for continuous improvement of the cybersecurity capabilities (e.g., vulnerability management, cybersecurity operation, incident response handling, and recovery handling).
- Test these plans at reasonable intervals that are appropriate for the organisation.

NIST SP 800-82r3 emphasises the importance of integrating OT BCP into the organisation's overall BCMS while aligning the OT BCP with the organisation's overall BCM framework to ensure a cohesive approach to business resilience. This involves the establishment of clear communication channels and coordination mechanisms between OT and IT teams, as well as other relevant stakeholders. It also points out that the organisation must allocate sufficient resources for OT BCP development, implementation, and maintenance to ensure its effectiveness.

Additionally the SP stresses the importance of continuous improvement in OT BCP through regular reviews and updates to reflect changes in OT systems, threats, and organisational priorities. Through this process it is important to document lessons learned from exercises and incidents to continuously enhance the BCP's effectiveness. It is also important to integrate emerging technologies, such as automation and analytics, to enhance the efficiency and effectiveness of OT BCP processes.

10.4 Summary

As listed in Figure 6, ISO 22301 provides a general framework for business continuity, while ISO 27031 focuses specifically on ICT readiness. NIST SP 800-82r3 goes a step further by providing OT-specific guidance on developing and implementing a robust BCP for OT environments.

The choice of which standard to use will depend on the specific needs of the organisation. If the organisation is primarily concerned with business continuity, then ISO 22301 is the right choice. However, if the organisation is also concerned with ICT readiness, then ISO 27031 is the right choice.

NIST SP 800-82r3 goes a step further by providing OT-specific guidance on developing and implementing a robust BCP for OT environments.

Feature	ISO 22301	ISO 27031	NIST SP 800-82r3
Scope	Overall business continuity, including OT	ICT readiness for business continuity, specifically focusing on IT systems and infrastructure supporting OT	OT-specific business continuity, addressing the unique security and operational challenges of OT environments
Focus	Maintaining business operations, including OT	Recovering ICT services and infrastructure, including those supporting OT	Protecting critical OT assets, minimising downtime, preventing data loss, and preserving operational continuity
Requirements	More detailed and prescriptive, with specific guidelines for developing and implementing BCP procedures	More flexible and adaptable, allowing organisations to tailor BCP strategies to their specific IT environment	OT-specific and prescriptive, providing guidance on addressing cyber security threats and disruptions specific to OT systems
OT Considerations	Provides general guidelines for addressing OT security in the context of overall business continuity	Limited OT-specific guidance, focusing on ICT security measures for OT systems	Comprehensive OT security guidance, addressing risk assessment, control implementation, incident response, and recovery strategies

Figure 6: Comparison table of ISO 22301, ISO 27031, and NIST SP 800-82r3

Figure 7 is a summary of the key OT-specific considerations for each standard listed.

Standard	OT-specific Considerations
ISO 22301	Consider the unique security and operational requirements of OT systems and processes
ISO 27031	Implement additional security controls to protect OT systems from cyber threats
NIST SP 800-82r3	Conduct regular risk assessments and implement appropriate security controls to mitigate identified risks
All Standards	Integrate OT BCP with the organisation's overall BCMS
All Standards	Ensure that OT personnel are trained in BCP procedures and incident response
All Standards	Continuously review and update OT BCP to reflect changes in threats, technologies, and organisational priorities

Figure 7: Summary of the key OT-specific considerations for each standard

11 Business Continuity Lifecycle

The Business Continuity Lifecycle is a five-step plan.

11.1 Analysis of Business

This initial phase if the lifecycle is an opportunity to look critically at the business to identify vulnerabilities. All-important processes within the business and between the business and customers will need to be evaluated for vulnerabilities.

11.2 Assessment of Risk

Having identified the possible risks it is now a matter of identifying the likelihood of the risk ever occurring and what the impact of such a risk on the business should it ever happen.

Figure 8: Business Continuity Lifecycle

11.3 Develop a Business Continuity Strategy

For each risk a strategy will be needed, such strategies will normally be determined by available budget, either:

- Accept the risk.
- Accept the risk but get a business continuity partner who can help in the event of an incident.
- Reduce the risk.
- Reduce the risk but get a business continuity partner who can help in the event of an incident.
- Reduce the risk adequately that a business continuity partner is not necessary.

11.4 Develop a Business Plan

Now that risks have been identified and a strategy to deal with them decided a full business plan will be needed. Such a plan should be simple because employees will need to act quickly and decisively after an incident.

11.5 Rehearse Plan

There is a military maxim that applies at this stage "*Train hard, fight easy*". The plan must be rehearsed so that employees will know exactly what to do in the event of an incident.

12 BCP Process

Figure 9: BCP Process

The BCP Process has the following four steps:

12.1 Project Scope and Planning

- Structured analysis of the whole business from the perspective of crisis management.
- Appointment of a BCP team with Senior Management approval. The team should consist of:
  - Representation from each department with responsibility for the company core systems
  - Representation from support departments
  - IT personnel with technical expertise in the core systems
  - Information Security officer
  - Legal representation with knowledge of the contractual requirements that may impact the plans
  - Senior management representative
- Identification of all resources available to the team for BCP.
- Understanding of the regulatory and legal situation that governs the organisation's response to an major event requiring a business continuity response.

12.2 Business Impact Assessment

The Business Impact Analysis (BIA) is performed to identify the key business processes and technology components that would suffer the greatest financial, operational, customer, and/or legal and regulatory loss in the event of a disaster. The main intent of a BIA is to identify all the critical resources, systems, facilities, records, etc., that are required for the continuity of the business. Additionally, the time it would take to recovery such resources will be identified.

- For each urgent 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 RPO must ensure that the ***Maximum Tolerable Data Loss (MTDL)*** for each activity is not exceeded. For example if the RPO is set to six hours, then backups must be continuously maintained within that time or more often, say every four hours, i.e. a daily backup will not suffice.

The RTO must ensure that the ***Maximum Tolerable Period of Disruption (MTPD)*** for each activity is not exceeded.

12.2.1 Risk analysis

Now that the recovery requirements are defined, an identification and documentation of potential risks should be undertaken. Identifying the risks give the opportunity to review each and define a specific set of work instructions. Here is a list of common risks:

- Terrorism
- Cyber attack
- Sabotage
- Disease
- Fire
- Flood
- Utility outage

12.2.2 Assessment of Likelihood

Now that we have identified risks what is the likelihood of these occurring? For each produce an ***Annualised Rate of Occurrence (ARO)***. How often is it likely that this event will occur in any year?

12.2.3 Assessment of Impact

Should an identified risk actually occur what is the likely impact of the event on the business ? Determine the ***Exposure Factor (EF)*** to the business as a percentage of the Assets Value (AV) and from these figures calculate the ***Single Loss Expectancy (SLE)***:

$$SLE = AV \times EF$$

From the earlier ARO figure, it is a simple matter to calculate the ***Annualised Loss Expectancy (ALE)***:

$$ALE = SLE \times ARO$$

12.2.4 Prioritisation of Resources

Taking all the risks analysed sort them in a descending list ordered by the ALE of each risk.

12.3 Continuity Planning

Figure 10: Continuity Planning

Having identified the risks, the impacts of these risks should they occur and the priority of resources to deal with risks should they occur we must now plan a strategy to minimise the impact that incidents that occur would have on assets. Continuity planning is in a number of levels:

12.3.1 Strategic Level

This is the identification of which of the risks will be considered in the BCP. Some risks for example may be deemed acceptable to the higher management.

12.3.2 Activity Level

At the Activity Level the complexity of interdependencies on services, business processes, data and technologies needs to be analysed and appropriate tactics chosen to address the needs of:

- People, workforce, skills and knowledge
- Premises
- Alternative Sites
- Infrastructure
  - IT Backbone
  - Servers
  - Workstations
- Information
  - Backup off site
- Stakeholders partners and contractors
  - Alternative partners and contractors

12.3.3 Approval of Plan

Support from top-level management is essential otherwise; the plan is very likely to fail under test.

12.3.4 Training

All personnel that may need to be involved in the plan need training and regular exercise in the execution of the plan. Consider mock exercises regularly much like a fire drill to ensure personnel are “on their toes”.

12.4 Documentation

12.4.1 Continuity Planning Goals

A list of goals for continuity planning. It should start with:

“To ensure the continuation of the business in the event of an emergency situation”

Fill out the remaining goals as necessary for the organisation.

12.4.2 Senior Executive Statement

This statement should come from the C-level management to indicate to all employees the importance of the BCP. The statement should reinforce that Business Continuity is every employee’s responsibility. It should also contain language as to the urgency of the BCP.

12.4.3 Timetable

This is the implementation timetable identified by the BCP team and agreed with upper management.

12.4.4 Priority List

This is a statement of the ordered priorities identified in the BIA.

12.4.5 Risk Assessment

This portion of the documentation captures the assessment performed in the BIA, it should include the analysis on each risk. All risks should have a status as to acceptance or mitigation and for the latter the process and provisions to be put in place.

12.4.6 Records

All vital data and records should be identified and the where they will be stored plus the backup processes and procedures for handling them. How and where.

12.4.7 'Action-on' Emergency Incident

Guidelines must be produced to document the immediate 'Actions-on' response procedures to each potential incident. Who is notified?

12.4.8 Change process

The BCP will need to be tweaked from time to time and the document should include the agreed mechanism for such change.

12.5 Testing

The BCP should document a testing programme and timetable designed to find the flaws in the BCP. The testing should include any elements of the plan that have been outsourced to a third party. Testing should include at different timescales the following types of test:

Document check

- BCP documents checked to ensure changes in operations or administrative changes in personnel or contact details for example are updated

Walk through test

- This is a paper exercise where necessary personnel are brought together to go through the plan to see if problems can be anticipated and fixed

Simulation

- Mock emergency created to test personnel and systems. Exercise conditions with little risk

Parallel test

- Run a full test at backup sites while still maintaining operations at the primary site. This is an important test particularly if the backup site is outsourced. It should be performed at least annually

Full

- Shut down primary site and move operations to the backup site. This is an infrequent test as it is complex and expensive to perform, yet it is the best form of test that can be carried out

13 Exercise #2: Business Continuity and Disaster Recovery

You are the Chief Information Security Officer (CISO) for a specialised Pharmaceutical company in Ireland. The company manufactures a combined calcium and vitamin medications for customers who underwent stomach surgery. The company specialises in direct sales and as such your company holds sensitive customer data that should not leave Ireland and strictly cannot leave the European Union (EU).

The company currently operates a pair of cloud hypervisor servers in a data centre in Limerick and a private data centre in Carlow. These servers offer HA Hipervisor to the virtual servers.

- Develop a BIA for 3 key business processes and technologies.
- Carry out a risk assessment and list 5 key risks.
- Order the risks.
- Describe the mechanism used to order the risks.
- Assuming the only risks are the 5 identified and develop a short BCP for them.
- Suggest a test regime for the BCP.

14 Bibliography

[1] ‘CIS Critical Security Controls’. Center for Internet Security. Accessed: Jan. 10, 2024. [Online]. Available: https://www.cisecurity.org/controls/v8
[2] ‘ISO 22301:2019, Security and resilience – Business continuity management systems – Requirements.’ ISO, Geneva, 2019. Accessed: Sep. 10, 2023. [Online]. Available: https://www.iso.org/standard/75106.html
[3] ‘ISO/IEC 27031:2011. Information technology — Security techniques — Guidelines for information and communication technology readiness for business continuity.’ ISO, Geneva, 2011. Accessed: Sep. 10, 2023. [Online]. Available: https://www.iso.org/standard/44374.html
[4] K. Stouffer *et al.*, ‘Guide to Operational Technology (OT) Security’, National Institute of Standards and Technology, NIST SP 800-82 Rev. 3, Sep. 2023. Accessed: Oct. 01, 2023. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-82r3