Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Gateway security guidance package: Gateway operations and management

First published: July 2022

Executive summary

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

- the adoption of new standards, protocols, or functionality
- the implementation of exemptions to security policy
- the development of new malicious actor tradecraft
- staff attrition
- the discovery of vulnerabilities
- operator error
- system load.

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

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

Organisations need to understand risks related to their supply chain. A big part of this is asset management which is a focus of this guidance. This includes themes such as needs analysis, procurement, deployment, platform hardening, administration, monitoring, maintenance, and decommissioning.

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

Security standards

Organisations should implement the documented better practice suggested by software and hardware vendors. Organisations should have comprehensive documentation on how to configure settings to achieve security outcomes, and information about how to verify that controls are in place and are operating effectively. In the absence of direct government security configuration guidance, organisations should refer to vendor guidance and industry better practice guidance from reputable sources.

Many information security assurance frameworks and better practices may be used, including:

- Infosec Registered Assessors Program (IRAP) assessments
- AS/NZS ISO/IEC 27000 suite
- 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 blueprints
- industry better practice guides.

When procuring products or services that perform a security enforcing function, organisations should consider selecting products that are manufactured by reputable suppliers, or suppliers that regularly undertake rigorous and independent security testing of their products, such as through the Common Criteria 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 processes, for example, by leveraging SCAP and STIGS, or cloud security blueprints. Refer to the *Platform hardening* section of this guidance.

Further information

- Common Criteria, *Certified Products*

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. Following vendor better practice guidance, it is important for products to have also undergone testing through processes such as Common Criteria (if an evaluated product is being used, deploy it in accordance with the security target). 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.

Architects and engineers should develop threat models to identify what attacks 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. Some examples are:

- What does security telemetry look like when operating effectively, and during control failures?
- When tags (e.g. VLAN, tenancy) are used to separate customer data, what telemetry can help to identify when a compromise has occurred across security domains?

Continuous assurance and validation

A continuous monitoring plan can assist an organisation in proactively identifying, prioritising and responding to vulnerabilities and other events which impact on gateway efficiency. Continuous monitoring includes activities to monitor systems for vulnerabilities. This includes identifying missing security patches, validating variances from configuration baselines, identifying and applying vendor better practice guidance, implementing continuous improvement activities, and leveraging new security capabilities provided by suppliers and vendors.

As part of a continuous assurance program, organisations should confirm that controls are in place and continue to operate effectively. By developing comprehensive tests to validate the gateway controls beyond traditional point-in-time audits, organisations will have a higher level of assurance that risks are being effectively managed. Regular threat modelling activities for the organisations systems and processes can highlight new attack vectors that can be controlled through new or enhanced mitigations, which in turn should be tested using automated or manual controls.

From the US National Institute of Standards and Technology (NIST):

Information security continuous monitoring (ISCM) is defined as maintaining ongoing awareness of information security, vulnerabilities, and threats to support organizational risk management decisions.

Further information

- NIST, SP 800-137: *Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations*
- ISACA, *Continuous Security Validation*

Threat modelling and other techniques can be used to help an organisation develop a suite of continuous validation tests. 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.

The use of infrastructure management automation – including Compliance-as-Code, for example, Open Security Controls Assessment Language (OSCAL) and SCAP; Infrastructure-as-Code; and Continuous Integration / Continuous Development (CI/CD) pipelines – can increase assurance that a given gateway is operating as designed over the life of the system.

Further information

- NIST, *OSCAL: the Open Security Controls Assessment Language*
- NIST, *Security Content Automation Protocol Validated Products and Modules*
- US Department of Defence, *STIGs Document Library*

OSCAL provides an automation mechanism to assess systems against a range of security control catalogues, such as the *Information security manual*, configuration baselines, system security plans, assessment plans and results. The Australian Signals Directorate (ASD) produces the ISM in the OSCAL format.

Continuous validation frameworks and tools

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

- SCAP (leveraging Extensible Configuration Checklist Description Format – XCCDF, and 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, ISO 9001)
- penetration testing by suitably skilled and experienced personnel
- automated vulnerability scanning
- manual validation techniques (e.g. manual testing of firewall rulesets)
- periodic load testing, stress testing, and performance testing
- CI/CD practices.

Further information

- NIST, *Security Content Automation Protocol*
- OpenSCAP, *SCAP Components*

In traditional infrastructure environments, gateway staff will need to work with other operational and business teams to conduct some of these tests.

Organisations need to document policies and processes for version control of security configurations, as well as software version tracking, for example, documenting OS and firmware versions, as well as the running of security configuration. Tools that perform this function should be isolated in management environments, with access strictly controlled. Organisations should consider contractual terms with service providers to ensure that they maintain visibility into supply chains, implementation, management and operational risks.

An organisation might use security evaluated products (such as Common Criteria) 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 default 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.

Further information

- IETF, *RFC 9116: A File Format to Aid in Security Vulnerability Disclosure*

- Securitytxt.org, *security.txt*

Threat modelling

Organisations should perform regular threat modelling activities to identify unmanaged or insufficiently mitigated risks in gateway systems as part of their continuous assurance activities.

Threat modelling is the activity of identifying and understanding the various security threats that could affect the confidentiality, integrity and availability of a system or solution. Threat modelling also assists architects, engineers and developers to prioritise and mitigate risks through technical solutions and controls.

Threat modelling as a process can be performed at various points throughout the lifecycle of a system, including after system deployment as part of the continuous improvement and upgrade activities. However, developing a good threat model in the early stages of the design process can help prevent architectural decisions that would introduce threats that may not have been identified until much later in the lifecycle.

Threat modelling is often combined with the development of a risk management plan, and therefore some threat modelling methodologies are designed to be aligned with the organisation's standards or controls framework.

Further information

- Threat Modelling Manifesto, *Principles*
- OWASP, *Threat Modelling Cheat Sheet*

Threat modelling process

Threat modelling is typically performed in four stages, with each stage based on one of these four questions:

- What are we working on?
- What can go wrong?
- What are we going to do about it?
- Did we do a good enough job?

Threat modelling should be a mental checklist for all staff working in gateway environments, and be part of the operational culture of gateway teams.

At times, beyond the application of threat modelling on a business-as-usual (BAU) basis, gateway implementations may benefit from threat modelling workshops aimed at producing a formal threat model description. 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.

The first stage should focus on developing a high-level understanding of the gateway solution, 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 stage will involve the application of threats to the model in order to describe how the gateway solution can be impacted by malicious actors. There are many ways to approach the development of threats, including the use of threat taxonomy databases or the development of attack trees. There is also software that assists with the development of diagrams and the generation of relevant threats. The output of this stage should be a list of threats that are applicable to the gateway solution, including some description of what the impact of those threats would be.

The third stage is where the model will start to incorporate controls and architectural principles that minimise or remove the realisation of threats or mitigate the impact of those threats. This guidance, along with the wider advice and guidance provided by ASD such as the ISM, can be used as a resource for 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 stage involves the threat model being circulated amongst the solutions' stakeholders and other required subject matter experts. This gives other parties the opportunity to review the list of threats and mitigations, and propose considerations and ideas that may have been missed in the previous steps. This step can lead back into any of the previous three steps, depending on what feedback is received.

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.

Cyberthreat intelligence

CTI should be leveraged as a security mechanism in gateway environments. Gateways should leverage 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 of the service allows 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).

Cyberthreat Intelligence Sharing Platform

ASD provides a Cyberthreat Intelligence Sharing (CTIS) platform that supports automated cyberthreat intelligence sharing. This service allows partners to bi-directionally share CTI in a common language via a secure machine-to-machine sharing mechanism. 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 Cybersecurity Partnership Program.

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

- consume CTI to inform organisations of indicators of compromise (IoC) or malicious behaviour via processes such as STIX and TAXII, 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 leverage 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.

Further information

- ICANN, *Reputation Block Lists: Protecting Users Everywhere*

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.

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).

Further information

- IETF, *RFC 9116: A File Format to Aid in Security Vulnerability Disclosure*
- Securitytxt.org, *security.txt*

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*.

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. JIT privileged access, short-lived SSO tokens). Where Active Directory is used within a gateway administration zone, it is strongly recommended that there are no domain trusts outside of this environment, and that access within this 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.

An organisation is 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 accounts should be limited to access only the systems needed, and these 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 Commonwealth entities is offered via ASD's Cybersecurity Partnership Program. ASD's Privileged User Training is a tailored two-day course that uses theory and practical exercises to provide privileged users with an in-depth look at how they can apply strategies to mitigate cybersecurity incidents in their day-to-day

work. Gateway staff, as privileged users, should undertake this or equivalent training prior to being granted privileged access.

Further information

- ASD, *Secure administration*
- ASD, *Guidelines for system management*
- Microsoft, *Enterprise access model*
- Microsoft, *Administration*
- Amazon Web Services, *Security best practices in IAM*

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 in addition to 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 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 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 on 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 cybersecurity 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 security controls implementations should trigger an IRAP re-assessment of the gateway. Significant changes in risk should similarly trigger a new ATO.
- The change management process (used to request changes to configuration to a service provider) should involve the critical analyse of the risks of implementing a 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 it relates to accountability in implementing changes within a gateway environment, and the ability 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' 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 Commonwealth 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.

Further information

- NIST, *SP 800-160 Vol. 1 Rev. 1: Engineering Trustworthy Secure Systems*
- NAA, *Information management policies*
- NAA, *AFDA Express Version 2 – Technology & Information Management*
- NAA, *Cloud computing and information management*

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 leverage 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 business benefit, but can also introduce risk. A variety of 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, there are many considerations beyond costs. The gateway is a collection of systems, and procurement processes provide a means to acquire repeatable and interchangeable objects that form part of the system architecture. Broadly, the following categories of activities should be considered:

- assurance of supply chain

- security considerations on product selection
- support considerations on product selection
- return merchandise authorisation (RMA) processes
- asset management.

Organisations should prioritise vendors that have demonstrated experience in delivering control capabilities and writing secure code. Organisations should work with such vendors to identify better practices for deployment, operations and maintenance. 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.

Supply chain assurance

Organisations need to identify and assess risks associated with a vendor's manufacture, supply, maintenance, support, and management of gateway hardware, software and services. Organisations should prioritise the assessments based on the criticality of the security policy enforcing functions, as they relate to the gateway's threat model.

Products and services that are being used to protect and enable information access within gateway systems need to come from trusted suppliers. This means that all aspects of the supply chain need to be understood, so that decision-makers and operational staff have confidence that no malicious or unauthorised manipulation of equipment or code has occurred.

When considering a gateway product, it is important to understand the organisation's practices relating to their supply chain risk management processes, such as when procuring or outsourcing functions. The scope of the supply chain includes the design, manufacture, delivery, deployment, validation, support and decommissioning of hardware, software and related services that are used within a system.

When procuring ICT systems and services, supply chain risks exist in the hardware, software, and support systems used to provide ICT services, and should be considered along with other logistic, deployment environment (data centre) and operational considerations. Organisations should validate that devices and software are deployed according to the vendor's recommendations. Where this is not possible or practical, the risk should be assessed, understood and managed.

Mature vendors that have established mitigation strategies for supply chain risk will use tamper evident seals on boxes and hardware that should be checked prior to the receipt of ICT infrastructure. An organisation's ICT, procurement, and logistics teams should work together to develop procedures for performing physical delivery validation checks prior to the receipt of goods (e.g. ensuring that tamper evident seals are intact on delivery, validating serial numbers, confirming manifests).

Secure transport and storage of unconfigured equipment should be factored into an organisation's supply chain management processes. Cold spares in storage should be protected from tampering to the same degree as a unit in production.

ASD publications on supply chain risks are:

- ASD, *Cyber supply chain risk management*
- ASD, *Identifying cyber supply chain risks*
- ASD, *How to manage your security when engaging a managed service provider*

- ASD, *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 an HCF-certified provider.

The Mitre ATT&CK framework identifies a number of supply chain compromises that can take place at any stage of the supply chain including:

- manipulation of development tools
- manipulation of a development environment
- manipulation of source code repositories (public or private)
- manipulation of source code in open-source dependencies
- manipulation of software update and distribution mechanisms
- compromised of system images (multiple cases of removable media infected at the factory)
- replacement of legitimate software with modified versions
- sales of modified or counterfeit products to legitimate distributors
- shipment interdiction.

Further information

- MITRE, *Supply Chain Compromise*

Security considerations on product selection

Products used should be sourced from trusted and reputable suppliers. Vendors should provide capabilities that provide defence through the use of better security defaults, support regular maintenance of the product or service, and ensure this is continuously improved over time to respond to changing threats. A product's purpose, deployment requirements, and integration requirements (including constraints and limitations) should be known prior to purchase. For gateway services, independent third-party assessments of security functionality should exist and, when configured as specified, the product should operate as intended.

Equipment and products

Organisations need to have a trusted or higher level of confidence in products that perform security functions so that they can securely perform their role (i.e. enforce security policy). ASD recommends that products that enforce security policy in gateways are independently tested, for example, against a Common Criteria (CC) 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.

Services

When considering the use of a service provider, the organisation should consider the extent to which it can contractually influence what the provider offers. This includes factors such as the ability to negotiate key performance indicators (KPIs), service level agreements (SLAs), the ability to integrate with other systems that support an organisation's strategies and requirements, and the ability to implement security policies.

When an organisation is not able to implement and validate security controls to appropriately manage risk in a service, then it should consider a change in service provider. Organisations should also evaluate all services carefully to avoid vendor lock-in.

To counter risks associated with device 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 OS signed by a trusted anchor.

Software

Vendors have started producing a Software Bill of Materials (SBOM) in order to add transparency to the vendor's supply chain. An SBOM is intended to support organisations with the rapid identification of vulnerabilities within a product or service, or those that are inherited by included libraries. 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 a number of platforms commonly used within gateway environments. An SBOM could allow organisations to identify their exposure to new public vulnerabilities ahead of vendor notification.

Systems

International standards and certifications vary in the level of assurance they provide, and none exist that completely align to the controls in the ISM. For this reason, when assessing a gateway service for use by an organisation, there is no substitute for that service to be assessed by an IRAP assessor against the controls in the ISM.

There are a multitude of 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 a number of industry frameworks and standards that can assist an organisation identify vendors that are assessing and managing risk. Assurance frameworks include:

- CC PPs
- ISO 27001, ISO 31000 and ISO 9001
- IRAP
- 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 in order to be compliant with that regulation (e.g. PCI DSS, CI/SONS).

Alternatives to third-party testing

Market and business pressures to rapidly adopt cloud services (which are rapidly evolving) 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 leverage modern security features available in later versions of the product, or were not tested for assurance of security functionality. Organisations have to make risk-based decisions to achieve objectives and deadlines, but unmanaged or unidentified risk may be reflected in IRAP assessments or other third-party audits that are conducted.

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 and secure configuration principles and practices. If not included in the vendor's IRAP report, organisations are encouraged to ask vendors for information about their threat modelling, use of Secure by Design and Secure by Default principles and practices, history of security patching, vulnerability disclosure policy, SBOM, and transparency to customers about cybersecurity incidents.

In reviewing the quality of the product, including through the IRAP process, additional considerations should include:

- Does the vendor and/or product have history of detecting and patching vulnerabilities quickly?
- Does the vendor provide transparency to the customer?
- Does the vendor usually meet SLAs, and are support and spare parts reality available?
- Does the vendor provide details on their secure code development processes, or share details of cybersecurity incidents with customers?
- Does the vendor provide enterprise features such as the ability to install a consumer's choice of external public key infrastructure (PKI)?
- Can a consumer refine the configuration of Transport Layer Security (TLS), and/or enable security features that help reduce a services 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 to communicate their threat model for how 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 in order 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 in order for the product to function, when performing security functions removes vendor guarantees, or where 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 the Essential Eight 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 target at least Maturity Level Two (preferably Three) when implementing the Essential Eight within their gateway environments. There are a number of other gateway and OS-related hardening activities documented in the *Strategies to mitigate cybersecurity 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
- restrict who has administrative privileges to gateway systems
- enable MFA for devices and systems within a gateway
- back up data and configuration of all gateway devices and systems.

The below table maps the Essential Eight to gateway operations.

Essential Eight maturity

Essential Eight mitigation strategy	In scope for gateways	Gateway considerations
Application control	Always in scope (for services within a customer's control).	In scope for management servers and administrative workstations exposed to email and web browsing threats (e.g. Windows based jump hosts and management hosts).
Patch applications	Always in scope – prioritise achieving Maturity Level Three.	Externally facing services and systems should be prioritised when making patching decisions.
Configure Microsoft Office macro settings	May be in scope – prioritise achieving Maturity Level Two.	If Microsoft Office is not installed, including on Linux hosts, the macro controls and hardening guidance are not applicable. Gateway security policies should specify if Microsoft Office is to be installed in management zones (not recommended). Gateway management zones that create or use Microsoft office files should only generate and export office files from that management zone (administrative functions should not require office files to be imported into management zones used to administer gateways).

User application hardening	May be in scope	If no Microsoft Office applications are installed, then the macro controls and hardening guidance are not applicable. If PDF software is not installed, then the hardening guidance is not applicable. Any internet access from gateway management zones should be by exception, and strictly to a controlled set of endpoints to meet a business need which cannot be addressed in any reasonable way. This position should be supported by both security policy and controls. Web browser hardening advice should be implemented. Gateway security policies should specify appropriate web browser use within management zones, noting that internet access from these zones is not recommended. Web browser, Microsoft Office and PDF software security settings should be monitored for changes by users (including privileged users). Gateway security policy should communicate expectations to users/administrators within the gateway environment. PowerShell hardening advice should be implemented. .NET Framework 3.5 and below is removed or disabled.
Restrict administrative privileges	Always in scope – prioritise achieving Maturity Level Three.	Windows Defender Credential Guard features are not applicable for Linux based systems.
Patch operating systems	Always in scope – prioritise achieving Maturity Level Three.	Patch operating system and firmware for all platforms used in a gateway.
Multi-factor authentication	Always in scope – prioritise achieving Maturity Level Three.	Maturity Level Three is expected for system administrators managing gateway systems. Internet facing systems for customers should also use MFA, but this use case falls outside of this gateway advice.
Regular backups	Always in scope – prioritise achieving Maturity Level Three.	

Further information

- ASD, *Essential Eight maturity model*
- ASD, *Strategies to mitigate cybersecurity incidents*
- ASD, *Hardening Linux workstations and servers*

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. Special caveats may include certain information not being stored in support systems, or 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. A number of vendors have been implementing licensing and basic online telemetry collection to themselves or to a nominated third party via this method. If such practises exist, then this collection should be identified during the initial procurement process prior to a purchase, rather than 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 in order 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, health monitoring). Organisations should consider if the product will force the adoption of support models that it 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 test that vendor support contracts have been set up properly before they are needed (for example, during a critical cybersecurity incident). This process could be as simple as calling the vendor and asking 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 cybersecurity incident response 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 with the ability to raise and escalate support issues with a vendor.

DR and BC processes need to articulate 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; and access to backup data such password vaults and knowledge bases like wikis and change and cybersecurity 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 (i.e. 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 may be wise to periodically test 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 prior to deployment. 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 download.

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 (e.g. spare parts, testing processes).

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 to identify 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 prior to 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 to ensuring an expedited return merchandise authorisation (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 vendors documented sanitisation processes. Multiple risk-based decisions will need to be made if a device cannot be sanitised successfully prior to 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 returned equipment needs to be wiped and sanitised as per the consuming organisation's risk management policies, and that this may require the return of non-functioning test equipment (i.e. the vendor may need to re-build an appliance).

Asset management

Organisations must ensure good record keeping processes using a CMDB or equivalent ledger in a Service Management System (SMS), which captures the following:

- serial numbers
- location of assets
- operating system, software, or API versions
- important contacts (level of support, support and account managers, and escalation paths)
- system names, system purpose, system network location, bill of software, etc.

CMDBs should also provide Business Impact Levels (BIL) and equipment classification as required, providing additional support for the lifecycle management of the asset. Such information helps identify backup criticality and requirements for DR and BC.

CMDBs need to be appropriately secured as they contain a rich collection of information that is useful for malicious actors, such as malicious insiders, and should be regularly validated against other available data such as vulnerability scanners. Vulnerability scanner results may identify assets which are not tracked in a CMDB.

Virtualisation offers useful abstraction layers, and provides more opportunities to automate workflows; however, these functions and features have attack surfaces that need to be managed. Software, APIs, containers and cloud services are assets that organisations need to manage throughout their operational life. When leveraging 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 that should be provided the same protections as the broader gateway. 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 cybersecurity 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 License Agreement (EULA) statements, and other service terms and conditions prior to 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.

Gateway administrator onboarding

Consider developing an onboarding process so that new gateway administrators get all of 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 prior to employment (ensuring staff have or obtain appropriate skills and clearances prior to being granted administrative access to gateway systems).

While not exhaustive, 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 on the change management and risk management processes
- information about monitoring platforms (health, performance and capacity)
- processes to register with key vendors (open support tickets and download security patches)
- implement development plans and training requirements.

System deployment checklist

Organisations should document all of the processes needed to implement the expected ICT governance processes.

A generic systems deployment checklist can be used by gateway operators to ensure that appropriate activities are conducted. At a unit level, this may include activities such as:

- applying asset stickers and protective markings
- upgrading the system
- updating the CMDB with location and other standard information
- configuring out-of-band access
- documenting the configuration
- applying hardened build and configuration standards.

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

Operational activities

There are a range of activities that 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).

Stocktakes

Asset mustering should include 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.

An organisation's asset stocktake procedures may require the physical verification of infrastructure, or may allow logical verification through an appropriate audit mechanism. Organisations should label devices with asset tags, apply protective markings to the device, and apply appropriate protective markings to removable media.

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 ITSA, 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. A non-exhaustive list of examples of these processes and the evidence produced include the following:

- POAM review: 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 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.
- Cybersecurity incident review: cybersecurity incident register maps to documentation of cybersecurity incidents and how they were managed.
- Vulnerability and patch management review: evidence of vulnerability management processes, routine patching, emergency patching, etc.
- Discovery review: review how gateway teams discovers assets, shadow IT, and possibly also review non-gateway teams.
- Change management review: validate that configuration changes were approved and implemented as per 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 organisations broader DR planning activities.

Vulnerability scanning

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

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

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

As vulnerability assessment activities are noisy, there is a risk that the organisations 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 are not exposed to the vulnerability assessment scanning system. It is useful to understand the security posture of systems within DMZs (such as management zones and enclaves) 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 (daily for organisations operating at Essential Eight Maturity Level Three), informing the system's AO of risks over time. Executive reporting dashboards and reports are useful prior to granting (or renewing) an ATO.

An organisation should 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 assessment 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 and should not focus on one of these at the expense of the other. Organisations should investigate 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, real-time user monitoring (RUM), etc. Organisations should perform health monitoring and APIs through encrypted means, and should assess if a combination of in-band or out-of-band monitoring is necessary.

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

- identifying capacity constrains (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 a number of 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 logs and telemetry that can be used to identify, triage and respond to cybersecurity incidents. Cybersecurity incident response teams require data to respond to cybersecurity 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.

As part of a log collection and retention strategy, organisations need to understand the value of the logs that can be generated by systems. 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 should identify:

- 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 cybersecurity incident
- authentication details (where relevant, e.g. proxy, VPN)
- integration with a SIEM (or SOAR)
- privileged administration related events
- time synchronisation (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
- the location and log ingestion APIs of log storage systems
- the location and log ingestion APIs of log analysis systems (e.g. SIEM)
- log integrity (when forensic or legal considerations are relevant).

The more capable an organisation's SOC gets, the larger volumes of data it can leverage. 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 should help a team derive meaning from service use. In mature organisations, these logs are used by multiple stakeholders, for example, for troubleshooting, security, billing, and other business use cases.

Operators of centralised and multi-tenancy gateways should already be aware that their customers want the logs related to their service uses. Services should attribute logs to a given organisation, and mechanisms for transport from the provider to their customers should be functional, reliable, and secure.

Threat modelling can help an organisation identify what is good and what is bad (i.e. use threat modelling to identify what may be important to log). More information on threat modelling can be found in the *Threat modelling* section within this document.

Organisations should have the technical capability and maturity to know when to enable and disable additional logging. Operational teams may, from time to time, need to increase system logging (e.g. debug logging). When enabling an increased level of logging, organisations need to be aware of the potential performance and cost implications of doing so (e.g. storage requirements, SIEM integration).

High levels of logging can have a negative impact on device performance. Excessive logging can have a negative impact on log storage systems and impact SIEMs. Organisations should consult vendor documentation to determine if alternative logging options can reduce this impact, for example, log-ship telemetry, and alternative streaming telemetry (e.g. TCP instead of UDP, delivery out of band).

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 should 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 into 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 valid as an ongoing process.

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 exemptions to this policy. Gateway administrators should work with security governance teams to regularly review exemptions to DPI policy. Organisations should conduct threat modelling and risk assess any DPI exemptions prior to them being implemented.

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 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.

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 increased levels of capability.

A gateway system may use pattern matching techniques on files, network flows, or observable behaviours. A gateway may use a variety of systems to detect malicious content or action, and to generate metadata that can be subsequently used to scan for IoC. File hashes (e.g. SHA256), fuzzy hashes (e.g. SSdeep) and other forms of pattern matching, for example, regex or Yara, can be particularly useful in gateway systems as these systems either generate telemetry and logs for a separate processor, or use native capabilities to assess data against lists of known bad.

Gateway sub-systems that provide APIs may allow an organisation to gain economies of scale, and consistent capabilities across various gateway service delivery systems. For example, web proxies, mail relays and file transfer solutions could all be supported by a common system that performs anti-malware, sandboxing and telemetry generation through an ICAP API (noting this may come cause architectural inflexibility).

Further information

- APNIC, *Threat hunting with Yara: The red pill approach*
- Github, *SSDeep Project*

As network protocols become increasingly encrypted, the ability to perform meaningful packet capture of data transiting a network will be limited. These capabilities can still be implemented in the majority of 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
- DLP 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 of 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. While not intending to be exhaustive, gateways should be used to enforce an organisation's security policy, in turn reducing risks related to the following types of attacks:

- malicious code
- DoS and 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 public identification 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 work through. A method of testing workflows is to assess – if a vulnerability was announced after hours on a Friday, what organisational processes would help facilitate the change request being implemented over a weekend (assuming 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, Infrastructure-as-Code (IaC), CI/CD pipelines, there are a plethora of 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 a lot of 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 starting to displace 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 roll-back strategies. Gateway teams that have previously been the ones to manage all of the 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 to this is having formally defined roles and responsibilities, consistently applied risk management and security policies, testing and validation processes, version control, cybersecurity incident management, and excellent visibility of the environment's configuration, appropriately delegated approval processes, and rollback strategies.

Automation

Many gateway functions can be automated through the use of 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 are all automatable.

There is still a need to understand the risks associated with making these changes, and there will still be important governance processes associated with managing a gateway environment that are not removed through automation.

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 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 stand-alone tenancy. Most gateway infrastructure is likely to have a cloud-based instance that can be used to build out 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.

Cloud service providers, 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 in order 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*
- 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, *Cybersecurity 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 use 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 systems 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 systems data plane.

Secure Boot and Trusted Boot

Organisations should enable file system encryption and Secure Boot or Trusted Boot where 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 (March 2021) has formally deprecated TLS 1.0, 1.1 (including DTLS 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 stock-takes 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 can be used to scan networks for interfaces within gateways.

Further information

- IETF, *RFC 8996: Deprecating TLS 1.0 and TLS 1.1*
- SSLyze, *SSL/TLS Scanning Tool*

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 public key cryptography to authenticate routing information on the internet. Organisations, particularly telecommunications carriers and large cloud providers, can leverage 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 cyberthreats, 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 for 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 so as to reflect this change.

Organisations should include relate security clauses in contracts, 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.

Further information

- NIST, *SP 1800-14: Protecting the Integrity of Internet Routing: Border Gateway Protocol (BGP) Route Origin Validation*
- manrs.org, *About MANRS*

Route 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 relating 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 BCP 38, and should align with BCP 84 and BCP 194.

Further information

- 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*

Public Key Infrastructure

PKIs are 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.

The use of a CA, particularly where the CA is operated by a third party, alters the nature of the trust arrangement between participants by introducing a third party CA who must also be trusted.

The use of a public CA would normally be implied where there is a need for a many-to-many interaction, or there is a need for clients on a public network, such as the internet, to be able to interact with government services with a reasonable degree of confidence regarding authenticity and without needing to trust a CA operated by a Commonwealth entity.

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 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 and 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 consumption of an organisation's staff, and where there are secure ways of distributing certificates from an internal, private CA, this may be a more appropriate solution.

Irrespective of any other technology, organisations should attempt to minimise the number of public CA's that issue certificates on their behalf. Organisations should also consider publishing details of such arrangements, so that other parties have some expectation regarding what CAs might issue certificates for the organisation (e.g. CAA records in DNS).

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.

Organisations should exercise extreme care in circumstances where they intend to issue certificates and:

- 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.

A system's KMP should document the parties involved in the management of the key material. Risk management plans should assess first party and third party risks. Procedural information should describe how key material is audited, stored, used, and revoked.

Organisations should be able to describe where and how keys are generated and stored, who has access to the material, and how the material is protected from unauthorised access and use.

The degree of control and documentation relating to keys should be commensurate to 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.

Wildcard certificates, while convenient, do increase the consequences of the loss of the associated private key as a malicious actor can now impersonate any server in the organisation. The use of wildcard certificates should be minimised. Keys associated with wild card certificates should be subject to a commensurate degree of control and audit. Use of SAN extensions is preferred over wild card certificates.

Where certificates are used for authentication, the supporting services (including CRLs or OCSP responders) should be available in order to inform the authentication decision. Organisations should also 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 an OCSP validation check.

Private keys should not normally be exportable after installed on a system. Products should support the secure storage of key material, for example a Hardware Security Module (HSM) or Trusted Platform Module (TPM). Gateway products and services should prevent the export of key material. If this is not available, an organisation should use password-protected key material. 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.

Private key files should be encrypted using unique passwords. Appropriate access and audit controls should be applied to both the encrypted key files and the passwords. Certificates are an important element in the configuration of services, not only is access to this element important, but restoration of certificates from an event should be considered as part of an organisation's disaster recovery plans.

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.

Organisations should keep track of certificates that they have been issued from a public CA. Where a certificate from a public CA is no longer required, the certificate should be revoked. Certificates with a short expiration date (validity) is another option to reduce risks associated with the loss of key material.

The usage of HSM provides assurance by offloading cryptographic functions and storage to a dedicated device. The usage of a HSM does not reduce the need for an organisation to manage their PKI infrastructure appropriately.

Organisations should use unique underlying key material for any certificates that they have issued. As per ISM-0507, 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.

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 incentives 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.

Further information

- DTA, *Gatekeeper Public Key Infrastructure Framework*
- DISR, *VANguard*
- NIST, *SP 800-57 Part 2 Rev. 1: Recommendation for Key Management: General Best Practices for Key Management*
- ASD, *Guidelines for cryptography*

Contact details

If you have any questions regarding this guidance you can write to us or call us on 1300 CYBER1 (1300 292 371).

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.

For more information, or to report a cybersecurity incident, contact us:

cyber.gov.au | 1300 CYBER1 (1300 292 371)

Australian Government

Australian Signals Directorate