Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Network segmentation and segregation – Anti-patterns

Table of contents

Introduction **1**

Audience and scope 2

Where network segmentation and segregation break down **2**

Anti-pattern 1 - The approved design is treated as the live environment 3

Anti-pattern 2 - Network location is treated as trusted 4

Anti-pattern 3 - Shared management paths undermine intended separation 5

Anti-pattern 4 - Controlled reachability is treated as controlled risk 6

Anti-pattern 5 - The ability to change policy is treated as rapid containment 8

Case study **10**

Investigation into a network compromise 10

Broader lesson 11

Further questions to consider **12**

Summary **13**

Further reading **13**

Annex **14**

Relevant anti-pattern ISM controls 14

Introduction

Network segmentation and segregation are established security practices, but their effectiveness can weaken as networks, services and business requirements change. New dependencies, inherited access, vendor connectivity, network address reuse and shared management platforms can gradually introduce communication paths that were not part of the original design. Over time, organisations may settle into recurring arrangements that appear workable but no longer provide the intended logical separation between systems and business functions.

This publication describes these arrangements as anti-patterns: repeated but ineffective responses to common network security problems. It examines five selected anti-patterns:

1. treating the approved design as the live environment
2. treating network location as trusted
3. allowing shared management paths to undermine intended separation
4. treating controlled reachability as controlled risk
5. treating the ability to change policy as rapid containment

These anti-patterns are not exhaustive or prioritised. Together, they examine whether organisations understand their current network state, limit inherited trust and shared management reach, recognise the limits of reachability controls and can enforce containment when required.

The publication supports ASD's *modern defensible architecture* by challenging implicit trust, encouraging explicit access decisions and recognising that separation must extend across data, management and control planes. Targeted microsegmentation can help apply these principles around selected workloads, services and resources, while complementing established network boundaries. The publication also draws on relevant Information Security Manual outcomes for network documentation, communication limited to business need, protection of network-management traffic and centralised logging.

Each anti-pattern explains a recurring problem, how it can be identified and the security outcome organisations should seek. It prompts organisations to examine relevant evidence, investigate material discrepancies and verify that authorised network communication remains under positive control. A sanitised case study illustrates how several anti-patterns can combine to create an attack path, while the 'Questions to continue the conversation' section supports technical investigation, prioritisation and continuing assurance. The publication does not prescribe a particular architecture, product or implementation method.

Audience and scope

This publication is written primarily for network, security and architecture professionals responsible for designing, operating or assuring network segmentation and segregation.

This publication focuses on how anti-patterns can undermine intended separation in production. It uses a selected case study to support threat-informed discussion and assurance, but does not prescribe a particular architecture, product or implementation method.

This document forms part of the network segmentation and segregation publication series:

- ***Network segmentation and segregation - Overview***, covering strategic, governance and assurance considerations.
- ***Network segmentation and segregation - Anti-patterns***, covering recurring operational failure conditions, threat-informed case-study analysis and continuing assurance.

Where network segmentation and segregation break down

An anti-pattern is an identified failure pattern: a repeated but ineffective technical implementation or operational response to a common problem. In network segmentation and segregation, anti-patterns are recurring conditions that undermine intended separation in production, often through legitimate operational change, accumulated complexity or inherited trust relationships rather than deliberate control failure.

How to use the indicators: one indicator may justify investigation, but no fixed number proves an anti-pattern exists. Consider the indicators together, assess their significance in your environment and confirm the condition using current technical evidence.

The following five anti-patterns are not exhaustive or prioritised. They were selected because they expose distinct failures in understanding current network state, constraining inherited trust and shared management reach, recognising the limits of reachability controls, and achieving timely containment. Together, they help organisations identify where the operating network has diverged from approved architecture, where apparently effective controls have become weakened or ineffective, and what evidence, accountability and remediation activities are needed to restore and maintain positive control over authorised network communication.

Anti-pattern 1 - The approved design is treated as the live environment

What is the anti-pattern?

Approved diagrams are treated as evidence of production implementation. Routine operational changes, vendor connectivity, temporary workloads and legacy dependencies can create communication paths that are not represented in approved designs or leave obsolete paths in place after their original purpose has passed. Current diagrams and configuration snapshots may also fail to represent historical network state, particularly where temporary interfaces, routes, tunnels or policy changes are subsequently removed.

Asset inventories and IP address management (IPAM) systems help establish which assets use particular addresses and who is accountable for them. However, IP addresses, hostnames and other identifiers can be reassigned over time while firewall rules, network objects or trust relationships remain unchanged. The control appears stable because the rule has not changed, but its effective subject has. This can allow inherited access, unintended trust relationships or undocumented communication paths to persist without being recognised through architecture documentation or configuration review alone.

How to identify this anti-pattern

- Network diagrams, connection registers, asset inventories and security-zone documentation have not been reviewed or updated often enough to keep pace with change and no longer reflect the current topology, routing behaviour, trust relationships or permitted communication paths.
- IPAM and asset management records do not reliably identify the system, owner, environment, business function and lifecycle status associated with each address, subject or network object.
- Address reassignment, system decommissioning, vendor offboarding and topology changes do not trigger review of associated firewall rules, object groups, routes and documented dependencies.
- Firewall rules, access control lists (ACLs) or routing policies remain active after the related system, service or business function has been decommissioned, or have no identifiable owner, current business purpose or recent legitimate traffic.
- Temporary interfaces, routes, tunnels, routing advertisements or policy changes are not retained in configuration history, change records or other evidence required to reconstruct historical network state.
- Rules, routes or network objects that appear obsolete cannot be removed because the dependent applications, communication flows, integrations and service owners are unknown.
- The organisation cannot reliably determine what communication paths existed at a previous point in time, or whether inherited access persisted after the original business need ceased.

How to address this anti-pattern

Establish a process that continuously validates the documented network against the operating network, rather than relying on periodic manual reviews. Use network management, firewall management, IPAM records, configuration management database (CMDB), and flow monitoring tools to automatically correlate asset inventories, network configurations, routing information, firewall policies, and observed communications. Configure key events, such as reassignments, domain name system (DNS) changes, new routes, firewall rule creations, system onboarding, and system decommissioning, to trigger review of associated network objects, firewall rules, dependencies and authorised communication paths. Investigate material differences between the documented and operating network, assign accountable owners and remediation dates, and record any residual risk requiring acceptance.

Anti-pattern 2 - Network location is treated as trusted

What is the anti-pattern?

Users, devices, workloads or vendors inherit access because they are located within an approved network zone, corporate network address range or remote-access pathway. Network location can inform an access decision, but it should not establish trust or grant broad access by itself.

ASD's *modern defensible architecture* promotes explicit access decisions based on verified identity, device posture, context and business need, rather than treating network location as trusted. Microsegmentation is an extension of network segmentation that enforces policy at a much finer granularity, such as the individual workload, application, service or device, rather than broad network zones. Its primary purpose is to limit attack propagation and reduce the consequences of compromise by restricting communications to specifically authorised interactions. Targeted microsegmentation can help confine access to well-understood and validated dependencies where enforcement can be tested and continuously assured. It complements, rather than replaces, established zones, routing, gateways and firewall controls.

How to identify this anti-pattern

- Remote Desktop Protocol (RDP) access to domain controllers is granted to anyone connected to the "IT Virtual Local Area Network (VLAN)" rather than to specific administrative systems.
- An always-on site-to-site virtual private network (VPN) from a managed service provider, vendor or partner is treated as trusted, allowing broad access to internal systems.
- Firewall rules allow more access than applications need.
- Whole networks are allowed access when only a few systems need to communicate.
- Being on a particular VLAN, subnet, or jump host is enough to gain access without additional destination-side authorisation.
- Alternative or legacy network paths exist that bypass normal security controls.
- Network monitoring cannot clearly distinguish expected traffic from unusual or suspicious traffic.

How to address this anti-pattern

Define authorised communication paths based on current business needs and system criticality. Remove access that exists because of legacy designs, broad address ranges, or historical trust relationships, and test proposed restrictions before enforcement. Where a legacy system, supplier connection, or service dependency prevents a path from being removed, record the path as an exception with an accountable owner, business reason, risk treatment, review date and planned removal or remediation date. Assess actual communication paths, including legacy, vendor, administrative and bypass routes, rather than relying solely on the intended routes shown in the design. Organisations need to be able to identify unnecessary network reachability, explain why it exists, and demonstrate when and how it will be reduced or eliminated.

Anti-pattern 3 - Shared management paths undermine intended separation

What is the anti-pattern?

Management systems such as identity services, privileged-access tools, software-deployment platforms, monitoring systems and vendor access solutions often need to reach multiple environments. When these systems have more access than they need, or are not adequately protected, they can create pathways across security boundaries. They may also allow malicious actors to bypass intended separation entirely, making management systems high-value targets.

Systems may appear separate but still depend on the same management network or management infrastructure. This means that a malicious actor who compromises a shared management system may be able to access, manage, or disrupt multiple environments, even if those environments are otherwise segregated. Out-of-band (OOB) management can help contain and recover from incidents, but only if it remains independent of the systems it is intended to protect. If it relies on the same accounts, identity services, remote-access solutions or management platforms as the production environment, it may be unavailable when needed most.

If shared management infrastructure is compromised, a malicious actor may be able to change configurations, modify security policies, deploy software or alter logs across multiple environments used to investigate the incident. As a result, a compromise of a single administrative platform can undermine the intended separation between systems.

How to identify this anti-pattern

- An edge device, such as a firewall or remote-access gateway, can initiate connections to management interfaces of internal systems.
- A management system can administer systems across multiple critical network zones or environments.
- There are management paths between environments that are not shown in the network segmentation and segregation design documentation.

- Administrative authority is broader than operational need, including standing supplier accounts, shared privileged identities or access to entire device groups.
- OOB management shares credentials, identity services, routing, remote-access gateways or other dependencies with the environment it is intended to recover.
- The OOB management path has not been tested following significant network, firewall, authentication, or platform changes.
- Multiple critical environments depend on the same management, authentication, monitoring or recovery platform without sufficient redundancy or an independently tested fallback.
- Management or audit logs are stored only on systems they are intended to monitor.
- Shared management platforms, administrative paths, or vendor connections do not have a clearly accountable owner or regular review process.

How to address this anti-pattern

Map management and vendor paths as part of the segregation design. Prevent edge devices, vendor connections and less-trusted systems from directly reaching the management interfaces of internal systems unless explicitly required. Constrain each management platform to the devices and network zones it must administer, and separate management traffic from application traffic using dedicated management networks, VLANs, private VLANs, security zones or equivalent controls.

Require authenticated management connections and apply role-based or attribute-based access controls to limit administrative actions according to role, device, service, environment and operational need. Assign accountable owners for management platforms, cross-zone administrative paths and vendor connections, and regularly review their reach, privileges and exceptions.

Retain management and audit logs outside the managed environment, monitor material configuration and software changes, and regularly test whether critical systems can be contained and recovered if shared management infrastructure becomes unavailable or untrusted.

Anti-pattern 4 - Controlled reachability is treated as controlled risk

What is the anti-pattern?

Network diagrams, documented traffic flows and connection testing are used to determine which systems can communicate across network boundaries. These checks can confirm that required connections are permitted and prohibited connections are blocked. However, they do not show whether traffic on an authorised path can be observed, copied, modified, redirected or disrupted.

An actor controlling a device on an authorised path may interfere with traffic without creating a new or prohibited connection. Legitimate communication may continue, and tests may confirm that blocked paths remain blocked, while the confidentiality, integrity or availability of traffic on permitted paths is affected.

The firewall or other enforcement point may therefore operate exactly as configured. The anti-pattern is treating that result as evidence that all risks associated with permitted traffic are also controlled.

How to identify this anti-pattern

- Network threat models and assurance testing assess whether a source can establish a connection to a destination system, but do not evaluate how permitted traffic could be observed, copied, modified, redirected, delayed, replayed or discarded while in transit.
- Successful testing of blocked network paths is incorrectly treated as evidence that confidentiality, integrity and availability of communications are protected across permitted network paths.
- Switch Port Analyzer (SPAN) or Encapsulated Remote SPAN (ERSPAN) sessions, packet capture mechanisms (for example, tcpdump or Wireshark), network tunnels (such as IPsec, SSH tunnels), policy-based routing configurations and other traffic-redirection mechanisms are not inventoried, formally approved or maintained within an authorised configuration baseline.
- Changes to dynamic routing protocols (for example, Open Shortest Path First (OSPF), Border Gateway Protocol (BGP)), forwarding tables, tunnel endpoints, traffic-mirror destinations, policy-routing rules or device-originated traffic flows are not continuously monitored and validated against an approved network state.
- Unexpected packet-capture activity, traffic-mirroring configurations, tunnel creation, routing changes, forwarding-path modifications, or redirection mechanisms are treated solely as operational or availability events and do not trigger security investigations to determine potential unauthorised observation, manipulation or exfiltration of network traffic.
- Evidence from network device configurations, centralised logging and telemetry platforms (such as SIEM, syslog, NetFlow), and observed network traffic is not correlated to identify discrepancies between expected and actual device behaviour, including unauthorised forwarding, mirroring, tunnelling or routing activities.

How to address this anti-pattern

Clearly define and document what the segmentation and segregation control demonstrates and where its assurance boundary ends. Network segmentation and segregation can demonstrate that communications are permitted or denied in accordance with policy, but it does not itself assure the confidentiality, integrity or availability of traffic traversing permitted network paths. Develop and maintain a documented framework for network attack-path analysis that identifies potential

paths for access, lateral movement, traffic observation, manipulation, redirection and privilege misuse, and retain analysis outputs as security evidence.

Supplement connectivity, segmentation, and segregation assessments with consequence-based threat analysis covering access, traffic observation, manipulation, disruption and compromise of policy, configuration or administrative authority. Monitor for unexpected changes to packet-capture, traffic-mirroring, tunnel, routing and forwarding configurations and investigate such changes as potential security events. Protect traffic using appropriate end-to-end controls where practicable, validate the integrity of network devices and management infrastructure and retain logs, telemetry and configuration evidence independently of the components being continuously assessed.

Anti-pattern 5 - The ability to change policy is treated as rapid containment

What is the anti-pattern?

The organisation assumes it can rapidly restrict network communications during an incident because it possesses policy management tools, privileged accounts, orchestration platforms or emergency change procedures. The ability to request or initiate a change is treated as equivalent to achieving effective containment.

In practice, operational, governance and technical constraints can delay containment. There may be no pre-approved response playbook, emergency change processes may introduce delays, or required approvers may be unavailable. Once authorised, the intended action must still be translated into network policy, validated, distributed, applied consistently across all relevant enforcement points, and independently verified. Delays, failed devices, management-plane disruption, configuration errors, policy-distribution failures, or untested emergency procedures can leave communication paths open, lock administrators out of critical systems, or cause unintended operational outages.

How to identify this anti-pattern

- Incident, outage and recovery plans identify what needs to be isolated but do not define the specific rule, route, zone, virtual routing and forwarding (VRF) or administrative-path changes required to achieve containment.
- Incident containment is not exercised through the complete workflow, from authorised decision through policy creation, deployment and independently verified enforcement across the relevant network.
- Deployment time is measured until a policy push is issued rather than until every relevant firewall, gateway and other enforcement point has successfully applied and enforced the change.

- Emergency changes depend on the same identity, orchestration and management and control-plane paths that may be unavailable, degraded or untrusted during an incident or operational failure.
- OOB and break-glass access mechanisms have not been exercised under time pressure or following significant routing, firewall, authentication or management-platform changes.
- Stale, unreachable, incompatible or failed enforcement points are excluded from containment reporting rather than recorded and managed as unresolved exposure.
- Containment playbooks do not preserve required administrative access, validate rollback behaviour, or account for asymmetric routing, service dependencies and shared infrastructure.
- Exercises test detection, response and recovery activities but do not measure containment convergence time, failed enforcement points, policy-deployment exceptions, or communication paths that remain available during change propagation.

How to address this anti-pattern

Treat emergency enforcement agility as a measurable security capability. Develop and maintain a small set of bounded, authorised and regularly exercised end-to-end containment playbooks for foreseeable scenarios. Protect the integrity and provenance of each playbook, validate it against current topology, network policies and service dependencies, preserve required management access, and verify rollback behaviour through exercises.

Measure the complete interval from an authorised containment decision to independently verified enforcement across all relevant enforcement points, including approval, policy preparation, distribution, convergence, exceptions and failed devices. Verify the live outcome using evidence independent of the orchestration platforms or enforcement points reporting successful deployment and treat any enforcement point that fails to implement the change as continuing exposure until independently resolved.

Case study

Investigation into a network compromise

An Australian Cyber Security Centre (ACSC) investigation identified that a malicious actor had compromised an internet-facing network edge device using legitimate network capabilities to extend and maintain access within the environment. The investigation identified activity affecting device configuration, authentication mechanisms, routing, packet capture, management connectivity and local audit records.

The malicious actor gained access to the device's running configuration through an exposed network-management service, established remote administrative access, and modified authentication and access-control settings. The malicious actor subsequently configured a temporary loopback IP address associated with an existing jump server, advertised the address through OSPF, and used existing management paths to access additional network infrastructure.

The following sanitised sequence highlights the network segmentation and segregation activity observed during the investigation. It does not represent the complete chronology of events and omits victim-specific technical detail, environmental context and other aspects of intrusion not directly relevant to the network security lessons discussed.

- T1133 - **External Remote Services:** The malicious actor changed the running configuration of an internet-facing network edge device and established remote administrative access.
- T1556 - **Modify Authentication Process:** The malicious actor changed the AAA method order so local authentication was evaluated before TACACS+. Legitimate administrative access remained available, making the change less apparent during cursory review.
- T1599 - **Network Boundary Bridging, assumed a trusted network position:** The malicious actor configured a temporary loopback interface using the same IP address as an existing jump server.
- T1599 - **Network Boundary Bridging, advertised the trusted address:** The actor advertised the temporary loopback address into an existing OSPF area and configured it as the source of outbound SSH traffic. This allowed downstream ACLs to evaluate the traffic as originating from an address associated with the jump server.
- T1599 - **Network Boundary Bridging, used a management routing context:** The malicious actor enumerated routes and reachable destinations in a management VRF and used the compromised network device to send ICMP, SSH and Telnet traffic towards additional network infrastructure.
- T1040 - **Network Sniffing:** The malicious actor captured encrypted TACACS+ authentication traffic traversing network devices. The investigation assessed that the actor likely obtained the shared key from device configuration, decrypted authentication packets and compromised additional credentials.

- **T1021 and T1021.004 - Remote Services: SSH:** The malicious actor used compromised service and administrator credentials to attempt connections to other network devices and supporting systems. The investigation confirmed successful lateral movement in some instances, while other records established only that a connection was attempted.
- **T1070.003 - Indicator Removal: Clear Command History:** The malicious actor withdrew the OSPF advertisement, removed the temporary loopback interface and cleared local logging. The temporary state would no longer appear in the running configuration, but centrally forwarded command-accounting records preserved evidence of the activity.

Broader lesson

The case study shows how multiple anti-patterns can combine to undermine the security objectives of network segmentation and segregation. A temporary loopback interface and OSPF advertisement created network state that was not represented by the approved design or architecture documentation. Subsequent removal of the configuration and deletion of local logs reduced the evidence on the affected device, while centralised command-accounting and logging systems preserved an independent record of the activity.

The malicious actor sourced traffic from an address associated with an existing jump server, causing downstream access control lists (ACLs) and other trust decisions to treat the traffic as originating from a trusted network location. Management VRFs and permitted remote-administration services then provided communication paths to additional infrastructure beyond the initially compromised device.

The case demonstrates four anti-patterns: the approved design is treated as the live environment; network location is treated as trusted; shared management paths undermine intended separation; and controlled reachability is treated as controlled risk. It does not directly demonstrate rapid containment.

Assurance activities should cover the data, management and control planes and assess whether the organisation can detect consequential configuration, routing and trust-boundary changes, reconstruct historical network state, identify deviations from approved architecture, and retain evidence independently of the devices and management systems being administered.

Further questions to consider

Network, security and architecture professionals may like to consider and use the following questions to continue the conversation within their organisation to support network segmentation and segregation implementation.

- Does the approved architecture accurately reflect the communication paths, dependencies and trust relationships that exist in the production environment today?
- Do the asset inventory and IPAM records reliably identify what is connected to the network, who is accountable for each asset and where it is expected to reside?
- Could configuration changes, temporary workarounds, legacy dependencies or operational practices create undocumented communication paths or leave obsolete access in place?
- Which security controls make trust or access decisions primarily on the basis of source IP address, network location or network segment membership?
- If a network management, monitoring, orchestration or automation platform were compromised, what systems, networks and administrative functions could a malicious actor subsequently access?
- When was the last time it was verified that the OOB management environment remained operational, independent and trustworthy during a simulated compromise or major service disruption?
- Can the Network Intrusion Detection Systems (NIDS) or Network Intrusion Prevention Systems (NIPS) adequately observe traffic traversing critical trust boundaries, management paths and critical network segments?
- Would unexpected changes involving OSPF, BGP, Multi-Protocol Label Switching (MPLS), VRF, routing policies or traffic-forwarding behaviour generate a security alert and trigger investigation?

Summary

Network segmentation and segregation can become less effective as production environments evolve. Approved architectures, diagrams and configuration snapshots may not accurately represent the communication paths, dependencies, inherited access or temporary network state that exist in practice.

Asset inventories and IPAM records provide important context, but they may not reflect changes in the effective subject of a rule or trust relationship over time. Network location alone should not be treated as a basis for trust, and shared management infrastructure can create communication paths across environments that otherwise appear separated.

These conditions can compound, allowing a malicious actor to alter their effective network position and exploit legitimate routing, management and administrative paths to extend access.

Organisations need to assess the data, management and control planes as a holistic system, reduce reliance on implicit trust, and verify that communications remain constrained to legitimate business needs. Consistent with ASD's *modern defensible architecture*, targeted microsegmentation can help operationalise Zero Trust by applying explicit, granular controls that are continuously validated around critical systems, services and management functions.

Further reading

- ASD's *Implementing network segmentation and segregation*
- ASD's *modern defensible architecture*
- ASD's *Information Security Manual*
- ASD's *mitigations for network defence*
- ASD's *CI Fortify: Advice for isolating vital systems*
- NIST's SP 800-215, Guide to a Secure Enterprise Network Landscape
- NIST's SP 800-207, Zero Trust Architecture
- UK National Cyber Security Centre's Zero trust architecture design principles: Common anti-patterns
- UK National Cyber Security Centre's Security architecture anti-patterns
- NSA's Network Infrastructure Security Guidance
- NSA's Network Device Integrity Methodology

Annex

Relevant anti-pattern ISM controls

Anti-pattern 1 - The approved design is treated as the live environment

- ISM-0518; Network documentation is developed, implemented and maintained.
- ISM-0516; Network documentation includes high-level network diagrams showing all connections into networks and logical network diagrams showing all critical servers, high-value servers, network devices and network security appliances.
- ISM-1182; Network access controls are implemented to limit the flow of network traffic within and between network segments to only that required for business purposes.
- ISM-1912; Network documentation includes device settings for all critical servers, high-value servers, network devices and network security appliances.

Anti-pattern 2 - Network location is treated as trusted

- ISM-1181; Networks are segregated into multiple network zones according to the criticality of servers, services and data.
- ISM-1182; Network access controls are implemented to limit the flow of network traffic within and between network segments to only that required for business purposes.

Anti-pattern 3 - Shared management paths undermine intended separation

- ISM-1006; Security measures are implemented to prevent unauthorised access to network management traffic.
- ISM-1577; An organisation's networks are segregated from their service providers' networks.
- ISM-1963; Security-relevant events for internet-facing network devices are centrally logged.
- ISM-1964; Security-relevant events for non-internet-facing network devices are centrally logged.

Anti-pattern 4 - Controlled reachability is treated as controlled risk

- ISM-1203; System owners, in consultation with each system's authorising officer, conduct a threat and risk assessment for each system.
- ISM-0009; System owners, in consultation with each system's authorising officer, identify any supplementary controls required based upon the unique nature of each system, its operating environment and the organisation's risk tolerances.

- ISM-1912; Network documentation includes device settings for all critical servers, high-value servers, network devices and network security appliances.
- ISM-1963; Security-relevant events for internet-facing network devices are centrally logged.
- ISM-1964; Security-relevant events for non-internet-facing network devices are centrally logged.

Anti-pattern 5 - The ability to change policy is treated as rapid containment

- ISM-0043; Systems have a cyber security incident response plan that covers the following:
  - guidelines on what constitutes a cyber security incident;
  - the types of cyber security incidents likely to be encountered and the expected response to each type;
  - how to report cyber security incidents, internally to an organisation and externally to relevant authorities;
  - other parties that need to be informed in the event of a cyber security incident;
  - the authority, or authorities, responsible for investigating and responding to cyber security incidents;
  - the criteria by which an investigation of a cyber security incident would be requested from a law enforcement agency, the Australian Signals Directorate or other relevant authority;
  - the steps necessary to ensure the integrity of evidence relating to a cyber security incident; and
  - system contingency measures or a reference to such details if they are in a separate document.
- ISM-0912; Systems have a change and configuration management plan that includes:
  - the establishment and maintenance of authorised baseline configurations for systems;
  - what constitutes routine and urgent changes to the configuration of systems;
  - how changes to the configuration of systems will be requested, tracked and documented;
  - who needs to be consulted prior to routine and urgent changes to the configuration of systems;
  - who needs to approve routine and urgent changes to the configuration of systems;
  - who needs to be notified of routine and urgent changes to the configuration of systems; and

    - what additional change management and configuration management processes and procedures need to be followed before, during and after routine and urgent changes to the configuration of systems.

- ISM-1784; The cyber security incident management policy, including the associated cyber security incident response plan, is exercised at least annually.
- ISM-1610; A method of emergency access to systems and their resources is documented and tested at least once when initially implemented and each time fundamental information technology infrastructure changes occur.

For more information, refer to ASD's *Information Security Manual*.

Disclaimer

The material in this guide is of a general nature and should not be regarded as legal advice or relied on for assistance in any particular circumstance or emergency situation. In any important matter, you should seek appropriate independent professional advice in relation to your own circumstances.

The Commonwealth accepts no responsibility or liability for any damage, loss or expense incurred as a result of the reliance on information contained in this guide.

Copyright

Use of the Coat of Arms

The terms under which the Coat of Arms can be used are detailed on the Department of the Prime Minister and Cabinet website
https://www.pmc.gov.au/resources/commonwealth-coat-arms-information-and-guidelines

Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre