Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Gateway Operations and Management

Gateway Security Guidance Package

First published: July 2022

Last updated: July 2025

Table of contents

Introduction

This document provides practical guidance on the secure operation and management of gateways. It is intended for engineers, operations and support teams to understand better practice approaches for operating, maintaining and disposing of gateways. It builds upon Gateway Security Guidance Package: Gateway Security Principles, which outlines the fundamental expectations for secure gateway design and operation. By providing operational context to these principles, the guidance helps organisations translate high-level security concepts into practical day-to-day processes. It offers best practices for configuring, maintaining, monitoring, and validating gateway environments to ensure they remain secure, resilient, and aligned with broader risk management strategies.

Gateway operations and management

A gateway's relative security capability and effectiveness will erode over time, typically impacting on controls, visibility, performance and comparative capabilities, unless gateways are properly maintained and operated. A variety of factors influence the ability of a gateway to operate as intended, including:

- support for new or updated standards, protocols, or functionality
- the implementation of exemptions to security policies
- advancements in malicious actor tradecraft
- challenges retaining critical corporate knowledge and capabilities
- the discovery of new vulnerabilities
- operator errors
- system load
- external third party dependencies and environmental factors.

Organisations need to continuously assess the effectiveness of their gateways to ensure that:

- security controls remain effective against current malicious actor tradecraft
- the gateway is supporting business requirements in a cost-effective way
- vulnerabilities are identified and remediated
- risks associated with configuration drift are monitored and managed
- new security features are enabled.

Organisations need to understand risks related to their supply chain and should consider Choosing secure and verifiable technologies guidance, which is designed to assist procuring organisations (including those procuring gateway services or products) to make informed, risk-based decisions within their own operational context. Organisations should also consider Foundations for modern defensible architecture, particularly Foundation 4: Reliable asset inventory, when considering supply chain risks.

Gateways help organisations implement a range of cyber security capabilities, such as deep packet inspection (DPI), on-demand packet capture, real-time reputation assessments, dynamic content categorisation and endpoint device posture validation. Gateways also provide the ability to ingest and action cyber threat intelligence (CTI) and to perform authentication.

Security standards

Organisations should implement documented better practice suggested by software and hardware vendors. Vendors should have comprehensive documentation on how to configure settings to achieve desired security outcomes, and information about how to verify that controls are in place and operating effectively. Where there is no direct government security configuration guidance, organisations should refer to vendor guidance and industry better practice guidance obtained from reputable sources.

There are various information security assurance frameworks, standards and hardening guidance organisations should consider using when assessing the cyber security posture of their gateways, such as:

- Infosec Registered Assessors Program (IRAP) assessments
- AS/NZS ISO/IEC 27000 suite
- NIST Cyber security Framework (CSF)
- Common Criteria (with appropriate Network Device Protection Profile)
- comprehensive penetration tests
- vulnerability assessments and vulnerability scanning
- Security Content Automation Protocol (SCAP) and Security Technical Implementation Guides (STIGs)
- cloud and platform blueprints and desired state configurations
- CIS Benchmarks (Centre for Internet Security)
- Industry better practice guides from cloud providers and vendors.

When procuring products or services that perform a function enforcing security policies, organisations should consider selecting products that are secure-by-design and secure-by-default. This includes choosing products manufactured by reputable suppliers that regularly undertake rigorous and independent security testing, such as through the Common Criteria (CC) using relevant protection profiles.

Vendor products targeting enterprise customers will often include documentation for configuration hardening. Organisations can streamline ongoing configuration validation processes using automated methods. For example, by using SCAP and STIGS, or cloud security blueprints. Refer to the 'Platform hardening' section of this guidance.

Vendors who understand the security requirements of enterprise customers typically have technical guidance, reference architectures and specialist training on how to securely design, deploy, and manage their platforms and services. Organisations should ensure that architecture, engineering, and operational teams have adequate training, and apply both vendor advice and this Gateway Security Guidance Package when designing, building and operating systems. In addition to implementing vendor better practice guidance, it is important for products performing a policy

enforcing role to have also undergone thorough and independent testing through processes such as the CC. An organisation's management and operations processes should align with broader system development lifecycle strategies to ensure that gateway systems retain, or enhance, their secure configuration over time.

An organisation might choose to use security evaluated products (such as CC) in their gateway because:

- They use that security product extensively elsewhere and would like to standardise on that product's control and visibility capabilities (e.g. a single vendor firewall management and centralised logging capability).
- The security value of the information means that the organisation wants to apply transparently proven and rigorously tested controls.
- The product supports additional visibility and/or verification of behaviour beyond the vendor's products, assertions and guarantees.
- The organisation needs a security function to assist implementation of a non-native business process or processing logic capability, or to centrally manage security policies in hybrid or multi-cloud environments.

Architects and engineers should develop threat models to identify what potential attacks that could be effective if gateway controls fail. Through this planning, they can develop an understanding of what logging and telemetry is necessary to identify normal and abnormal traffic flowing through the gateway. Organisations assessing the degree of risk of a control failure, and the consequences of such a failure, may determine a Cross Domain Solution (CDS) is necessary in a gateway system.

Continuous assurance and validation

A continuous monitoring plan can assist an organisation in proactively identifying, prioritising and responding to vulnerabilities and events that may impact on gateway efficiency. Continuous monitoring includes activities, such as:

- identifying unpatched devices
- validating deviations from configuration baselines
- applying vendor hardening guidance and better practices
- implementing continuous improvement initiatives
- reviewing system event logs
- using newly released security capabilities from vendors or cloud providers.

These activities apply across both traditional on-premises and cloud-based gateway environments. For cloud-based environments, this includes monitoring through native tooling and ensuring alignment with the shared responsibility model.

As part of a continuous assurance program, organisations should confirm that controls are in place and continue to operate effectively over time. By developing and executing comprehensive tests to validate the gateway security controls beyond traditional point-in-time audits, organisations can gain a higher level of assurance that risks are being effectively monitored and managed. Regular threat modelling activities focused on organisational systems and processes can identify relevant attack vectors. These vectors that can be addressed through new or enhanced mitigations, which should be tested using automated or manual controls mechanisms. More information on security continuous monitoring can be found on NIST's SP 800-137: Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations.

The ability of an organisation to anticipate, identify and respond to cyber security incidents will be dependent on many factors. For further guidance on incident response planning, refer to ASD's *Information Security Manual* (ISM) Guidelines for cybersecurity incidents and NIST's SP 800-61r3 Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile.

Threat modelling and other incident detection techniques can be used to help an organisation develop a suite of continuous validation tests for gateways. The use of automation to validate security configuration (e.g. SCAP and STIGs), combined with organisation and system-specific unit tests, can be used to identify where configuration is no longer 'in pattern', highlighting where security policy enforcement is no longer working effectively, or where there is increased exposure to risk.

Continuous validation frameworks and tools

The use of automation – including Compliance-as-Code, Open Security Controls Assessment Language (OSCAL) and SCAP, Infrastructure-as-Code (IaC) and Continuous Integration / Continuous Development (CI/CD) pipelines – can increase assurance that a gateway is operating as designed over the life of the system. In traditional on-premises infrastructure environments, gateway teams may need to work with other operational and business teams to conduct some of these tests.

OSCAL provides an automation mechanism to assess systems against a range of security control catalogues, such as the ISM, configuration baselines, system security plans, assessment plans and results. ASD provides the ISM in the OSCAL format.

Examples of frameworks, processes, tools and techniques that assist with continuous validation include:

- SCAP, using Extensible Configuration Checklist Description Format – Extensible Configuration Checklist Description Format (XCCDF) and system-specific STIGS, to test infrastructure or software implementations against vendor guidance
- automated scripts to test security and operational functionality against a known baseline
- periodic audit and assessment activities (IRAP, ISO 27001:2022, ISO 9001:2015)
- penetration testing by suitably skilled, experienced and reputable personnel
- automated vulnerability scanning, with the potential integration of AI enhancement and/or enrichment where possible
- manual validation techniques (e.g. manual testing of firewall rulesets)
- periodic load testing, stress testing, and performance testing
- CI/CD pipeline validation practices.

Organisations need to document policies, processes and procedures for version control of system configurations, as well as software version tracking. For example, documenting operating system and firmware versions, and deployed configurations. Tools that perform this function should be isolated in management environments, with access strictly controlled.

Threat modelling

Architects and engineers should develop threat models to identify potential attacks that could be effective if gateway security controls fail, or gateway sub-systems are compromised. Through this modelling, they can develop an understanding of the logging and telemetry required to identify normal and abnormal traffic flowing through a gateway. Threat modelling is a structured risk assessment process that identifies possible threats or attacks on a gateway system or service, to then inform its design and development with targeted and effective mitigation strategies. A threat model is scoped to a defined boundary, identifies possible threats, provides mitigation strategies and can be validated as part of a continuous process.

Threat modelling as a process should be performed at various points throughout the lifecycle of a gateway system or service. It should be integrated into the development and design of any gateway technology system or service, and embedded, as business as usual, into your risk management processes, upgrade activities, or CI/CD processes. This approach enables architects, developers, engineers, and security analysts to:

- consider the system or service context, its purpose, functionality, and operational mechanism
- adopt an adversarial mindset to anticipate potential threats and vulnerabilities during development
- identify and justify gateway security requirements early
- clearly communicate cyber security risks and their potential impact.

Threat modelling process

While threat modelling can be applied as part of routine operational processes, gateway implementations may benefit from dedicated threat modelling workshops. Theses workshops should be aimed at producing a formal threat model description that informs the design, implementation, and operation of gateways. Participants should include individuals from a variety of different backgrounds and job functions, including business owners and technical subject matter experts. The different views and objectives of the participants will be the drivers for identifying the boundaries and information flows.

Threat modelling is typically performed in a four-step process (**Figure 1** outlines the four steps of the threat modelling process), with each step based on one of these four questions:

- What is the boundary?
- What can go wrong?
- What can be done?
- How successful?

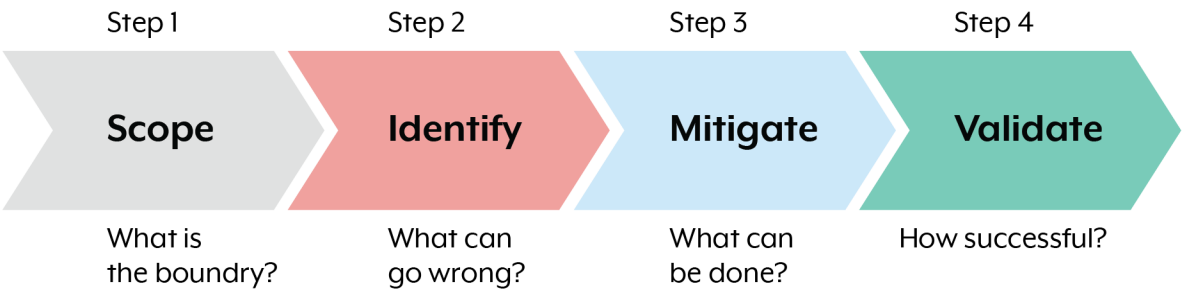

Figure 1: Threat Modelling steps

The first step focuses on defining the scope and developing a high-level understanding of the gateway system, including the flow of information and data, the parties involved and the system boundaries. For multi-tenant gateways, consider how the data flow and trust boundaries are different between the tenants and external connections. The output of this stage should be a diagram or conceptual understanding of the main components and how they interact within different trust boundaries.

The second step involves identifying potential weaknesses and threats to the gateway, and documents how the gateway system and its services could be impacted by malicious actors. There are many ways to approach threat development, including using threat taxonomy databases or by developing attack trees. Frameworks such as PASTA, STRIDE, OCTIVE, and others can be used to build a threat matrix. To triage and categorise threats, consider using Common Attack Pattern Enumeration and Classification, MITRE ATT&CK, and Common Weakness Enumeration. The output of this step should be a list of threats applicable to the gateway solution, including some description of what the impact of those threats would be.

The third step incorporates security controls and architectural principles that minimise or eliminate the realisation of threats, or mitigates their impact. This guidance, along with the broader advice from ASD, such as the ISM, can be used as a resource when selecting appropriate controls. The output of this stage should be the combination of all the previous stages into one document that details the gateway threats and how the solution design will treat those threats through a combination of control and design principles.

The fourth step measures how successful and effective the threat modelling process has been, and serves to validate its outcomes. The threat model is circulated amongst the solution's stakeholders and subject matter experts to communicate the identified threats and mitigations clearly. This step enables collaboration to identify anything that may have been overlooked or missed, and supports continuous monitoring and review of the threat landscape. It also helps highlight areas for improvements and identify ways of measuring the effectiveness of the implemented controls. Based on the feedback received, this step may require revisiting any of the previous three steps to strength the overall model.

The finished threat model from the workshop should be provided as a reference to the necessary resources of the gateway solution design as well as the operational teams. The controls and design principles should be inherited into system security plans and risk management plans. Developing future threat models should begin with reviewing previous threat models. For more information on threat modelling refer to: Threat Modelling Manifesto's Principles, OWASP's Threat Modelling Cheat Sheet, and NIST's 800 154, Guide to Data-Centric System Threat Modelling.

Cyber threat intelligence

CTI should be used as a security mechanism in gateway environments. Gateways should use reputation and dynamic categorisation services and should be able to ingest other sources of CTI. CTI can be implemented in many gateway components, including firewalls, forward and reverse proxies, mail relays, recursive Domain Name System (DNS) resolvers, and virtual private network (VPN) services.

Organisations should use their gateways to derive intelligence from their operation, generating data that allows security analysts to derive CTI. Gateways must support this by forwarding relevant logs, telemetry and data to an organisation's Security Operations Centre (SOC). Organisations should ensure that contract terms for the service allow them to access the gateway-related logs, telemetry, and data of the services provided by their cloud service providers (CSPs) and managed service providers (MSPs).

Cyber Threat Intelligence Sharing Platform

ASD provides a Cyber Threat Intelligence Sharing (CTIS) service that supports a two-way sharing platform that enables government and industry partners to receive and share information about malicious cyber activity at machine speed. Access to the CTIS platform is facilitated through the cyber.gov.au web portal and requires entities to sign a confidentiality deed and enrol in ASD's Cyber Security Partnership Program.

Organisations may consume and contribute to CTI as part of their cyber security operations through the following activities:

- consume CTI that informs them of indicators of compromise (IoC) or malicious behaviour via processes such as Structured Threat Information eXpression (STIX), Trusted Automated exchange Indicator Information (TAXII), Malware Information Sharing Platform (MISP), or other dissemination methods
- using CTIS to share observations of malicious activity, IoC, or malicious actor behaviour (tactics, techniques, and procedures)
- entities may use this information exchange in a range of activities depending on context (this might include automated blocking of URL hosts, file hashes, search activities, or even simply a confirmation of whether similar CTI have, or have not, been observed).

Reputation and dynamic categorisation

Support for reputation and categorisation, such as Reputation Block Lists (RBLs) and domain name categorisation, is often natively supported in even the most rudimentary web proxy and mail relay infrastructure. These capabilities can be used in gateway systems to make and enforce security policy decisions that allow or deny access to and from endpoints external to a given security domain.

Gateway infrastructure should use features such as near-time dynamic domain name categorisation, automated CTI ingestion, behavioural analytics, geo-location, and endpoint posture assessments as part of their gateway services.

Reputation and dynamic categorisation services can be implemented in many gateway components, including firewalls, forward and reverse proxies, mail relays, recursive DNS resolvers, and remote access services.

Where available, reputation and dynamic categorisation services should be used as part of a gateway's decision-making logic (security policy enforcement capability) when deciding to grant or deny access to a requested resource.

Organisations should take care to ensure that their own infrastructure does not get misused and subsequently listed on an RBL. Gateway operators should be aware of the better practices required to prevent a negative categorisation, how to monitor the reputation of gateway resources, how to manage risk and how to recover from such an event.

Organisations should have processes in place to allow for the reporting of security issues, including a public Vulnerability Disclosure Program (including /.well-known/security.txt), as well as the traditional reporting mechanisms for various platforms, for example, websites (contact us) and email (postmaster). For further information refer to the IETF's RFC 9116: A File Format to Aid in Security Vulnerability Disclosure.

Secure administration

Privileged access allows administrators to perform their duties such as establishing and making changes to servers, networking devices, user workstations and user accounts. Privileged access or credentials are often seen as the 'keys to the kingdom' as they allow the bearers to have access and control over many different assets within a network. As the control plane becomes more distributed, the importance of strong authentication increases.

ASD has developed secure administration advice to provide guidance and architectural patterns related to secure administration, covering the ISM controls:

- privileged access control
- multi-factor authentication (MFA)
- use of privileged access workstation logging and auditing
- network segmentation
- jump boxes.

Organisations should not expose management interfaces to the internet unless they have been designed and assessed for this purpose. In most cases, management interfaces should not be exposed outside of an organisation's security domain and are preferably only accessible from management zones and follow better practice to separate a system's control planes from a system's data planes. There is value in MFA-based pre-authentication of users via a gateway solution (for example, by using a reverse proxy or identity-based firewall rules) prior to presenting a management interface through a gateway. Further guidance on separating privileged operating environments and administrative infrastructure can be found in the ISM's Guidelines for System Management and ASD's Secure Administration.

Organisations should configure the strongest MFA option that is available within externally hosted administration portals (such as websites used to manage cloud, and registrar and registry consoles) in accordance with the Essential Eight.

Organisations should explore authentication systems that offer stronger-than-default protection of administrator accounts (e.g. Just In Time privileged access, short-lived Single Sign-On tokens). Where Active Directory is used within a gateway administration zone, it is strongly recommended that there are no domain trusts outside of that environment, and that access within that environment is limited to those with a need-to-access justification and that the justification and access is regularly reviewed to see if it is still required.

When implementing MFA solutions, it is recommended that organisations achieve a Maturity Level Three (for MFA) and that phishing-resistant MFA solutions are used, such as Fast Identity Online 2 (FIDO2).

Special consideration should be given to how an organisation implements read-only audit access to gateway management systems. This access is frequently needed for monitoring activities (such as policy and configuration audits) and may be needed to ensure transparency of system configuration. Care should be taken to ensure that this access cannot be used as a data exfiltration path out of gateway management zones. Organisations should develop processes to validate whether read-only access is implemented correctly, and sufficiently restrictive. Consider pushing audit evidence out of the management zone, rather than allowing users to directly download data.

Organisations are often required to place a high degree of trust in gateway system administrators, as they have privileged access to systems that process a range of sensitive data. The principle of role separation and least privilege should be followed (e.g. a person with internal administrative access should not have administrative privileges in a gateway). Organisations should ensure that staff with privileged access to systems have undergone appropriate background checks commensurate with the classification of data that the privileged access would facilitate. This is documented in the *Protective Security Policy Framework (PSPF).*

Organisations should also appropriately manage risk that is related to privileged non-person entity (NPE) accounts (service accounts). User and NPE account access should be limited to only the systems needed, and those accounts should not have access to other environments outside of the management zone used to administer the gateway.

Anomaly detection should be performed across authentication events, and any actions taken to administer a gateway. These capabilities should be prioritised as part of a continuous improvement activity.

Privileged User Training for NCEs is offered via ASD's Cyber Security Partnership Program. ASD's Privileged User Training is a tailored 2-day course that uses theory and practical exercises to provide privileged users with an in-depth look at how they can apply ASD's cyber security advice in their day-to-day work. Gateway staff, as privileged users, should undertake this or equivalent training before being granted privileged access.

Gateway administrator onboarding

Consider developing an onboarding process so that new gateway administrators get all the user access and briefings required to perform their role. Onboarding processes can help to ensure that staff clearances are verified, and that any development requirements are understood before

employment (ensuring staff have, or obtain, appropriate skills and clearances before being granted administrative access to gateway systems).

Onboarding processes include:

- clearance validations, and security briefings (where required)
- briefings on SOPs and work instructions
- the location of information repositories
- an overview of the roles and responsibilities of the team and individuals
- information about key contacts, their roles, and escalation paths
- information about the gateway system, services, and customers
- an outline of the change management and risk management processes
- information about health and performance monitoring platforms (cloud control plane dashboards and observability platforms)
- processes to register with key vendors (open support tickets and download security patches)
- document development plans and training requirements.

Administrative processes

Organisations that want to develop and operate gateway capabilities should be aware of the ICT Systems Development lifecycle. This process places emphasis on initial planning, analysis and design that precede the building, operating and maintenance activities. These processes should be undertaken by teams with architecture and engineering skills, supported by the organisation's governance frameworks and processes.

Organisations that want to consume gateway services still need their ICT processes, change management, and records management procedures to be formally documented and reviewed during the IRAP process.

While the processes will vary between traditional and cloud-delivered capabilities, organisations need to have the ability to assess supply chain risk, including platform-related risks and risks introduced through the gateway that may be protocol or service-specific.

Health and performance monitoring processes should be used to identify actual or potential service outages and should help organisations assess the performance of a service against a service level agreement. These monitoring activities can also help organisations undertake capacity-planning activities. Monitoring processes may require a deep understanding of how network-service dependence impacts service performance and availability.

Change management

- An organisation's change management processes should be supported by a configuration management database (CMDB). An up-to-date CMDB also assists security teams in identifying and engaging with system managers responsible for the system maintenance of a platform, service, or function (as both a preventative function, and as a cyber security incident response function).
- Changes to how a gateway is designed, implemented, managed or operated should be scrutinised for the introduction of risk (such as a loss of security visibility or control, unplanned system outages, and introduction of process flaws).
- Significant system changes, or changes to the implementation of security controls, should trigger an IRAP re-assessment of the gateway. Significant changes in risk should similarly trigger an authorisation to operate (ATO).
- The change management process (used to request changes to configuration to a service provider) should involve a critical analysis of the risks of implementing the proposed change, identifying and accepting risk, validation that the change was successful, and supporting rollback where necessary. Proposed changes should be peer reviewed by relevant subject matter experts.
- Version control and configuration management, particularly as they relate to accountability in implementing changes within a gateway environment, and the ability to validate the roll back of a system to the previous state in the event of a failed change, are all critical elements of change management.

Records management

- It is critical that the design, operation, maintenance, and changes to a gateway are well documented, and that this documentation exists for the life of the system.
- Organisations that manage gateways should capture institutional knowledge through recordkeeping and change management processes and decisions. Administrators implementing changes to a gateway system should capture the 'why' as well as the 'how' and the 'who' of approved changes. The classification of a gateway may be classified higher than an organisation's standard record keeping and service management (ticketing) systems (e.g. PROTECTED gateway with OFFICIAL internal networks).
- For Australian Government entities, the documentation, configuration, and the change management approval process of the gateway environment are subject to the recordkeeping requirements of the *Archives Act 1983*.
- Organisations should work with their Records Managers to ensure that the systems used to support a gateway meet the requirements of the Archives Act.

An organisation's obligation to protect data does not change because of a change in technology stack or service delivery model. Organisations that want to use new technology and service delivery mechanisms to make wholesale changes to how their gateway services are provided, may need to undertake significant assessment of their governance and operational strategies, policies, standards, procedures and skills. While there are significant benefits associated with adopting new technology, organisations need to ensure that their business processes and staff capabilities also evolve.

Organisations undertaking significant changes to a technology stack or business processes also need to monitor emerging technologies that can provide benefits, but can also introduce risk. Various organisations produce better practice guidance that is relevant to specific vendors, architecture, and service delivery models. Organisations should undertake research into activities that improve business outcomes (better service delivery models) or help mitigate risks associated with a technology or vendor.

Asset management lifecycle

When purchasing services, products, and equipment – whether physical or virtual, there are many considerations other than costs. The gateway, in both traditional and cloud-native forms, is a collection of systems and control planes. Procurement and service acquisition processes should aim to acquire repeatable, interoperable, and policy-aligned components that form part of the gateway system architecture. Broadly, the following categories of activities should be considered:

- assurance of supply chain (including service providers and third-party APIs)
- security considerations on product selection (including CSP-native services and virtual appliances)
- support considerations (including SLAs, patch cadence, and integration with observability tools)
- return merchandise authorisation (RMA) processes
- asset management (for physical, virtual and ephemeral assets).

Organisations should prioritise vendors and providers that have demonstrated capabilities in delivering secure, resilient control functions. This can include cloud-native vendors offering: reference architectures, IaC templates, validated SaaS connectors, and zero-trust architecture patterns. Organisations should collaborate with such vendors to identify better practices for deployment, operations and lifecycle management. A trust-but-verify approach is simplified where automated configuration against a blueprint or hardening guide is made available by the vendor or a trusted industry partner.

In addition to the STIGs, organisations should consider using relevant CIS Benchmarks, especially those tailored to gateway cloud platforms. These benchmarks support a consistent approach to security configurations across traditional and virtualised stacks.

Physical and personnel management

Whether in traditional on-premises environments or within cloud service models, gateway infrastructure must be protected from both physical and human threats. This includes:

- controlled access hardware in traditional deployments and secure identity governance in cloud environments
- environmental protections
- role-based controls for all personnel
- secure onboarding/offboarding for administrators
- continuous vetting and logging for all privileged access sessions.

Supply chain assurance

Procuring organisations need to identify and assess risks associated with the supply chain of both hardware and software of vendors gateway products and their supporting services. The scope of the supply chain includes the design, software development, physical manufacturing, delivery, deployment, validation, support, maintenance, remote or local management, and decommissioning.

Additional security considerations should be validated for personal and physical security in the vendor supply chain, including logistics, deployment environment (data centre) and operational considerations. Organisations can apply the recommendations in the Personal and Physical Security sections to their supply chain vendors.

When considering a gateway product or service, it is important for organisations to understand their own internal supply chain risk management processes, such as when procuring or outsourcing functions. Organisations should prioritise their assessments of 3rd party risk based on the criticality of the security policy enforcing functions, as it relates to the gateway's threat model. Organisations should ensure that the critical security functions they require to meet their organisational needs will be appropriately risk managed by the vendor including associated supply chain risks for those functions.

Gateway components need to be procured from trusted vendors with their own secure supply chain management processes. This means all aspects of supply chain risk need to be understood requiring clear and transparent reporting from vendors. Clear information from vendors will support decision-makers and operational staff to have confidence in selected products and that no malicious or unauthorised actions in the supply chain have caused compromise. Organisations should consider contractual terms with service providers to ensure that they maintain their supply chains, implementation, management and operational risks.

Mature vendors will have established mitigation strategies for secure transport of physical and digital products. Physical mitigations will include items such as the use of tamper evident seals on boxes and hardware, that must be checked before the receipt of products. Organisations should develop procedures for performing physical delivery validation checks before the receipt of goods

(e.g. ensuring that tamper evident seals are intact on delivery, validating serial numbers and confirming manifests). Digital products should be transported via secure digital channels. Mitigations will include items that can be used to verify the authenticity and integrity of a product such as:

- digital signatures
- using verifiable chains of trust, and
- secure hashes provided through secure secondary channel mechanisms (such as email or safe hand).

The Mitre ATT&CK framework identifies several supply chain compromises that can take place at any stage of the supply chain.

For more detailed guides on supply chain assurance refer to ASD's:

- Choosing secure and verifiable technologies
- Cyber Supply Chain Risk Management
- Identifying Cyber Supply Chain Risks
- How to Manage Your Security When Engaging a Managed Service Provider
- Questions to Ask Managed Service Providers.

Commonwealth entities procuring gateway services must consider the Department of Home Affairs' Hosting Certification Framework (HCF) and ensure all sensitive and classified government data and associated infrastructure rated at the classification level of PROTECTED is hosted by a HCF-certified provider.

Security considerations on product selection

Organisations should procure products from trusted and reputable vendors. Vendors should produce evidence of the capabilities that they have developed that improve their product security through better security defaults, support, and regular maintenance, and ensure these are continuously improved over time in response to changing threats. Vendors should provide full transparency about a product's purpose, deployment requirements, and integration requirements (including constraints and limitations) to procuring organisations before purchase. Procuring organisations should request independent third-party assessments of a product to assist in the risk decision-making process.

Service providers

When considering the use of a service provider, the organisation should consider the extent to which it can contractually enforce what the provider offers. This includes factors such as, the ability to:

- negotiate key performance indicators (KPIs)

- negotiate service level agreements (SLAs)
- integrate with other systems that support an organisation's strategies and requirements
- implement security policies.

When organisations are not able to validate security controls to appropriately manage risk with a service provider, then it should consider changing providers. Organisations should carefully evaluate all services offered to avoid vendor lock-in.

Software

Organisations should preference vendors who have committed to developing their gateway software to be secure-by-design and secure by default.

Secure-by-design is a proactive, security-focused approach taken by software manufacturers during the development of digital products and services that requires the purposeful alignment of cyber security goals across all levels of the manufacturing organisation. Secure-by-design requires that manufacturers consider cyber threats from the outset to enable mitigations through thoughtful design, development, architecture and security measures.

Secure-by-default refers to products that are secure 'out of the box' with little to no additional security setup or configuration required upon deployment. It means security measures designed to protect consumers against the most likely and prevalent threats are built into a product or service 'by default' at no additional cost. Most products will allow them to be configured by the procuring organisation. Vendors should ensure that products are delivered with the most secure settings configured by default, and provide clear advice about the potential risks and mitigations an organisation will need to consider if they deviate from secure defaults.

To counter risks associated with device software tampering, organisations should select vendor products that have implemented secure boot. This validation can be achieved via means such as hardware BIOS or firmware boot signing, with the Operating System (OS) signed by a trusted anchor.

Forensic capabilities

Forensic capabilities are essential for organisations to be able to conduct efficient incident response (IR) activities. Vendors should provide the following capabilities with their products to support the procuring organisation's IR activities:

- secure log generation and forwarding
- real-time telemetry
- forensic capture of volatile and non-volatile storage
- granular audit trails

Software bill of materials

A software bill of materials (SBOM) is intended to support organisations with the rapid identification of vulnerabilities within a product or service, or those that are inherited through included libraries. Vendors have started producing a SBOM in order to add transparency to the vendor's supply chain. Gateway services provided to an organisation, including self-administered and those offered by an MSP or a CSP, should include an SBOM. Organisations should consider the potential value of including SBOM requirements in contract clauses for gateway services.

After a vulnerability becomes known, an SBOM can be used by an organisation to identify products and services that may contain vulnerabilities inherited through the software supply chain. For example, the Log4j vulnerability (CVE-2021-44228), affected several platforms commonly used within gateway environments. An SBOM could allow organisations to identify their exposure to new public vulnerabilities ahead of vendor notification.

Product evaluation

Organisations need to be able trust or have a high level of confidence in a product that performs security functions. To achieve this ASD recommends products that enforce security policies in gateways are independently evaluated. Gateways are IRAP assessed using the ISM controls framework, and can in turn be built with policy enforcing components assessed through the Australian Information Security Evaluation Program (AISEP) using the Common Criteria (CC).

If using CC, evaluations should be conducted using the Protection Profile (PP) with tests conducted by Common Criteria Recognition Arrangement (CCRA) members. The CC PPs most pertinent to gateways are 'Boundary Protection Devices and Systems' and 'Network and Network-Related Devices and Systems'. These CC PPs were designed to comprehensively and systematically test vendor claims that a security product works as designed.

Other international standards and certifications vary in the level of assurance they provide, and none completely align to the controls in the ISM. For this reason, the best way to assess a gateway service is to have it assessed by an IRAP assessor against the controls in the ISM.

There are a multitude of other international standards and certifications that MSPs or CSPs can conform to and be certified against. While a prioritised list of certifications that may be relevant for the design and operation of a gateway is beyond the scope of this guidance, there are several industry frameworks and standards that can assist an organisation identify vendors that are assessing and managing risk. Assurance frameworks include:

- ISO 27001:2022
- ISO 31000:2018
- ISO 9001:2015
- Federal Information Processing Standard (FIPS) 140-3.

Organisations such as the Centre for Internet Security, NIST, NSA, and Cloud Security Alliance also play a part in codifying better practices (e.g. blueprint configurations and auditing processes).

Certain regulatory requirements may require organisations to implement specific controls or perform certain actions to be compliant with that regulation (e.g. PCI DSS, and CI/SONS).

Alternatives to third-party assessment

Market and business pressures to rapidly adopt cloud services has resulted in many new ICT solutions being designed that may not have undergone rigorous, systematic and independent testing. The pace of adoption of new technology, and the timeframes required to undergo rigorous security testing, may result in solutions being deployed that either do not use modern security features available in later versions of the product, or were not tested for assurance of security functionality. Organisations should review IRAP and other third-party audits that have been conducted to make risk-based procurement decisions. These reports can be used to identify unmanaged or inherent service delivery risk.

To counter risks surrounding the lack of formal functionality testing of security features, it is recommended that organisations choose suppliers and vendors that have committed to secure programming practices. If not included in the vendor's assessment reports, organisations are encouraged to ask vendors for information about the vendor's threat modelling and secure-by-design principles, history of security patching, vulnerability disclosure policy, SBOM, and transparency to customers about cyber security incidents.

In reviewing the quality of the product, including through the assessment processes, additional considerations should include:

- Does the vendor have history of detecting and patching vulnerabilities quickly?
- Does the vendor provide transparency to the customer?
    - Does the vendor provide detail on their secure code development processes, or share details of cyber security incidents with customers?
    - Does the vendor provider a SBOM or HBOM, or provide information about their supply chain risk management processes?
- Does the vendor meet SLAs?
- Does the vendor provide support that fits organisational needs?
- Does the vendor have spare parts readily available?
- Does the vendor provide enterprise features such as the ability to install a consumer's choice of external public key infrastructure (PKI)?
- Can customers refine the configuration of Transport Layer Security (TLS), and/or enable security features that help reduce a service's attack surface, or enforce security policy?

Organisations should review vendor documentation to understand how software and platforms auto-update (e.g. signatures, configuration, patches/updates). Vendors should be able advise how their software patching processes are protected from compromise.

Organisations should test and verify vendor claims that:

- have a poor rationale for why they need the level of network and/or internet access specified in their documentation
- require security or enterprise architecture principles to be bypassed to function, or requires controls to be disabled
- a product does not integrate with existing enterprise infrastructure.
- Concerns should be raised with the vendor if:
- an organisation's controls need to be bypassed for the product to function
- performing security functions removes vendor guarantees
- the level of privileged access is beyond what would typically be granted to a vendor, such as a user or service account.

Essential Eight in gateways

Elements of ASD's Essential Eight mitigation strategies are applicable to most gateway systems, containing principles that reflect better practice with a gateway system. Gateways can also provide Essential Eight controls to protect Microsoft Windows-based internet-connected networks.

ASD recommends that organisations should implement Essential Eight Maturity Level Three within their gateway environments. There are a number of other gateway and OS-related hardening activities documented in the Strategies to Mitigate Cyber security Incidents that are highly effective in preventing network attacks, that should also be implemented within gateways, management demilitarized zones (DMZs), jump hosts, and management workstations. The residual risks associated with operating at Maturity Level Two should be well understood by organisations. Organisations should consider incorporating contract terms to ensure these risks are managed when gateway services are provided and managed by a third party.

Organisations should apply the following Essential Eight concepts to gateway devices and services:

- apply security patches to all gateway systems (operating systems, firmware and applications)
- restrict who has administrative privileges to gateway systems
- enable MFA for devices and systems within a gateway
- backup data and configuration of all gateway devices and systems.

Support considerations on product selection

Data security

Support contracts should stipulate how the vendor will treat customer data provided to them as part of the support process. Where applicable, compliance with the Privacy Act should be

considered to ensure appropriate data protection measures are in place. Special caveats may include certain information not being stored in support systems and not retained by the vendor after an issue has been resolved. Organisations should consider requesting clauses requiring the vendor to confirm in writing that customer data has been deleted at the end of a support case. One example is when an IP router memory dump may contain cryptographic keys or other sensitive data, and there is no valid sanitation process. In this case, a necessary secondary action is to revoke the existing key material and reset administrative passwords when required.

Organisations should think about the releasability requirements for any government information and how it would be applied in each support case. If data is to be provided, how is this information to be securely provided to the vendor? Will this information be removed from their support systems after the issue has been resolved?

Remote access

Organisations should consider direct and competent supervision of access granted to external parties, who should only gain access through an organisation's approved remote access solutions. Organisations should consider what method and level of access may be needed for a vendor to provide support and ensure that this is codified in contract agreements. If support staff require physical or remote access to production systems, consider the need for local and cleared staff, and the effort required to escort staff. Where it is necessary to provide remote access to third-party support staff (e.g. shadowing via shared desktop software), consider if there is a need to provide read-write access to the system, or if it is simply necessary for them to view the configuration and guide staff through a troubleshooting process.

Organisations should consider requirements in the PSPF when determining if it is necessary to provide a third party with remote access to a gateway system (as opposed to read-only visibility through a remote access solution). Organisations providing third-party vendors with user accounts should follow standard organisational processes, including observing personnel clearance requirements.

The internet provides a valuable support channel for vendors. Several vendors have been implementing licensing and basic online telemetry collection to themselves or to a nominated third party via this method. If such practices exist, then this collection should be identified during the initial procurement process before a purchase, rather than 'discovered' after the fact.

Transparency and access to system information and telemetry

To provide support, a vendor may need to access logs and other data. In addition to ensuring an appropriately secure transfer mechanism, organisations also need to consider the security implications of providing the information requested by vendor support teams. For example, build scripts, IaC configuration, copies of running configuration, such as firewall configuration or an operating system or memory dump, may contain information that should be sanitised to reduce operational risk or compliance burden. Release of any information to the vendor needs to be authorised, before being provided.

External dependencies

A product may also be dependent on externally hosted services as part of the hardware (e.g. administration, telemetry, and health monitoring). Organisations should consider if the product will force the adoption of support models that they would not otherwise consider.

Maintenance contracts with suppliers and third-party vendors need to include functional SLAs for hardware and software support. These requirements should support an organisation's disaster recovery (DR) and business continuity (BC) requirements, for example, aligning SLAs (24x7 vs business day support) with the business recovery time objectives. While a gateway may normally provide a high-availability architecture, organisations should ensure that delays in the RMA processes do not delay the restoration of the gateway's high-availability state in the event of a hardware failure. When assessing supply chain support for the replacement of hardware, organisations should consider the performance of the vendor (e.g. Were SLAs of the support contract met? Have RMA processes occurred smoothly and swiftly? Are sanitisation and other deprovisioning processes hindered by vendor or supplier constraints?).

Test support arrangements

Organisations should conduct tests to ensure that vendor support contracts have been set up properly before they are needed (for example, during a critical cyber security incident or system outage). This process could be as simple as calling the vendor (outside of business hours if 24/7 support is part of the contract), confirming if a serial number is showing as under support, and that the member of staff is authorised to raise a support ticket.

Organisations should document the process for raising support calls with vendors. Identifying standard contact methods and escalation paths (e.g. vendor account managers) will ensure that the process of raising support requests is efficient. The worst possible time to test these cyber security IR procedures for the first time is during an outage. Store these details offline with other DR and BC documentation. Verify that any hardware replacements provided to you under RMA are also covered under your support contract. Be aware that procurement teams may inadvertently be nominated as the primary contact for support and maintenance purposes. Discuss organisational requirements with both internal procurement teams and the vendor account manager to ensure your organisation's operational requirements are clearly understood. It may be prudent to nominate multiple staff members who have the ability to raise and escalate support issues with a vendor.

DR and BC processes need to document what information is needed to rebuild and restore each component of a gateway to its current running configuration. A range of materials will be needed, including:

- build and configuration documentation, commonly referred to as 'As Built / As Configured' (AB/AC) documentation
- copies of running configurations
- software licences
- cable registers

- relevant firewall rules
- access to backup data such as password vaults and knowledge bases, like wikis and change and cyber security incident management systems used by gateway administrators.

Specific care needs to be taken with key material (associated with TLS, IPsec, SSH keys, and certificate escrow), otherwise encrypted data and/or services may never be recoverable. In the case of BC (that is, starting from scratch), an organisation should consider the order of system re-build; for example, re-build the management zone, gateway core, and then prioritise rebuilding DMZ capabilities.

Control validation and continuous assurance

Organisations deploying systems that enforce security policy should develop tests to ensure that security policy enforcement remains effective over time. These tests should include unit tests, integration tests, performance tests and end-to-end function tests of a service. Organisations should also test gateway controls to ensure that security policy remains effective after changes are made to the system, preferably through automated processes used to validate the change management process.

Controls should be tested after changes are made to security policy or software updates. It is a wise precaution to periodically retest controls as a continuous assurance activity. If using a CSP or MSP, ask the vendor what sort of continuous security assurance activities should be conducted by each party under a shared responsibility model. Organisations should preference vendors that are prepared to work with clients to develop tests for their gateway services.

Consider creating conceptual architecture diagrams (instead of the actual diagrams) for when there is a need to open a support case. This is also useful for other purposes, such as when working with developers or onboarding and familiarising new or lower-cleared staff.

Commissioning hardware

A better practice is to factory reset and re-install vendor software on devices prior to commissioning them. Enable secure boot on all gateway infrastructure that supports this feature. Preference should be given to vendors that support modern secure and trusted boot features backed by a trusted platform module. Verify software that has been manually downloaded before deploying it. File hashes and vendor software signatures should be validated. Vendors should supply digital signatures or hashes for binary files used to update systems, providing a form of integrity assurance. Vendors should also provide security mechanisms to ensure software downloaded and installed by the OS and applications is verified as originating from the vendor. File checksums and digital signatures should be validated through a secondary path rather than using a checksum or signature on the same website as the file was downloaded from.

Consider the need to perform a clean installation of a platform's operating system prior to commissioning new hardware. It is also very useful to have like-for-like hardware in a representative test environment.

Decommissioning and return merchandise authorisation

It is a good practice to power down infrastructure a couple of days before its removal from racks. Ideally, an organisation should perform in-place device sanitisation (validate that a backup of the running configuration exists). Statements or letters of volatility may also be informative in identifying appropriate memory sanitisation procedures. If in-place sanitisation is impractical, organisations should ensure that secure storage and transport processes are in place to protect deprovisioned infrastructure that is awaiting sanitisation. Note that there are operational risks associated with moving configured devices between facilities (e.g. theft from cars). Update asset registers as part of the decommissioning process, noting that the deprovisioning process will vary for different classes of devices and different types of service.

It is good practice to develop and test sanitisation processes before the deployment of new hardware platforms. As hardware failures are not uncommon early in a product's deployment, knowing how to perform these processes is useful in ensuring an expedited RMA process. It is also useful for decommissioning products at their end of life.

Hardware vendors often have RMA processes to replace faulty equipment. Requirements for RMA processes should be formalised in contracts. An organisation's operational staff should ensure that they have documented their RMA processes and have tested and verified a vendor's documented sanitisation processes. Multiple risk-based decisions will need to be made if a device cannot be sanitised successfully before returning it under an RMA process. Organisations should ensure that their standard operating procedures (SOPs) treat the deployment of any asset replacements as rigorously as their initial deployment (such as wiping, updating, re-configuring AB/AC documentation), as it is not uncommon for equipment returned to the customer under an RMA process to have system configuration from another customer.

Consider sanitisation processes when conducting proof-of-concept trials. Be clear with vendors that any returned equipment needs to be wiped and sanitised in accordance with the consuming organisation's risk management policies, and that this may require the return of non-functioning test equipment (that is, the vendor may need to re-build an appliance).

Operational activities

A range of activities should be conducted on a regular basis to ensure systems remain secure. These activities should be documented in SOPs. A plan of actions and milestones (POAM) should identify who is responsible for conducting tasks, the task schedule or deadline, and the management reporting lines for these activities (and any related findings).

Asset management

Organisations should ensure robust record keeping processes using a CMDB or equivalent ledger within a Service Management System (SMS). This should capture key metadata for both physical and virtual/cloud-native assets, including:

- serial numbers
- location of assets
- operating system, software, container image, or API versions
- important vendor contacts (support levels, account managers, and escalation paths)
- system identifiers, purpose, network locations
- software and hardware bill of materials
- classification and business impact levels (BIL)

CMDBs support lifecycle asset management and are essential in traditional, hybrid, and cloud-native environments including those using Security Service Edge (SSE) and Content Delivery Network services. They provide Business Impact Levels (BIL) and equipment classification. SSE services often span distributed infrastructure, APIs, and ephemeral workloads, all of which require governance as part of asset inventory. Such information helps identify backup criticality and requirements for DR and BC.

As CMDBs often contain a rich collection of sensitive system and service mappings information that is useful for malicious actors, such as malicious insiders, appropriate role-based access controls should be applied. Data holdings within a CMDB should be regularly validated against other available data such as vulnerability scanners and cloud-native inventory tools. These checks can highlight rogue or unmanaged systems.

Virtualisation offers useful abstraction layers, and provides more opportunities to automate workflows. Software, APIs, containers and cloud services are assets that organisations need to manage throughout their operational life. When using automation, consider the need to automate the creation, maintenance and deletion of appropriate CMDB records. CMDBs can also be enriched through other sources of information, such as through asset, service discovery, and vulnerability scans.

Systems that support gateway operations should be protected. CMDBs, version control systems (configuration repositories), CI/CD pipelines that support IaC, software libraries and ‘gold images’

are important systems and data that should be provided the same protections as the core gateway infrastructure. Risks related to availability and system integrity should be carefully considered.

Consider the use of shared and group mailboxes for vendor and product registration and support. Important support notifications are less likely to be missed using a shared mailbox. Vendors frequently notify customers of product announcements (e.g. new versions, end of sale/support/life, security announcements) that operational teams need to be aware of. By linking at least one support account with a shared mailbox, support staff can raise tickets in the event of a cyber security incident.

There is a need to understand the vendor's usage conditions as they may have stipulated legal terms and conditions on the product's usage or support. An organisation's legal representative should review the warranty and End User Licence Agreement (EULA) statements, and other service terms and conditions before procuring gateway hardware, software and services.

An often-overlooked asset is certificate or web services key material. These are used to allow servers and clients to authenticate and negotiate secure communications, such as TLS. These have a lifecycle and need to be managed, otherwise the organisation risks unplanned service outages due to expired validity periods.

Asset stocktakes

Asset mustering should include regular stocktakes to track a range of assets, including physical infrastructure, operating systems and software (including software versions), user and machine accounts (both local and centrally managed), cryptographic key material, and system backups. Stocktakes may involve physical verification for on-premises assets or logical verification through automated audit tools for cloud-native and virtual assets. Best practices include:

- labelling devices and removable media with assets tags and protective markings
- securing cryptographic key material and verifying system backups
- tracking software versions
- regular audits.

For more information on asset management refer to ASD's Foundations for modern defensible architecture: foundation 4 – reliable asset inventory.

System deployment checklist

Organisations should maintain and routinely update system deployment checklists to support consistent implementation of ICT governance processes across both traditional and cloud-native environments. These checklists should be adaptable to different system types, including physical hardware or cloud-based services. Deployment activities may include:

- applying asset stickers and protective markings for physical systems
- registering logical assets in the CMDB, tagging resources, and linking identity providers for cloud-based systems
- applying initial system or service upgrades and validating baseline versions
- updating the CMDB or service inventory with relevant metadata (e.g., environment, service tier, owner)
- configuring out-of-band management or cloud-native alternatives
- documenting the configuration in version-controlled repositories (AB/AC)
- applying hardened build and configuration standards
- validating IaC policy templates (e.g. SSE policy, Cloud Access Service Broker (CASB) rules, Zero Trust Network Access (ZTNA) access groups, Secure Web Gateway (SWG) Categories)

As there are many steps in putting a unit into production, a system deployment and configuration checklist verification performed by a second person can help with quality assurance.

Spot checks

Spot-check activities should ensure that appropriate physical on-site security processes are being followed, and that physical barriers are in place (e.g. alarm systems in place, doors, cages and racks are locked), and that an organisation's assurance processes are applied to gateway operational processes. Note that these types of checks may not be possible in a CSP, but where possible, performance of these checks should be in contract terms related to data centre facilities that host an organisation's gateways. Verification activities are typically undertaken by an origination's IT Security Advisor, but ad hoc validation processes may be undertaken by operational staff. The processes to undertake all of these spot-check activities should be documented in SOPs, supported by the gateway systems' security policy.

Checks may also be undertaken to ensure that systems previously deployed are brought into alignment with the current system deployment standards. This helps ensure that systems are consistently deployed throughout the gateway environment. Where possible, automation should be used to conduct technical verification.

ISM guidance specifically advises that organisations perform network cable audits. These checks can be supplemented with tools designed to assist with the physical cabling of a gateway environment (using automation to streamline auditing processes).

An IRAP assessment should validate that an organisation's Information Security Management Systems processes are followed. Organisations should retain records of spot checks and any observations or findings. Examples of these processes and the evidence they can produce include the following:

- Review the system POAM for evidence that identifies tasks have been accomplished.

- Access review: evidence that there has been a review of who still has access to the environment and a continued need for it.
- Policy Enforcement Point (PEP) ruleset review: evidence that reviews of firewall, proxy and other PEP rules are occurring, to ensure gateways are enforcing expected policies.
- Log review: evidence that logs are regularly reviewed and anomalies investigated.
- Cyber security incident review: the cyber security incident register links to documentation of cyber security incidents and how they were managed.
- Vulnerability and patch management review: provides evidence of vulnerability management processes, routine patching, emergency patching, etc.
- Discovery review: shows how gateway teams discover assets, shadow IT, and possibly also reviews non-gateway teams.
- Change management review: validates that configuration changes were approved and implemented according to organisational change management policies.

Disaster recovery

Organisations frequently fail to undertake effective DR exercises. Simple failover testing (to simulate a device failure) should be conducted regularly. More complex testing (such as simulating a data centre outage in a high availability environment) can require extensive negotiation with internal and external stakeholders. These activities should occur in alignment with the organisation's broader DR planning activities.

Vulnerability scanning

It is a good practice to conduct vulnerability scanning activities at regular intervals, including before and after making system changes, to ensure the system's security is effective during maintenance. Organisations may want to reduce the administrative burden of these activities through automation.

Vulnerability scanning should be conducted after applying the latest signature set. Vulnerability scanning of certain infrastructures (e.g. firewalls, Subject Alternative Name [SAN] certificates, or backup infrastructure) may not be comprehensive if the system or service being scanned is not supported by signature or platform-related tests.

Vulnerability scanning can result in false positives. These should be verified, and then fine-tuning can be implemented to reduce further false positives.

As vulnerability scanning activities are often noisy, there is a risk that the organisation's SIEM will be flooded with security events. Again, a level of tuning may be desirable, depending on the costs associated with these additional logs – there may be value in shortening log retention requirements for low-value logs. The vulnerability scanning scope needs to be carefully considered.

Commercial vulnerability scanners often require authenticated scans (either agent or agentless) to correctly identify non-compliance with patching and configuration standards, as well as identifying vulnerabilities that exist within a platform that is not exposed to the vulnerability scanning system. It is useful to understand the security posture of systems within DMZs (such as management zones and enclaves) that are being scanned.

Vulnerability scanning also assists with asset discovery – discovering unknown devices and services within a gateway. Organisations should conduct vulnerability assessment and vulnerability scanning on a regular basis, as these inform the system's AO of risks over time. The tests should be daily for organisations operating at Essential Eight Maturity Level Three. Executive reporting dashboards and reports are useful to have before granting (or renewing) an ATO.

An organisation should also consider the value of conducting SCAP and STIG scans as a compliance validation activity (tracking baseline changes over time), particularly prior to deploying a system into a gateway environment. Vulnerability scanning is particularly useful where security or operational teams do not have complete visibility and control over an organisation's ICT environment(s). Where an organisation uses a service provider, they should request ongoing visibility of the results of vulnerability scans relating to the managed components that protect their infrastructure and services as part of a continuous monitoring program.

Health monitoring

Organisations need to monitor the health and performance of services, as well as the assets that help provide them, ensuring neither is overlooked. Consider using platform APIs, telemetry, or other health monitoring interfaces that the vendor provides for this purpose, such as streaming telemetry, Simple Network Management Protocol (SNMP), integration support with monitoring tools, and real-time user monitoring (RUM). Use encrypted health monitoring and APIs protocols, and evaluate whether a combination of in-band and out-of-band monitoring is necessary.

Health monitoring of gateway services (and gateway components) is needed for a number of reasons:

- identifying capacity constraints (capacity planning)
- monitoring reliability of service and underlying infrastructure (SLAs and KPIs)
- monitoring performance.

Gateway services need to be functioning well in order to meet organisational needs. Health monitoring systems should incorporate performance metrics to ensure that services meet these needs.

The importance of visibility

Gateways should provide near real-time operational visibility to several teams within an organisation. Operations and SOC teams need visibility to ensure that systems, including security systems such as gateways, are operating correctly. A gateway should generate security-relevant logs and telemetry that can be used to identify, triage and respond to cyber security incidents.

Cyber security IR teams require data to respond to cyber security incidents in a timely and efficient way. The ISM recommends that access to all logs relating to an organisation's data and services is documented in contractual arrangements. Useful security-related events to centrally log include:

- data packets and flows permitted through gateways
- data packets and flows attempting to leave gateways
- real-time alerts for attempted intrusions.

As part of a log collection and retention strategy, it helps to consider the usefulness of different types of logs. The ISM has guidelines on what logs should be collected.

Event logs are integral to event monitoring activities – they should be retained for the life of systems, or potentially longer, where it is practical for an agency to do so and appropriate to the agency's risk management framework for information systems. The recommended retention rates for DNS and web proxy traffic is a minimum of 18 months.

Organisations should ensure that gateway logs are obtained from CSPs and MSPs as part of that organisations log analysis strategy. The ISM provides guidance that logs should be centrally stored.

Organisations should appropriately classify and protect log repositories. Note that aggregated log data (and other forms of metadata) is unlikely to be classified higher than the data passing through a gateway.

Log collection strategy

A log collection strategy, whether traditional or cloud-native, relies on knowing what's important to capture, how to store it securely, and how to make use of it. Key considerations include:

- what to log
- what to capture (e.g. headers, payloads of get requests)
- how to collect logs
- how to adjust log and telemetry generation during a cyber security incident
- authentication details (where relevant, e.g. proxy, VPN)
- system integration with a SIEM or Security Orchestration, Automation, and Response (SOAR), incident response and alerting
- privileged administration related events
- time synchronisation (consistent timestamps)
- the structure and format of log data (preferably aligned with a commonly used log schema)
- record keeping requirements, including log retention and disposal procedures (legal and regulatory needs)

- the location and log ingestion APIs of log storage systems
- the location and log ingestion APIs of log analysis systems (e.g. SIEM or SOAR)
- log integrity (when forensic or legal considerations are relevant).

The more capable an organisation's SOC gets, the larger volumes of data it can use. Security teams working in SOC environments need to work closely with business and governance teams to ensure that any required privacy impact assessments are undertaken to ensure compliance with legislation. Inexperience and process-related immaturity can lead to more collection and data retention than is required, having a ripple-effect on the total cost of ownership.

Logs and data provide support for multiple use cases security, troubleshooting, and billing—and should be structured to allow attribution in multi-tenant gateways or shared environments. Ensure reliable, secure logs are delivered from CSPs/MSPs to customer systems.

Threat modelling helps determine what's important to log and should be reviewed regularly. System administrators should be able to tune logging levels as needed, while staying aware of potential storage and in-line network performance impacts. Alternatives like out-of-band or stream-based telemetry may help reduce overhead. More information on threat modelling can be found in the 'Threat modelling' section within this document (see page 6).

While typically not a substantial cost, organisations should be aware that service providers may charge transport, storage or analysis fees (particularly shipping from a tenancy). Organisations must consider these operational costs as part of the enterprise architecture function. Log storage is a cost of doing business, but organisations should not store logs longer than necessary, as defined by the NAA or required by operational security teams, or other legislative or regulatory requirements. Organisations should understand the log retention policies of their suppliers as CSPs and MSPs may have different retention policies, which may not align with an organisations record keeping requirements defined by the NAA.

As logs from gateway systems can contain sensitive information, it is important to ensure that access to logs is restricted to authorised personnel. Logs from various gateway systems should be forwarded for centralised storage, classified appropriately, with appropriate role-based access controls (RBAC) and audit logging (for governance oversight purposes). These factors should be considered during the initial design of log storage and analysis solutions. As logs are used for legal purposes (e.g. to establish a forensic timeline), organisations should ensure log integrity, and have processes in place to prevent and detect tampering. Logs that have the potential to contain sensitive information should be encrypted when stored at rest. It is recommended that gateway administrators do not have write access to gateway logging systems. Appropriate RBAC models to manage insider threat should be a design consideration of logging and analysis systems.

The absence of logs from systems should be a concern to organisations. Organisations should have processes in place to identify when systems have stopped generating logs and telemetry. An organisation's SIEM should be configured to analyse logs and other telemetry (both real-time and historical analysis) for matches against IoCs.

SOC teams and gateway administrators should be in regular contact to ensure gateway systems are optimally tuned. For example, it is not uncommon for the structure of log data to change

unexpectedly during a system upgrade, resulting in a need to restructure the format of log data being generated (gateway side) or requiring SIEM log parsing engines reconfigured. To determine if a gateway system is generating logs, a gateway administrator can generate log events. The absence of these events can be detected by a SIEM, which then generates an alert indicating potential loss of logging capability. By conducting war-gaming (purple-teaming) activities, these teams can increase the effectiveness of communications and improve IR.

Gateway administrators should have a high degree of visibility of the file systems and configuration of gateway servers. Gateway administrators should perform regular file integrity monitoring and configuration validation against known and approved baselines, typically by monitoring for changes in file hashes or text-based configuration files. Changes to a system's configuration baseline should be investigated and confirmed as valid as an ongoing process. More detailed information on log collection and retention can be found in ASD's Best practices for event logging and threat detection, and the ISM's Guidelines for system monitoring.

Traffic payload inspection

The Gateway Security Guidance Package: Gateway Technology Guides document of this package describes in detail the desired level of visibility and control for specific network protocols that transit into and out of an organisation's security domain. These capabilities are described at a high level below.

Organisations that perform security inspection of network traffic have an obligation to inform staff of their organisational security policies in order to be compliant with the *Telecommunications (Interception and Access) Act 1979*, section 7.2.aaa. An organisation's security policies should state that DPI will be conducted by default, with formal processes to assess the risk of exceptions to this policy. Gateway administrators should work with security governance teams to regularly review exceptions to DPI policy. Organisations should conduct threat modelling and risk assess any DPI exceptions before implementing them.

Gateway capabilities should help an organisation implement their security policies. This may include denying access to known bad content or connections from sources of malicious traffic. To function in this way, gateways at some point must decrypt and inspect traffic. This capability may come through a service that relays, proxies or forwards traffic, such as recursive DNS resolvers, mail relays, and forward and reverse proxies. This capability may be transparent to users, such as through intrusion prevention systems (IPS).

Visibility and policy enforcement

Solutions such as explicit proxies, mail relays, and the chaining of recursive resolvers are required to retain policy enforcement for TLS version 1.3, and version 1.2 with Perfect Forward Secrecy (PFS). TLS version 1.3 poses longer-term security challenges for a network-based intrusion detection system (NIDS) and network-based intrusion prevention systems (NIPS), and other 'bump in the wire' security solutions, such as transparent proxies, as it natively supports perfect forward security.

Systems that decrypt TLS traffic require appropriate PKI management and related key management processes. The risks and privacy impacts associated with this activity need to be

formally documented by the organisations undertaking these activities. Processes should be documented in key management plans (KMPs), as part of the gateway security policies and related SOPs. For more on PKI, refer to *Platform hardening* section of this guidance (See page 37.)

An organisation's enterprise IT, and operational technology (OT) needs may be significantly different, in turn requiring different security policies to be enforced through a gateway. Each system consuming a gateway service should be assessed as IT/OT, with appropriate internal architectures and risk management processes developed from this designation. Building management may consist of a variety of systems that may need to communicate with either the internal network, the internet, or both. Examples may include cardex, air, CCTV, lighting, etc. These systems may operate in separate security domains, and at times it may be necessary and appropriate to integrate (e.g. traffic flows from OT to IT, mediated by a gateway).

A gateway system's ability to detect malicious content can be enhanced by performing sandboxing and off-device analysis through the Internet Content Adaptation Protocol (ICAP) or through other capabilities. Open-source sandbox technologies can provide baseline capabilities (or prove value through a proof of concept), and commercial offerings have greater levels of capability.

A gateway system may use pattern-matching techniques on files, network flows, or observable behaviours to detect malicious activity. It may employ aa variety of detection systems to identify malicious content or actions, and to generate metadata that can be later used to scan for IoCs.

File hashes (e.g. SHA256), fuzzy hashes (e.g. ssdeep) and other forms of pattern-matching (e.g. regex or Yara) can be particularly useful in gateway systems. These systems either generate telemetry and logs for analysis by external systems, or use native capabilities to assess data against lists of known-malicious artefacts. Gateway subsystems that expose APIs can enable an organisation to achieve economies of scale, and maintain consistent capabilities across different gateway service delivery systems. For example, web proxies, mail relays and file transfer solutions could all be integrated with a common system that performs anti-malware, sandboxing and telemetry generation through an ICAP API – though this may cause architectural inflexibility.

As network protocols become increasingly encrypted, the ability to perform meaningful packet capture of data transiting a network has become increasingly limited. These capabilities can still be implemented in most gateways that can support a proxy function. However, the ability to decrypt packet capture is likely to be limited to specific gateway functions in the future. NCEs should have, or develop, the ability to store decryptable packet capture, or payload data, for seven days, noting that this capability may only be practical through a gateway system.

By default, gateways should block content that cannot be decrypted, or where content scanning cannot be performed with the desired level of assurance.

Organisations need to identify the various audiences and the value of the reporting that is provided to them. ICT teams need to identify their stakeholders and what information they need to perform their role. For example, senior staff do not need every alert, and a gateway engineer who is on call may require a different set of alerts at different times of the day.

Monitoring and reporting should consider the following:

- operational performance

- current availability status (red, orange, green)
- trends in the number and type of operational threats detected and blocked
- volumetric data
- data loss prevention statistics
- how CTI can be generated and shared.

Anomaly detection

Monitoring of data flow characteristics can assist an organisation to measure performance and identify security issues. An organisation may need to analyse data over time to determine what flow characteristics look abnormal, including through the use of AI and ML tools. An organisation should implement appropriate monitoring that can identify abnormal events requiring further investigation.

Anomaly detection should be used to identify unusual or rare events that may indicate that a security compromise has occurred. Anomalous behaviour can be identified through a number of systems, including gateways, other network systems, and client and server endpoints.

Gateway systems should provide data and telemetry to SOCs in order to detect anomalous behaviours in an organisation's systems and services. Events and data may be analysed through a combination of statistical modelling (e.g. through a SIEM's AI/ML modelling), and human analysis (e.g. SOC IR analyst).

Gateway protecting capabilities

Gateways systems exist to mitigate a broad range of risks. However, each gateway service has attack surfaces that can be exploited by malicious actors. Nevertheless, gateways should be used to enforce an organisation's security policy, in turn reducing risks related to the following types of attacks:

- malicious code
- denial-of-service (DoS) and Distributed denial-of-service (DDoS) attacks against exposed services (e.g. web servers, DNS, VPN, and email)
- phishing and spam
- command and control
- insider threats
- data breaches (including deliberate exfiltration)
- unauthorised access.

Gateway maintenance

Platform patching

The window between the public identification of the existence and nature of a vulnerability – and its exploitation – has been reduced to days or hours in some cases. It is imperative that an organisation's change management processes support its system administrators in implementing security patching. An organisation's emergency and out-of-session approval processes should be low-friction and not require substantial effort to follow. If a vulnerability was announced after hours on a Friday, what organisational processes would help facilitate the security patching change request being implemented over a weekend (assuming waiting for remediation during business hours was unacceptable)?

Platform upgrades

OS and platform upgrades typically introduce new or updated features and functionality, which may include:

- test environments
- unit testing
- planning
- change management
- new vendor better-practice guidance
- business functionality testing
- assessing if more extensive testing should occur, such as performance testing.

Consider the need for supplementary training (formal and informal) for administrators. Validate health and performance monitoring, particularly if re-deploying on the same hardware.

Organisations should have a continuous improvement process. This includes implementing new capabilities and features developed by vendors that assist an organisation improve the effective and efficient management of controls. This may be prioritised, based on the organisation's threat modelling.

SecDevOps

IT operations

Traditional ways of managing infrastructure, including gateways, have been evolving over time. Automation and orchestration tooling has evolved, and there are now many ways to manage

gateway infrastructure. From NetOps to GitOps and DevOps, IaC, CI/CD pipelines, there are many options that can be used to manage gateway infrastructure.

The underlying concept of being able to implement services consistently, reliably and with reduced business risk is putting pressure on teams that are traditionally used to working on a defined set of change requests in a given change window in which to implement approved changes.

The DevOps approach of making smaller but more regular changes is a popular alternative to more traditional change management processes. Teams that have never worked as part of a DevOps process have a steep learning curve to develop an understanding of the concepts, processes, toolchains, testing and rollback strategies. Gateway teams that have previously been the ones to manage all gateway functions may need training and guidance as they transition to new ways of working. These ways of working are likely to involve closer and more involved contact with development teams.

As with many new concepts, the advice is to start the journey slowly, with caution and a small blast radius, and incrementally build capability, experience and confidence over time.

Governance

To achieve good SecDevOps outcomes, an organisation needs good governance processes. Key activities to achieve good governance of gateways include:

- formally defined roles and responsibilities
- consistently applied risk management and security policies
- testing and validation processes
- version control
- cyber security incident management
- comprehensive visibility of environment configurations
- delegated approval processes, and
- rollback strategies.

Automation

Many gateway functions can be automated using APIs. This includes common activities to release new software in a gateway. For example, implementing and removing firewall rules, implementing DNS changes, configuring reverse proxies, web server configuration, and routing table updates can all be automated.

There is still a need to understand the risks associated with making these changes, and there are important governance processes that would need to be applied to the application of automation to the management of a gateway environment.

Administrative concepts and processes have been carried over into DevOps, such as the control plane, privileged access, configuration version control, RBAC, and the generation of security and service-related telemetry.

The familiar concept of ensuring that a running configuration should align with documented build standards is also replicated through concepts such IaC, CI/CD pipelines, and continuous validation. Organisations may choose toolsets that use either imperative or declarative approaches to automation (or perhaps a combination).

Training

One of the benefits of the cloud is that gateway teams can gain experience using the processes in a stand-alone tenancy. Most gateway infrastructure is likely to have a cloud-based instance that can be used to build lab environments, allowing staff to build capability in a low-risk environment. Lab environments should not be externally attributable back to the organisation and should use dummy data where available. Lab environments can be a safe place to develop security processes such as threat modelling, integration, and management tool sets.

CSPs, and software and hardware vendors, often publish information on how to use their API's and may offer training in a broad range of tools and processes that support automation.

Organisations should consider a range of vendor-neutral training in the concepts and implementation of SecDevOps and related tool chains.

Platform hardening

Platform hardening reduces the attack surface of devices and services by applying settings beyond default settings. This is typically done to implement better practice settings to meet regulatory or industry specific security requirements.

In addition to the guidance developed by platform and software vendors, as well as MSPs and CSPs, organisations can reference OS and platform hardening guidance from the following organisations (alphabetical order):

- ASD, Essential Eight Maturity Model
- ASD, Hardening Microsoft Windows 10 and Windows 11 Workstations
- ASD, Guidelines for system hardening
- CCCS, Top 10 IT security actions to protect Internet connected networks and information (ITSM.10.089)
- Center for Internet Security (CIS), CIS Benchmark List
- NIST, Checklist Repository
- NSA, Cyber security Advisories & Guidance

- NCSC, Device Security Guidance
- NSA, Network Infrastructure Security Guidance (focused on hardening Cisco platforms)
- OWASP, Projects
- US Department of Defence, Security Content Automation Protocol (SCAP).

Note, SCAP and STIG both identify vulnerabilities, and identify deviations against vendor better practice, using XCCDF.

Protocol encryption

Organisations should use encrypted protocols services to protect network communications related to the administration and monitoring of gateway systems. Organisations should also disable all clear text management services (e.g. Telnet, HTTP, FTP, SNMP 1/2c) to ensure that sensitive information cannot be easily obtained by a malicious actor in a position to capture network traffic.

Out-of-band management

Administrative access to physical gateway components should be managed using dedicated management networks, noting that network and server infrastructure typically has dedicated console or management ports for this purpose. By using dedicated networks, organisations can retain administrative access to systems in the event of network disruptions that impact on a gateway system's data plane.

Secure Boot and Trusted Boot

Organisations should enable file system encryption and Secure Boot or Trusted Boot if the vendor makes this available. This capability is also available in non-traditional operating systems (e.g. routers and firewalls) and should form part of a procurement assessment.

Deprecated protocols

The Internet Engineering Task Force's (IETF) Best Current Practice 195 has formally deprecated TLS 1.0, 1.1 (including Datagram Transport Layer Security variants). Web proxies, reverse proxies, Secure Sockets Layer Virtual Private Network (SSL VPN), and mail relays should be hardened to disallow connections that use these TLS versions.

Organisations should conduct regular stocktakes of TLS implementations within ICT infrastructure, focusing on identifying deprecated implementations of TLS, and inadequately hardened implementations of TLS. Vulnerability scanning tools and open-source tools (e.g. SSLyze) can be used to scan networks for interfaces within gateways.

Regional Internet Registry

The Asia–Pacific Network Information Centre (APNIC) is the Regional Internet Registry (RIR) responsible for providing internet number resources (IPv4 and IPv6 address space, and Autonomous System Numbers – ASNs) within the Asia–Pacific region.

Organisations should consider if there are benefits to monitoring global route changes observed for the IP space they announce through Border Gateway Protocol (BGP).

Organisation account holders to RIR portals should configure the strongest MFA option that is available within administration consoles in accordance with the Essential Eight.

Contact details

Organisations who are assigned internet number resources from an RIR need to maintain appropriate and accurate contact details within their RIRs administration portals, in accordance with the RIR's policies. SOPs should be followed to ensure these details (including Whois contact details, and the authorised list of contacts for corporate, technical, and billing roles) are maintained during times of organisational change (e.g. staff departures or machinery-of-government changes). Consider the benefits of a group mailbox for contact details, instead of individual staff email addresses.

Border Gateway Protocol

Organisations should periodically audit their internet number resource assets.

Route Origin Authorisations and Resource Public Key Infrastructure

Resource Public Key Infrastructure (RPKI) enables resource holders of IP address space to explicitly authorise who can make BGP routing advertisements for that IP address space through a Route Origin Authorisation (ROA). Resource holders configure ROA with a list of the prefixes that an ASN is authorised to announce.

RPKI uses asymmetric cryptography to authenticate routing information on the internet. Organisations, particularly telecommunications carriers and large cloud providers, can use RPKI to verify routing information they receive, transmit and use in routing calculations. By monitoring, publishing, and enforcing RPKI information, an organisation may reduce BGP-related cyber threats, such as:

- DDoS attacks
- accidental or deliberate redirection or rerouting of their internet traffic
- undermining of IP address-based reputational services
- limiting routing instability on the internet.

RPKI ROA record(s) are published by network owners, and describe routes in terms of network/prefix and Border Gateway Protocol Autonomous Systems (BGP AS) from which they are expected to originate.

In isolation, creating ROAs does not prevent prefix hijacking, as carriers need to implement RPKI Route Origin Validation (ROV) in order to discard invalid updates and allow ROA to be fully effective. When network operators make changes to the IP address space length advertised through BGP route updates, updates to the relevant ROA are also needed to reflect this change.

Organisations should include related security clauses as part of their procurement and contract requirements, and should assess carrier claims about routing security, such as their implementations against Mutually Agreed Norms for Routing Security (MANRS), when making procurement decisions.

RPKI allows organisations to provide additional route authentication information to assist carriers in making appropriate forwarding/security decisions. It is recommended that network operators configure ROAs for all of the IP address space they are allocated, or manage on behalf of their clients, including where they are not advertised externally. The scope and impact of hijack attacks against the IP address space that an organisation owns is reduced by organisations configuring ROA, and carriers implementing ROV.

In cases where an organisation chooses to deprioritise invalid routes to meet specific operational requirements, as opposed to rejecting them, such decisions should be regularly reviewed – noting valid routes should always be prioritised over those that are not.

For more information refer to the ASD ISM's Guidelines for gateways, NIST's SP 1800-14: Protecting the Integrity of Internet Routing: Border Gateway Protocol (BGP) Route Origin Validation, IETF's RFC 7454: BGP Operations and Security and MANRS's About MANRS publications.

Router filtering

The border routers in gateways should implement ingress filtering to prevent a range of network attacks. Organisations should prefer network carriers, MSPs and CSPs that support IETF RFCs related to ingress filtering of invalid traffic.

Organisations should also ensure that their gateways apply egress filtering to prevent Bogon (invalid) traffic originating from their networks.

Service provider offerings are expected to align with IETF's BCP 38, and should align with BCP 84 and BCP 194.

For more information refer to the IETF's RFC 2827: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, RFC 3704: Ingress Filtering for Multihomed Networks, and RFC 7454: BGP Operations and Security publications.

Public Key Infrastructure

PKI is used to protect a range of gateway services, including web servers (e.g. TLS), email (e.g. TLS, S/MIME and DKIM), DNS (e.g. TLS and mTLS), remote management (e.g. SSH and RDP), network access control (e.g. 802.1x), and full disk encryption (e.g. BitLocker and Linux Unified Key Setup). The ISM contains advice and controls relating to PKI backed systems.

Certificates and Certificate Authorities (CA) are a way of making asymmetric key cryptography scalable by providing a secure and reliable means of verifying the public key that is used by other parties. However, the use of a CA, or the use of a public CA may not always be the most appropriate method, particularly if the CA is operated by a third party, or alters the nature of the trust arrangement between participants by introducing a third-party CA who must be trusted. In this scenario, organisations should use additional verification techniques to ensure the CA can be trusted, that is, obtaining the CA identifier (the thumbprint) through a secure alternate channel.

Vendors should attempt to minimise the number of public CAs that issue certificates on their behalf or are trusted within their products. Vendors should publish the details of the CAs they consume, so the procuring organisation can make informed decisions about a vendor's choice of CAs.

Organisations should be aware that in many cases, the process of renewing a certificate does not change the underlying key material to which the certificate refers. Organisations should consider generating new key pairs when generating certificate signing requests (CSR).

Limitations of PKI in gateway products

In many cases, the use of certificates may not improve the scalability or reliability of a service. In the case of a point-to-point encryption service under common administrative control, the use of manually exchanged Rivest-Shamir-Adleman (RSA) keys may be a more appropriate option because:

- there is no need to place any trust in a third-party CA
- reliability may be improved by virtue of there being no fixed expiry date on the material
- security and reliability may be improved because there is no need to try to incorporate processes or communications paths that would provide access to Certificate Revocation List (CRL)'s, Online Certificate Status Protocol (OCSP) service, etc.

In other cases, where a service such as a VPN, is intended solely for use by an organisation's staff, and where there are secure ways of distributing certificates from an internal, private CA, the VPN may be a more appropriate solution.

Organisations should exercise extreme care in circumstances where they intend to issue certificates, including when:

- the underlying key pairs were generated or are controlled by a third party, or

- the security posture or management arrangement of the appliances where the key pairs are to be used is unknown or subject to frequent change.

PKI management

A system's KMP should document the parties involved in the management of key material, with risk management plans assessing first party and third-party risks. The ISM recommends organisations should develop a KMP that codifies the operational practices of generating and managing certificates. A KMP should contain elements of the proceeding advice and other unique characteristics should be addressed in that document, and related SOPs. Procedural information should describe how key material is generated, stored, used, audited and revoked. Organisations should be able to describe who has access to the material, how the material is protected from unauthorised access and use, and what to do in the event the organisation's disaster recovery plan is enacted. Organisations should ensure they audit all certificates they have been issued from both public and private CAs.

The KMP should capture information on all variables used in generating certificates (e.g. algorithms, key strength, Common Name (CN), organisational unit (OU), usage restrictions, validity period). Certificates need to be recoverable or replaceable under DR and BC plans. If using a PKI hierarchy, the trust relationship should be documented and the processes of distributing keychains for servers and clients should be captured.

Certificates should be treated and managed as an accountable asset. Using certificates for client authentication will provide additional complexities to manage the lifecycle of certificates (expiry, refresh, and re-distribution should be well understood).

Organisations should understand the operational impacts of not replacing certificates before they expire, and should consider what business processes should exist to monitor for impending expiry. For example, health monitoring alerts may complement manual processes (such as a team tasking tool) to prevent unintended service outages caused by expired certificates.

The degree of control and documentation relating to keys should be commensurate with the importance and security value of the services with which the keys are used. Keys used in labs and testing environments may need little or no control or recordkeeping. However, there should be controls and processes that prevent the re-use of such keys in more sensitive or production instances. Where a certificate from a CA is no longer required, the certificate must be revoked. Using certificates with a short expiration date (validity) is another way to reduce risk.

Wildcard certificates, while convenient, increase the impact if the associated private key is compromised – as a malicious actor can now impersonate any entity that fits the wild card format. Wildcard certificates should only be used if no other viable process is available. Use of subject alternate names (SAN) extensions is preferred over wild card certificates.

Where certificates are used for authentication, the supporting services (including CRLs or OCSP) should be available in order to inform the authentication decision. Organisations should consider what actions should be taken in circumstances where these validation services are not available. For example, organisations should determine if a VPN connection attempt should be denied if the VPN server or connecting client cannot perform a CRL or OCSP validation check.

Gateway products and services should prevent the export of key material after they are installed or generated on a system. Products should support the secure storage of key material. For example, they could use a Hardware Security Module (HSM), Trusted Platform Module (TPM) or software security module (SSM). Using these provides assurance by offloading cryptographic functions and storage to a dedicated device or software. The usage does not reduce the need for an organisation to manage their PKI infrastructure appropriately.

If dedicated devices or software are not available, key material must be encrypted using a unique password. Appropriate access and audit controls should be applied to both the encrypted key files and the password. Note, there may be availability risks associated with the manual entry of a password to make the key material available for use. The use of products that will only accept or generate keys with no associated password should be avoided. If required key files and material should only be kept or stored in an unencrypted format for the minimal amount of time needed where no other approach is possible. Such tasks should be subject to appropriate oversight to prevent deliberate or inadvertent misuse of the key file. For more information refer to the DTA's Gatekeeper Public Key Infrastructure Framework, DISR's VANguard, ASD's Guidelines for Cryptography, and NIST's SP 800-57 Part 2 Rev. 1: Recommendation for Key Management: General Best Practices for Key Management publications.

Further information

Further information on topics covered in this section can be found in the following cyber security guidelines (listed in alphabetical order):

- Amazon Web Services, Security best practices in IAM
- APNIC, Threat hunting with Yara: The red pill approach
- ASD, Cyber Supply Chain Risk Management
- ASD, Essential Eight Maturity Model
- ASD, Guidelines for Cryptography
- ASD, Guidelines for system hardening
- ASD, Guidelines for System Management
- ASD, Hardening Microsoft Windows 10 and Windows 11 Workstations
- ASD, How to Manage Your Security When Engaging a Managed Service Provider
- ASD, Identifying Cyber Supply Chain Risks
- ASD, Questions to Ask Managed Service Providers
- ASD, Secure Administration

- ASD, Vulnerability disclosure programs explained
- ASD's Choosing secure and verifiable technologies
- ASD's Foundations for modern defensible architecture
- ASD's Information Security Manual
- CCCS, Top 10 IT security actions to protect Internet connected networks and information (ITSM.10.089)
- Center for Internet Security (CIS), CIS Benchmark List
- Common Criteria, Certified Products
- DISR, VANguard
- DTA, Gatekeeper Public Key Infrastructure Framework
- Github, SSDeep Project
- ICANN, Reputation Block Lists: Protecting Users Everywhere
- IETF, RFC 2827: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing
- IETF, RFC 3704: Ingress Filtering for Multihomed Networks
- IETF, RFC 7454: BGP Operations and Security
- IETF, RFC 8996: Deprecating TLS 1.0 and TLS 1.1
- IETF, RFC 9116: A File Format to Aid in Security Vulnerability Disclosure
- ISACA, Continuous Security Validation
- manrs.org, About MANRS
- Microsoft, Administration
- Microsoft, Enterprise access model
- MITRE, Supply Chain Compromise
- NAA, AFDA Express Version 2 – Technology & Information Management
- NAA, Cloud computing and information management
- NAA, Information management policies
- NCSC, Device Security Guidance
- NIST 800 154, Guide to Data-Centric System Threat Modeling
- NIST, Checklist Repository

- NIST, OSCAL: the Open Security Controls Assessment Language
- NIST, Security Content Automation Protocol
- NIST, Security Content Automation Protocol Validated Products and Modules
- NIST, SP 1800-14: Protecting the Integrity of Internet Routing: Border Gateway Protocol (BGP) Route Origin Validation
- NIST, SP 800-137: Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations
- NIST, SP 800-160 Vol. 1 Rev. 1: Engineering Trustworthy Secure Systems
- NIST, SP 800-57 Part 2 Rev. 1: Recommendation for Key Management: General Best Practices for Key Management
- NSA, Cyber security Advisories & Guidance
- NSA, Network Infrastructure Security Guidance (focused on hardening Cisco platforms)
- OpenSCAP, SCAP Components
- OWASP, Projects
- OWASP, Threat Modelling Cheat Sheet
- Securitytxt.org, security.txt
- SSLyze, SSL/TLS Scanning Tool
- Threat Modelling Manifesto, Principles
- US Department of Defence, STIGs Document Library

Table of Abbreviations

ACSC	Australian Cyber Security Centre
AISEP	Australian Information Security Evaluation Program
APNIC	Asia-Pacific Network Information Centre
ASD	Australian Signals Directorate
ASN	Autonomous System Number
ATO	Authorisation to Operate
BC	Business Continuity
BGP	Border Gateway Protocol
BIL	Business Impact Levels
CA	Certificate Authorities
CASB	Cloud Access Service Broker
CC	Common Criteria
CCCS	Canadian Centre for Cyber Security
CCRA	Common Criteria Recognition Arrangement
CDS	Cross Domain Solution
CIS	Center for Internet Security
CMDB	Configuration Management Database
CN	Common Name
CRL	Certificate Revocation List
CSP	Cloud Service Providers
CSR	Certificate Signing Requests
CTI	Cyber Threat Intelligence
CTIS	Cyber Threat Intelligence Sharing
DDoS	Distributed Denial-of-Service
DLP	Data Loss Prevention
DMZ	De-Militarised Zone

DNS	Domain Name System
DPI	Deep Packet Inspection
DR	Disaster Recovery
DTA	Digital Transformation Agency
DTLS	Datagram Transport Layer Security
EULA	End User Licence Agreement
HBOM	Hardware Bill of Materials
HCF	Hosting Certification Framework
HSM	Hardware Security Module
IaC	Infrastructure-as-Code
ICAP	Internet Content Adaptation Protocol
ICT	Information and Communications Technology
IETF	Internet Engineering Task Force
IoC	Indicators of Compromise
IPS	Intrusion Prevention Systems
IR	Incident Response
IRAP	Infosec Registered Assessors Program
ISCM	Information Security Continuous Monitoring
ISM	Information Security Manual
IT	Information Technology
ITSA	Information Technology Security Advisor
KMP	Key Management Plan
MANRS	Mutually Agreed Norms for Routing Security
MFA	Multi-Factor Authentication
ML	Machine Learning
MSP	Managed Service Providers
NAA	National Archives of Australia
NIDS	Network-based Intrusion Detection System
NIPS	Network-based Intrusion Prevention Systems

NPE	Non-Person Entity
OCSP	Online Certificate Status Protocol
OS	Operating System
OSCAL	Open Security Controls Assessment Language
OT	Operational Technology
OU	Organisational Unit
PEP	Policy Enforcement Point
PFS	Perfect Forward Secrecy
PKI	Public Key Infrastructure
POAM	Plan of Actions and Milestones
PP	Protection Profile
PSPF	Protective Security Policy Framework
RBAC	Role-based access controls
RBL	Reputation Block Lists
RIR	Regional Internet Registry
RMA	Return Merchandise Authorisation
ROA	Route Origin Authorisation
ROV	Route Origin Validation
RPKI	Resource Public Key Infrastructure
RSA	Rivest-Shamir-Adleman
RUM	Real-time User Monitoring
SaaS	Software-as-a-Service
SAN	Subject Alternative Name
SBOM	Software Bill of Materials
SCAP	Security Content Automation Protocol
SIEM	Security Information and Event Management
SMS	Service Management System
SNMP	Simple Network Management Protocol
SOAR	Security Orchestration Automation and Response

SSH	Secure Shell
SOC	Security Operations Centre
SSE	Security Service Edge
SSM	Software Security Module
STIGs	Security Technical Implementation Guides
SWG	Secure Web Gateway
TLS	Transport Layer Security
TPM	Trusted Platform Module
VPN	Virtual Private Network
XCCDF	Extensible Configuration Checklist Description Format
ZTNA	Zero Trust Network Access

Contact us

Following substantial updates to the Gateway Guidance in July 2025, ASD's ACSC welcomes feedback to ensure it remains clear, relevant and useful. If you have any questions or feedback, you can write to us or call us on 1300 CYBER1 (1300 292 371).

The Gateway Guidance is being released in parallel with the Department of Home Affairs' *Australian Government Gateway Security Standard*. We encourage interested stakeholders to provide feedback on the Gateway Standard directly to the Department of Home Affairs.

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