Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Gateway security principles

Gateway Security Guidance Package

First published: July 2022

Last updated: July 2025

Table of contents

Introduction **1**

Key terms and concepts **2**

Definition of a gateway 2

Cross Domain Solutions 3

Gateway architecture 3

Objectives of a gateway 4

Visibility 4

Detection 4

Prevention 5

Protection 5

Definition of a security domain 5

Policy enforcement point 7

Gateway security principles **9**

Security management is continuous 9

Risk is continuously managed 10

The invisible cannot be protected 10

Gateways protect organisations and staff 10

Plan for security flaws 10

Balance business and security 11

Risk cannot be outsourced 11

Placement of gateways **12**

Cloud-based gateways **13**

Cloud-native capabilities 13

Service-native integrations 13

Inter-organisation collaboration within a single cloud service provider 13

Security zones 15

Security Service Edge 16
Shared responsibility and trust 17
Risk considerations for outsourcing gateway services 18
Roles and responsibilities 18
Tenancy considerations 19
Cyber security incident reporting 20
Gateway visibility and telemetry 21
Cyber security incident response 24
Architect for maintenance 26
Cyber threat intelligence 26
Priority services for security visibility 26
More information 28
Contact us 28

Introduction

Gateways play a vital role in securing networks by managing and controlling data flows between different security domains. As key boundary systems, they enforce security policies and help protect an organisations' systems from external threats.

This guidance outlines the core security principles for effectively designing, implementing and managing gateways. It is intended for security, architecture and engineering teams who are responsible for designing or operating gateway solutions in their organisation.

The Australian Signals Directorate's Australian Cyber Security Centre (ASD's ACSC) developed these principles as part of a broader suite of guidance designed to support organisations in making informed risk-based decisions throughout the lifecycle of their gateway systems.

Although tailored for Australian Government entities and their service providers, this guidance is relevant to any organisation looking to strengthen the security posture of its gateway infrastructure.

Key terms and concepts

Definition of a gateway

The definition of a gateway in the *Information Security Manual* (ISM), 'Gateways securely manage data flows between connected networks from different security domains'.

Gateways are a set of capabilities that enable an organisation to securely:

- provide services to external parties
- exchange information with others
- support remote work
- operate across trusted and untrusted networks (including the internet).

A gateway is a network boundary solution responsible for controlling data flow into and out of an organisation's ICT environment or security domain. This position in the network means that gateways are critical implementation points. They provide a broad range of security capabilities that enforce an organisation's security policies before allowing access into or out of the organisation's network.

A gateway should:

- apply risk mitigations and cyber security controls to data flow between security domains
- provide visibility of transiting data according to an organisation's policies.

All stakeholders involved in designing, procuring, operating, maintaining and disposing an organisation's gateways need to understand the design principles and objectives of gateway controls.

Organisations can use gateway solutions through various delivery models, including on-premises, cloud-native, hybrid, or managed service provider (MSP) models. With cloud-native or MSP models, an organisation may use a gateway as an abstracted set of security services and capabilities rather than as specific equipment or functionality.

A gateway is typically comprised of a collection of physical, virtual and logical components that work together to provide the gateway's core networking and security services. Factors that can influence an organisation's gateway design include:

- threat posed by connecting to external networks
- business and operational requirements and strategies
- technical capabilities
- confidentiality, integrity, availability and privacy requirements

- threat modelling, risk appetite and risk management strategies
- integrations between self-hosted and cloud services
- sourcing and service delivery models
- location and number of data centres, staff and office locations
- availability of staff to design and sustain gateway capabilities
- operational visibility requirements.

Cross Domain Solutions

Organisations can use a Cross Domain Solution (CDS) as part of an internet gateway. A CDS is a system capable of implementing comprehensive data flow security policies with a high level of trust between two or more different security domains. CDSs are implemented between SECRET or TOP SECRET networks and any other networks belonging to different security domains. They can also be used for networks that operate at or below PROTECTED level that connect to the internet, where high-assurance security policy enforcement capability is needed to manage risk.

Refer to the ISM and resources for Cross Domain Solutions when gateway solutions contain at least one security domain classified SECRET or TOP SECRET, or classified at PROTECTED or below where a high-assurance solution is required to manage identified threats and resulting harm.

Gateway architecture

Gateways are often viewed as a single integrated solution that combines various services. There are several architectural approaches to consider:

- **Monolithic**: provides all gateway security functions through one centrally managed system (e.g. a secure internet gateway).
- **Disaggregated**: provides service-specific gateway functions through discrete but interoperable systems that do not share a common control plane but do share a common security policy.
- **Hybrid**: provides all required gateway services through a mixture of central and disaggregated service offerings and control planes.

Regardless of the architectural approach, it is critical that gateway services and their cyber security capabilities evolve to support an organisation's changing business and risk management needs. Previous Australian Government gateway policies and frameworks applicable to non-corporate Commonwealth entities (NCEs)[1] have historically advocated for the routing of internet traffic

[1] Non-corporate and corporate Commonwealth entities are government bodies that are subject to *the Public Governance, Performance and Accountability Act 2013*. More information: Non-corporate Commonwealth entity (NCE) | Department of Finance

through a small number of well-controlled secure internet gateways (SIGs), also described as monolithic gateways, which provided all gateway security functions through one centrally managed system. These SIGs quickly became the default place to concentrate cyber security capabilities as they controlled all traffic between a trusted internal network and the untrusted internet. However, this approach introduced constraints for system and network architecture, topology and security controls.

While a monolithic gateway concentrates security functions through a single centrally managed system, a disaggregated gateway separates gateway functions into discrete but interoperable systems. This allows for service-specific gateway functions to be delivered through separate systems. The availability of modern service delivery and consumption models, such as cloud, software-defined wide area network (SD-WAN), and remote work, have highlighted that the traditional monolithic gateway is no longer the only network security model available to NCEs. Therefore, a hybrid gateway provides gateway services by combining monolithic and disaggregated architecture.

The opportunities offered by advances in underlying gateway technology, services and capabilities, including those offered under a cloud-native consumption model, are evolving the broad range of architectures adopted by organisations to deploy their gateway solutions. In some cases, data is no longer routed through an organisation's existing gateway. The consequence here is that, where it is routed, it is not in a form that can be readily assessed by existing security tools. Hybrid and cloud-native gateways, combined with new ways of working, mean that gateway architectures will look different than before.

Objectives of a gateway

Without clearly understanding the capabilities or limitations of its gateways, an organisation cannot accurately or truly understand and manage its operational or cyber security risks. Senior and executive decision-makers, as the accountable persons or authorities in their organisation, are responsible for ensuring all relevant stakeholders clearly understand their organisation's gateway design principles and objectives.

Generally, an organisation has the following core cyber security objectives and functions for a gateway.

Visibility

- Improve Operational and cyber security visibility through generating and forwarding security-related telemetry.
- Understand and observe data flows.

Detection

- Monitor data flows for anomalies, suspicious activity and policy violations.
- Respond to detected cyber security threats, incidents or anomalies.

Prevention

- Enforce cyber security policies and prevent data breaches through:
  - implementing technical cyber security controls
  - authenticating, authorising and accounting (for specified services)
- Reduce an organisation's attack surface by only permitting approved data flows.

Protection

- Supporting the resilience critical business services (e.g. website hosting and browsing, email and remote access).
- Implementing compensating controls and mitigations for known vulnerabilities pending the application of security patches or the updates made available by vendors.
- Limiting or containing the impact of any compromise or incident.

Definition of a security domain

The ISM defines a security domain as:

> A system or collection of systems operating under a consistent security policy that defines the classification, releasability and special handling caveats for data processed within the domain.

A security domain may be a collection of ICT services that have a degree of commonality, which is used to justify the collection being treated as a single homogeneous group when it comes to security. Therefore, as ICT footprints evolve, organisations must assess several policy factors when defining their security domains. Policy factors may include:

- purpose
- ownership or sovereignty (Australian or foreign owned or controlled)
- consistent implementation of administrative and security controls
- consistent security and operational visibility
- data value and sensitivity (security or otherwise)
- threats and risks to which they are exposed (risk profile)
- interdependency on other systems (e.g. data and application processing interfaces)
- data classification, business impact level, information management markers and other caveats
- ability to perform cyber security incident response (e.g. manage a data spill)

- encryption.

Examples of different security domains include:

- an organisation's OFFICIAL and PROTECTED networks
- PROTECTED networks of two different entities
- multiple government tenancies within a shared service provider, such as those hosted by MSPs or cloud service providers (CSPs).

This means that systems that are operated by different organisations, or systems that operate at different classifications, are in different security domains.

Most organisations will usually have at least two security domains (trusted/internal and untrusted/external). Security domains can span across data centres and segmentation models (such as those offered by CSPs), depending on risk tolerance and technical design. Note, as outlined in the *Protective Security Policy Framework* (PSPF), the security policy for a security domain includes:

- security governance
- personal security
- physical security
- information security.

When using cloud or managed services, an organisation should assess whether these services form part of an existing security domain or should be in a separate security domain. This should be informed by business and compliance requirements, risk appetite, enforcement capabilities and security guides for consumers and providers.

Security domains do not need to align with specific network topologies. For example, virtual private networks (VPNs) can unify multiple locations under a single domain as long as the VPN is isolated from the transport network by using appropriate controls.

Gateways play a critical role in protecting security domains. Organisations should design controls to reduce or eliminate the attack surface associated with data flow between domains. This includes using both traditional on-premises and modern security technologies such as:

- proxies and web application firewalls (WAFs)
- host-based firewalls
- software-defined networking (SDN)
- behavioural analytics
- API-based logic enforcements
- Security Service Edge (SSE) solutions.

Consumers and service providers may have different perspectives on how to define a security domain. Figure 1 shows that multiple customers may share a virtual environment managed by a third party. As each organisation (consumer) is responsible for protecting its own systems and data, it should determine its own security domains. However, where a service provider offers multi-tenancy gateway services, the service provider may retain administrative and policy control over the system. Therefore, in cases where there is no clear demarcation of where the security domain boundary exists, an organisation should consider treating these environments as separate security domains until satisfied that there is equivalent security policy enforcement and operational visibility.

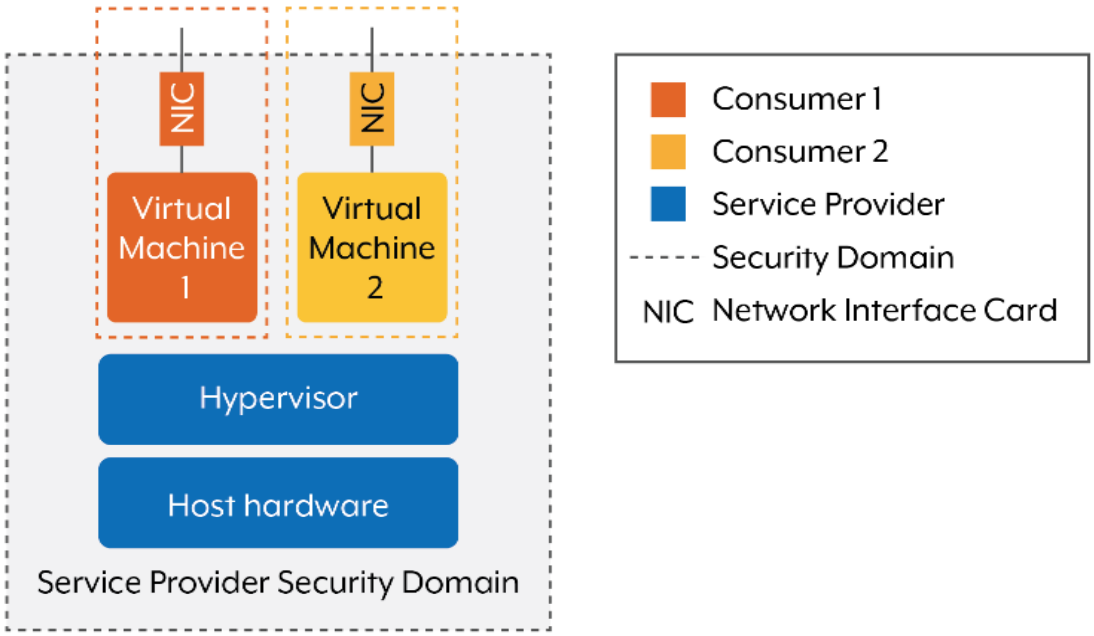

Figure 1: multiple security domain perspectives in a multi-tenant environment

Policy enforcement point

A policy enforcement point (PEP) may be a hardware or software component, security device, integrated appliance, tool, function or application that enforces an organisation's security policy.

Figure 2 shows that PEPs may operate synchronously or asynchronously and can be implemented at various layers of the Open Systems Interconnection (OSI) Model. Together, they support a defence-in-depth approach to securing data flow across security domains. A PEP can be considered a gateway capability responsible for enforcing security policy as data traverses between domains. Often, multiple PEPs are deployed in sequence (chained) to meet the complete set of gateway enforcement requirements.

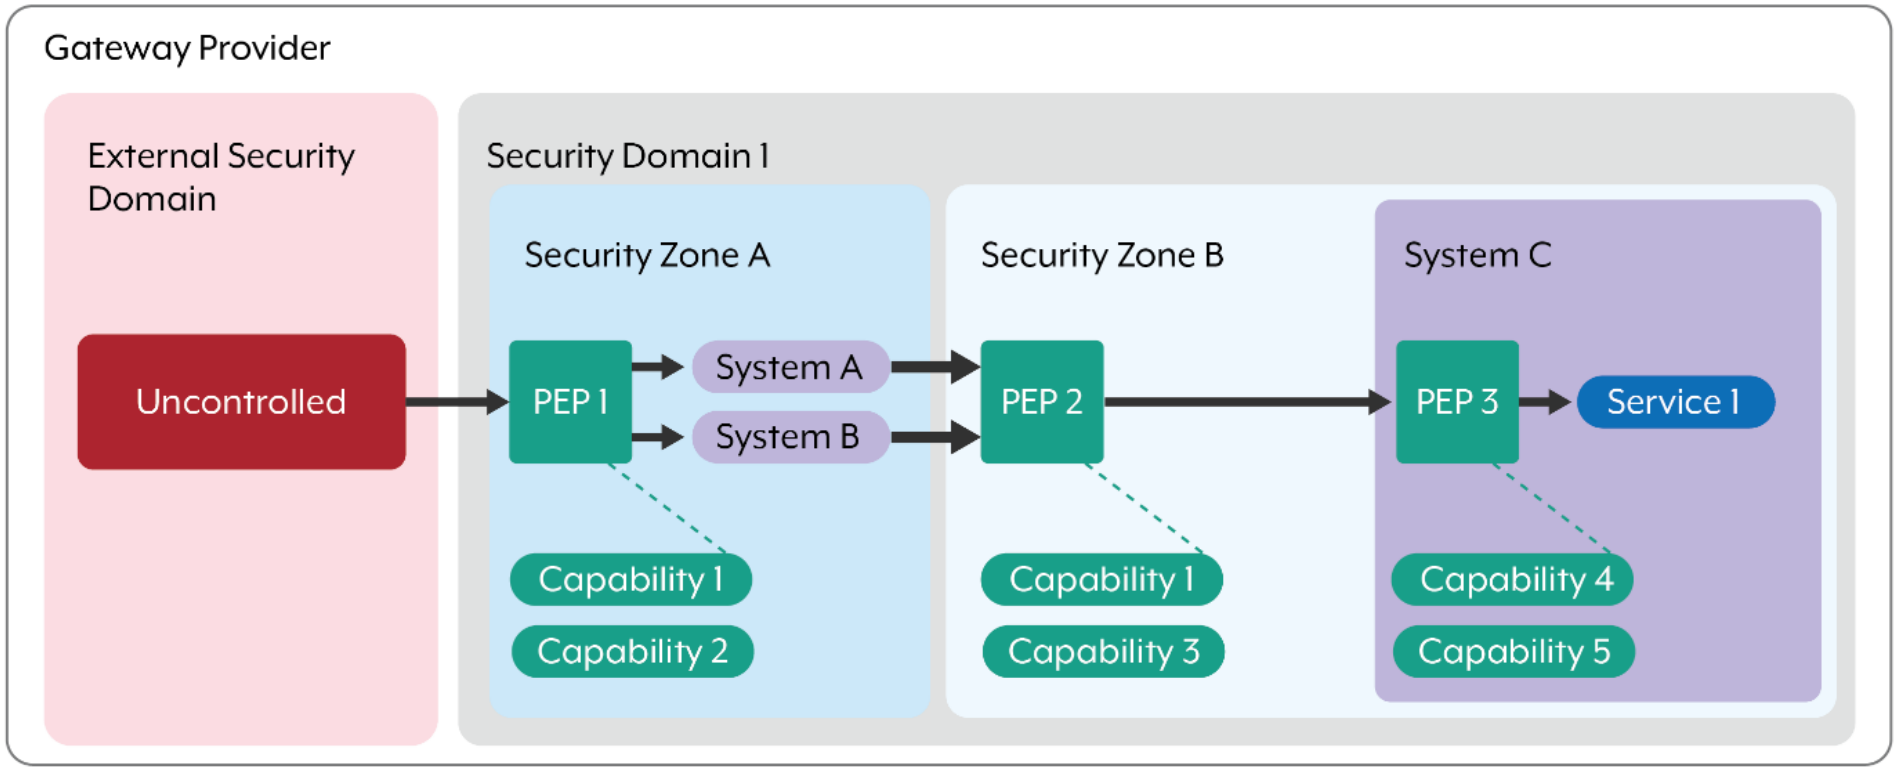

Figure 2: Chaining PEPs

It is important to note that duplicating the same control in multiple locations does not automatically provide defence-in-depth. For example, configuring a firewall or router with access control lists (ACLs) act as a PEP, but each interface enforces a distinct policy between security zones. Simply replicating the same rule sets does not enhance layered security. PEP capabilities can be implemented in various places, including:

- gateway components (e.g. CDSs, application proxies, client or server agents)
- endpoint configurations
- network infrastructure devices.

Inadequate policy enforcement can expose entities to significant risks, including:

- unauthorised data modification or exfiltration
- covert communication channels
- compromise of trusted systems, such as a firewall policy or edge device audit logs
- service disruptions or denial-of-service
- lateral movement within internal networks.

Note: the concept of a PEP has been adapted from the National Institute of Standards and Technology's Special Publication 800-207, Zero Trust Architecture.

Gateway security principles

ASD's ACSC has developed governance-related gateway security principles that an organisation should be aware of and consider when implementing or using a gateway. These principles should be applied when designing, procuring, operating, maintaining and disposing of a gateway.

Security management is continuous

Principle: Gateway security should be actively maintained, continuously improved and regularly reassessed to remain effective in a dynamic threat landscape.

Tactics, techniques, mitigations, technologies and better practices evolve over time. Gateways should form part of an organisation's defence activities, which include using ongoing threat intelligence, monitoring and proactively implementing mitigations. Risk assessments for security domains should be periodically revisited and systems should be reviewed for fitness-for-purpose. Operational teams should track updates from trusted sources (e.g. ISM and vendor advisories) and incorporate the changes into gateway controls.

Risk is continuously managed

Principle: Gateway design and operation should be guided by risk-based decision-making, informed by threat modelling and aligned with the organisation's broader risk management framework.

Threat modelling supports identifying risks and mitigating gaps. Using frameworks like MITRE ATT&CK helps organisations map gateway risks to specific adversarial behaviours, to controls and detection strategies, and to identify where ISM-aligned mitigations are needed. However, no framework is exhaustive. Each organisation must conduct its own risk analysis tailored to its systems, environment and tolerance. Design decisions must reflect obligations under the *Public Governance, Performance and Accountability Act 2013* (PGPA Act) and be documented within the entity's risk oversight structures. For more information on threat modelling, refer to Gateway security guidance package: Gateway operations and management.

The invisible cannot be protected

Principle: Organisations should maintain comprehensive visibility of data flows between security domains to enable accurate policy enforcement and risk management. Visibility can be provided either through gateways within the given security domain, or a combination of gateway capabilities within the security domain combined with appropriate and trusted visibility sources from external security domain(s).

An organisation should maintain visibility of all data entering and exiting its security domain(s). Inbound and outbound traffic should be visible through local gateway enforcement or trusted external visibility sources, aligned with the organisation's risk posture. Encryption, while necessary, can create inspection blind spots. Where decryption and inspection are not feasible, compensating controls (e.g. endpoint control capabilities and segmentation) should be deployed to manage the risk appropriately. For more information on continuous and actionable monitoring to aid visibility, refer to Foundations for modern defensible architecture.

Gateways protect organisations and staff

Principle: Gateways should prevent unauthorised information flow, protect users from malicious content, and enforce organisational policy using default-deny policies.

A gateway should be positioned appropriately to inspect data flows with sufficient context to permit or block them based on defined policies. By default, all traffic should be denied unless explicitly allowed. Controls should help protect both systems and personnel by addressing the risks of insider threat, unintentional policy breaches and external attacks.

Plan for security flaws

Principle: Gateways should be designed to tolerate control failures and allow for rapid deployment of compensating measures to maintain resilience.

Systems such as edge devices and user endpoints are inherently exposed and more likely to be targeted or compromised. Therefore, security design should focus on hardening and anomaly detection based on known adversary tradecrafts.

Compensating controls, including architectural measures, can limit the ability of a threat actor to maintain persistence and move laterally. All systems, including gateways, are susceptible to vulnerabilities. Defence-in-depth should be considered and applied to minimises adversary lateral movement and to ensure no single point of failure leads to an undetectable compromise.

Prior to implementing new systems, organisations should undertake table-top exercises in order to develop or update incident response capabilities, such as cyber security incident response plans (CIRP), and test them. To the extent reasonably practicable, control design and placement should anticipate subversion of critical components noting that the organisation may have to plan for and respond to a number of techniques used by malicious actors.

All systems are susceptible to bugs and security issues. Plans should be made in advance on how to deal with them. Pre-emptively building cyber resilience and response capabilities into systems can make it easier to deal with events as they arise. For example, being able to apply patches quickly, or disable a device known to be actively exploited.

Gateways can support resilience by restricting traffic or applying monitoring rules while remediation occurs, ensuring continuity during security events. For more information on resilient networks, refer to Foundations for modern defensible architecture.

Balance business and security

Principle: Gateways should enforce consistent security policies that also support diverse needs and operational contexts across security domains.

Security should be balanced with operational efficiency. Effective security balances business risk appetite and security policy objectives to empower business units to meet their operational objectives. Each security domain will have unique systems, security requirements and business objectives.

Environments such as development, sandboxing or research domains may require different levels of control. An organisation should demonstrate consistency in its approach to risk. Security policy exemptions, and inconsistent architectural and security capabilities in different systems will introduce risk to an organisation.

An organisation's gateway should apply security policy consistently. An organisation should consider what additional governance and oversight mechanisms are needed where multiple gateways are deployed, or where security policies are enforced through different technology stacks.

Risk cannot be outsourced

Principle: An organisation owns its security risk, even if the organisation outsources or transfers, in full or in part, the implementation responsibilities.

The Department of Home Affairs' Protective Security Guidance for Executives states that procuring goods and services does not transfer the operational risk from the Commonwealth. Security controls are intended to reduce an organisation's risk, but an organisation cannot eliminate all risks. This means that the extent of the organisation's security implementation responsibilities will vary depending on the type of deployment methodologies, services and contractual terms involved. Nevertheless, the organisation will always *own* the risk. An organisation inherits risk from dependent and underlying systems, including dependent services across multiple providers (e.g. MSPs, hyper-converged and multi-vendor cloud deployments).

Each organisation will have different responsibilities, threats, risks, legislative requirements, and obligations within its service delivery ecosystem. Actions should not be taken, whether contractual or otherwise, that prevent an organisation from meeting its responsibilities and obligations. All suppliers, service providers, and organisations have unique systems and responsibilities. This means an organisation must look at their risks by understanding the entire ecosystem in order to successfully use visibility, detection and prevention capabilities.

An organisation that enters into a contract that is not consistent with the ISM and the guidance provided in this document may be at risk of not fulfilling the requirements of the PSPF and the *PGPA Act*.

Placement of gateways

Organisations should view gateways as a suite of capabilities rather than a fixed network architecture. This approach provides greater flexibility in designing solutions tailored to specific environments while still maintaining a strong focus on risk reduction.

A gateway should remain the sole authorised pathway for transferring information into and out of a security domain. Depending on business and operational needs, an organisation may implement multiple gateways, such as for:

- internet connectivity
- cloud services
- voice and unified communications
- partner or supplier integrations.

Gateway designs should align with business requirements, particularly in terms of availability, confidentiality and integrity. Not all services require high availability, but critical business functions may demand greater resilience and throughput. In such cases, deploying a specialised or segmented gateway may be more effective than a single, centralised design.

All gateway-related design decisions should be documented, including the rationale behind them. These decisions should:

- reflect the principles and controls outlined in this guidance
- capture the implementation architecture and associated risk treatments

- be regularly reviewed by system owners and the authorising officer responsible for accepting residual risks.

Organisations should leverage the expertise of internal enterprise architects, and engineering and operations teams, and the advice of external experts (such as consultants, service providers or IRAP assessors).

Cloud-based gateways

The following sections should be read in conjunction with the Cloud assessment and authorisation and Cloud assessment and authorisation FAQ. These publications provide an introduction to several key cloud concepts including the shared responsibility model and IRAP assessments of cloud-based gateway capabilities. The guidance is also potentially applicable to other outsourcing scenarios, including most MSP offerings.

Cloud-native capabilities

As cloud architectures and API-driven integrations become more prevalent, organisations will have new opportunities to embed security capabilities directly into data flows. These capabilities can enhance both security posture and user experience.

Service-native integrations

Cloud-native gateways (a type of service-native integration) should leverage cloud-native services rather than retrofitting traditional, monolithic architectures. Service-native integrations enable stronger, context-aware security outcomes by allowing two-way interactions between security services and business applications. The following are some key benefits:

- Continuous protection instead of point-in-time protection. For example, a traditional gateway may only apply signature-based detection against payloads at the time data enters the security domain; however, a service-native integration could scan all the data stored within the service whenever a new signature is added.
- Availability risk. For example, instead of potentially silently dropping an inbound email at the gateway, a user could be provided with an application interface to gain visibility across blocked emails and have the option of taking appropriate streamlined action against false positives.
- Integrated security context. This provides a broader and more integrated view of security as the controls are applied within the context of the services fulfilling business objectives.

Inter-organisation collaboration within a single cloud service provider

When two organisations (Enterprise A and Enterprise B) use the same underlying cloud service, and data flows occur between two different security domains (refer to **Figure 3**), then both organisation's security policies must be enforced. This results in both organisations having visibility

and control of data flows. It is therefore recommended that relevant application logs and context related to the specific data flows traversing both security domains are made continuously available to both participating organisations. Notably, both Enterprise A and Enterprise B could apply their own additional capabilities to the data flows through their respective PEPs. For example, additional message filtering or content conversion security capabilities via an API integration.

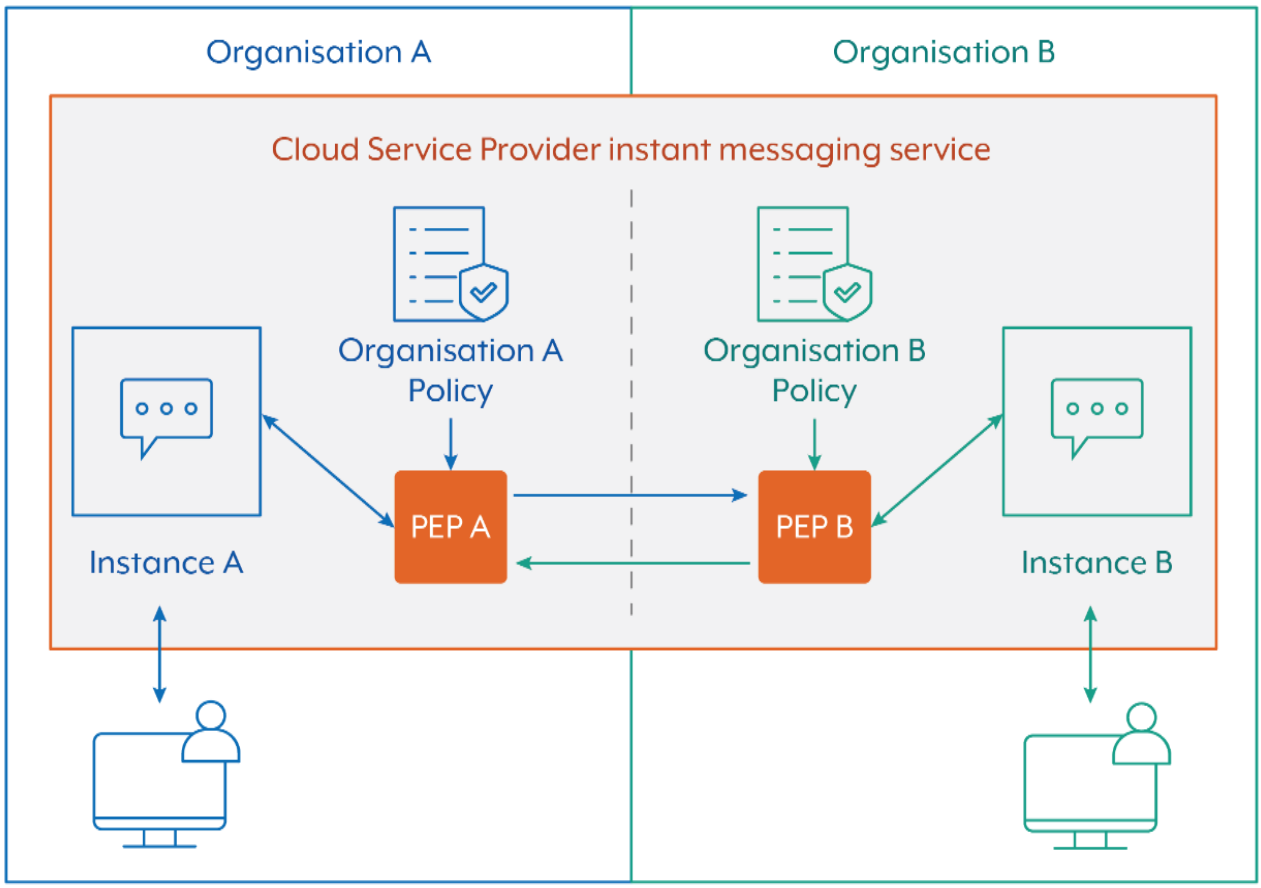

Figure 3: Inter-organisation collaboration within the one cloud service

Figure 4 shows a similar example, but here the two consumer organisations (Enterprise A and Enterprise B) are relying on a shared common security domain, such as found in some Whole of Australian Government (WoAG) services. In this example, the WoAG providing organisation has decided to deploy an additional PEP that can enforce the security policy of all consumer organisations at the same time. It would be expected that the relevant information from the providing organisation is shared with the consumer organisations, and that the consumer organisations would have direct control over their own security policy.

Each consumer organisation is responsible for its own risk assessment and granting authorisation to operate (ATO) of the common service. In effect, it authorises each consumer organisation's data to be transmitted, shared or processed by the shared service, while ensuring that the common service and PEP are in line with its own risk tolerances. For simplicity, additional PEPs within the consumer organisations have been omitted but these organisations can deploy them.

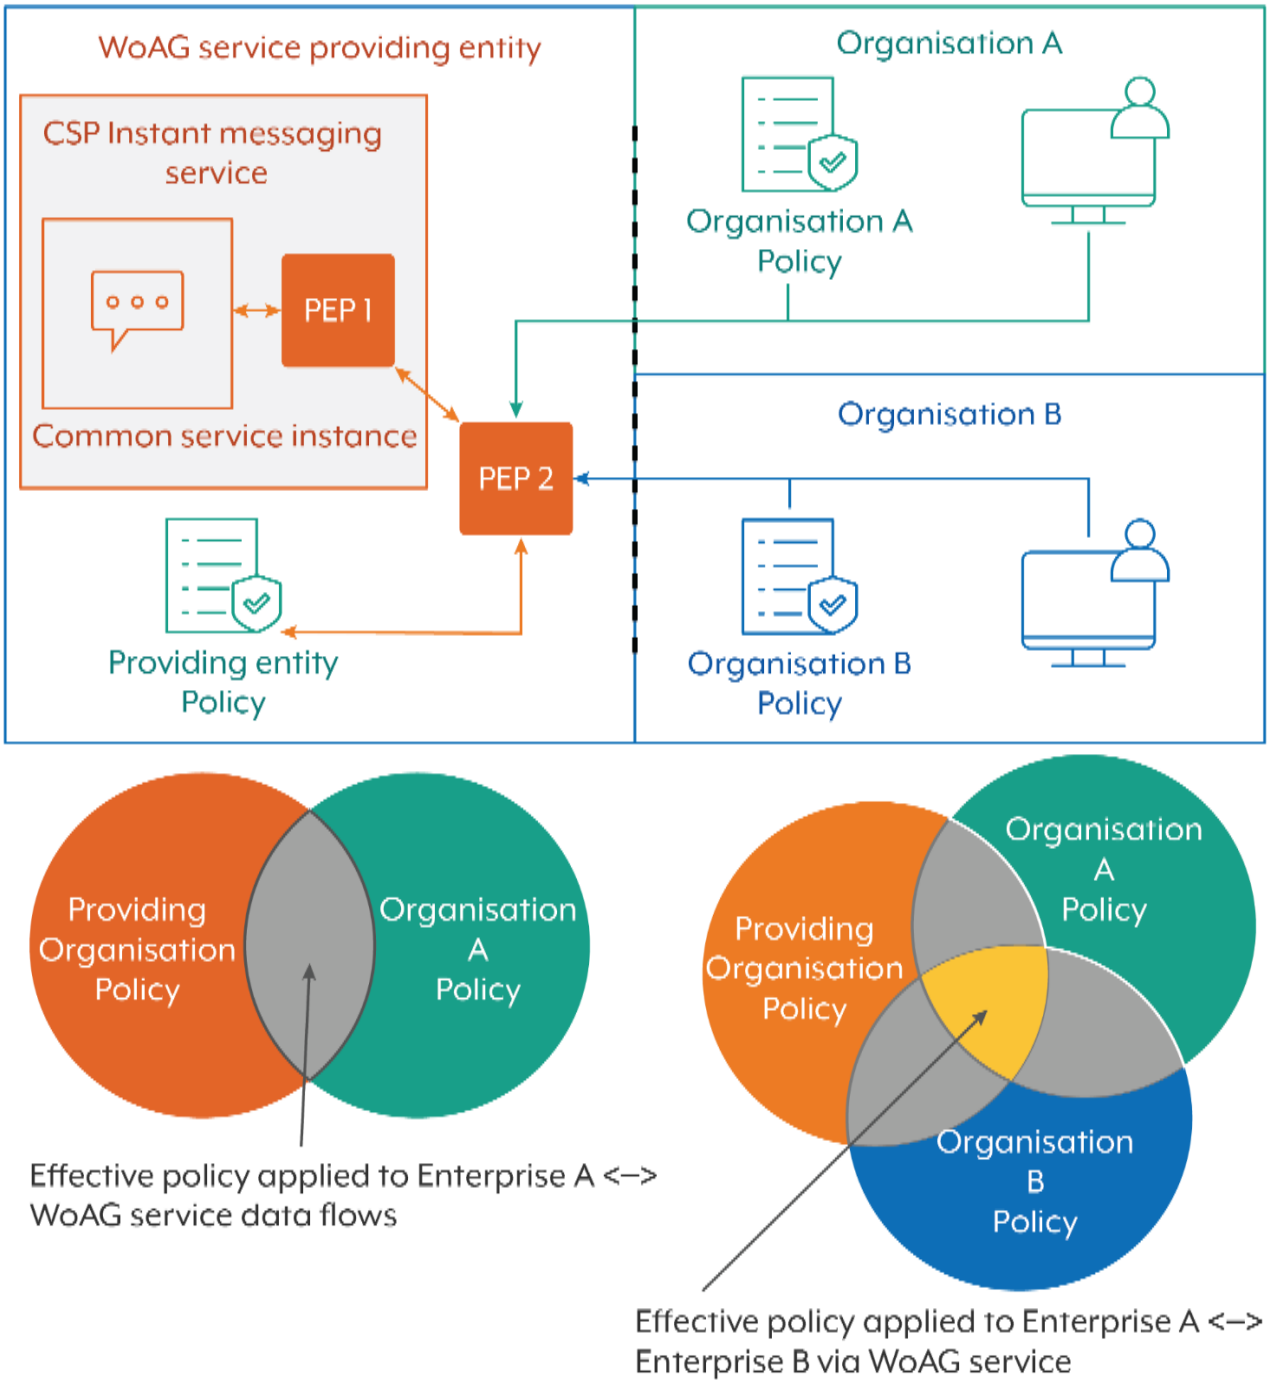

Figure 4: Shared common security domain

Security zones

Security zones are logical network segments used to group systems and services based on trust, sensitivity or function. In gateway architectures, they help enforce separation between environments and enable the application of targeted security controls.

Gateways enforce these controls at the boundaries of each security zone, providing inspection, filtering and logging to manage and secure inter-zone data flows. Importantly, different security zones can have different security controls tailored to the risk and business context of the zone.

This flexibility is particularly valuable in enabling specific scenarios while maintaining appropriate security postures. **Figure 5** shows an example application.

- In security zone #1, the PEP applies all the appropriate protections to safeguard sensitive enterprise data held within zone #1.
- In security zone #2, the PEP provides logging capabilities but does not inspect or decrypt traffic (e.g. transport layer security between a virtual browser and the internet). This can be done, as there is no sensitive enterprise data within the lower security zone to protect.

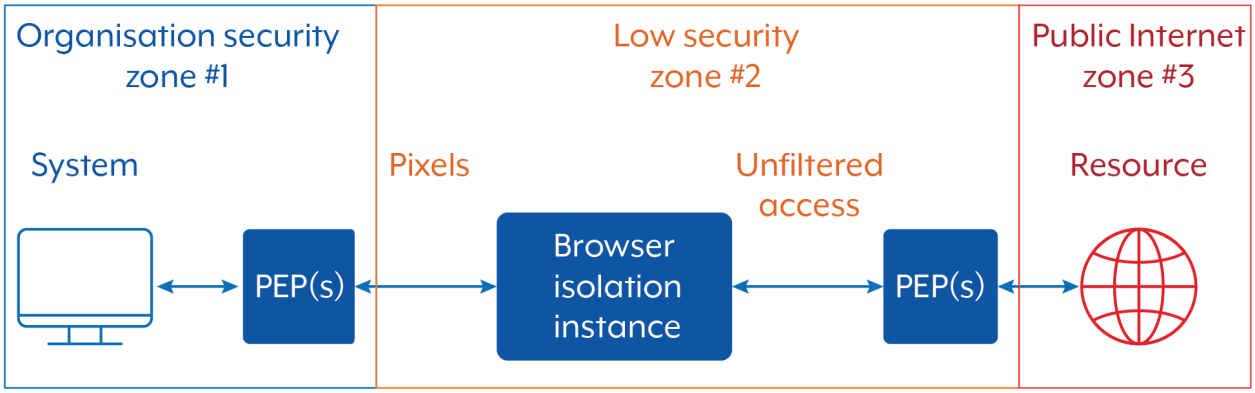

Figure 5: Security zones

This pattern is also applicable to sandbox, training or development environments where sensitive data is not present. However, data moving from a low security zone to a high security zone must be subject to strong security controls. For example, any code from the lower security zone would need to be treated as untrusted. Refer to the non-persistent virtualised sandboxed environment strategy from Strategies to mitigate cyber security Incidents: Mitigation details. For more information on network segmentation, refer to Implementing Network Segmentation and Segregation.

Security Service Edge

SSE is the collective term for several different cloud-based security services that can be used to protect data. SSE capabilities can be used to enforce gateway policy associated with a security domain, with common services including:

- cloud access security brokers
- secure web gateways
- Firewall-as-a-Service
- Zero trust network access.

SSE services are typically deployed as part of a hybrid architecture (on-premises and cloud), directly enforcing gateway policy for different workloads. They can also complement other PEPs at

security boundaries within the domain. Isolated cloud-based deployments (with no on-premises integration) can also have gateway policy entirely enforced by SSE services.

Overlaying SSE services onto SD-WAN virtualisation is a key aspect of Secure Access Service Edge (SASE) frameworks. SASE combines network and security services into a unified, cloud-delivered network security capability that can provide a flexible and effective means of meeting security requirements, optimising user experience and saving cost.

Shared responsibility and trust

In providing cloud PEP solutions, there are shared responsibilities between the different parties managing all aspects of the cloud solution. For example, one party may be predominantly responsible for security capabilities but, then again, the different aspects of security capabilities may be shared between parties. In all cases, the cloud consumer must retain visibility across data flows and further secure data flows as they see fit. For example, filtering data flows containing files based on a YARA rule, or HTTP requests based on HTTP headers. Using MSPs and CSPs requires organisations to place some degree of trust in the service provider and take steps to ensure that an outsourced service provider is managing shared risks appropriately.

Gateway security is a shared responsibility between a gateway provider, the gateway consumer, and any other third parties who are involved in providing the complete gateway solution, including cloud platforms.

It is important to understand that PEPs can be effectively layered together across data flows. For example, one cloud service may use the security capabilities provided by another service offered from the same CSP, or the cloud consumer may deploy its own additional capabilities between the client and the cloud service resource. When layering PEP capabilities, it is important that they do not lower the security baseline provided by the CSP as this can introduce new weaknesses. The connection between the two security domains also needs to be secure. Furthermore, cloud consumers should also consider what else impacts a desired architecture when evaluating it, such as user experience or availability. **Figure 6** outlines how a cloud-native PEP may need to be supplemented by a customer-deployed PEPs in order to appropriately enforce security domain separation.

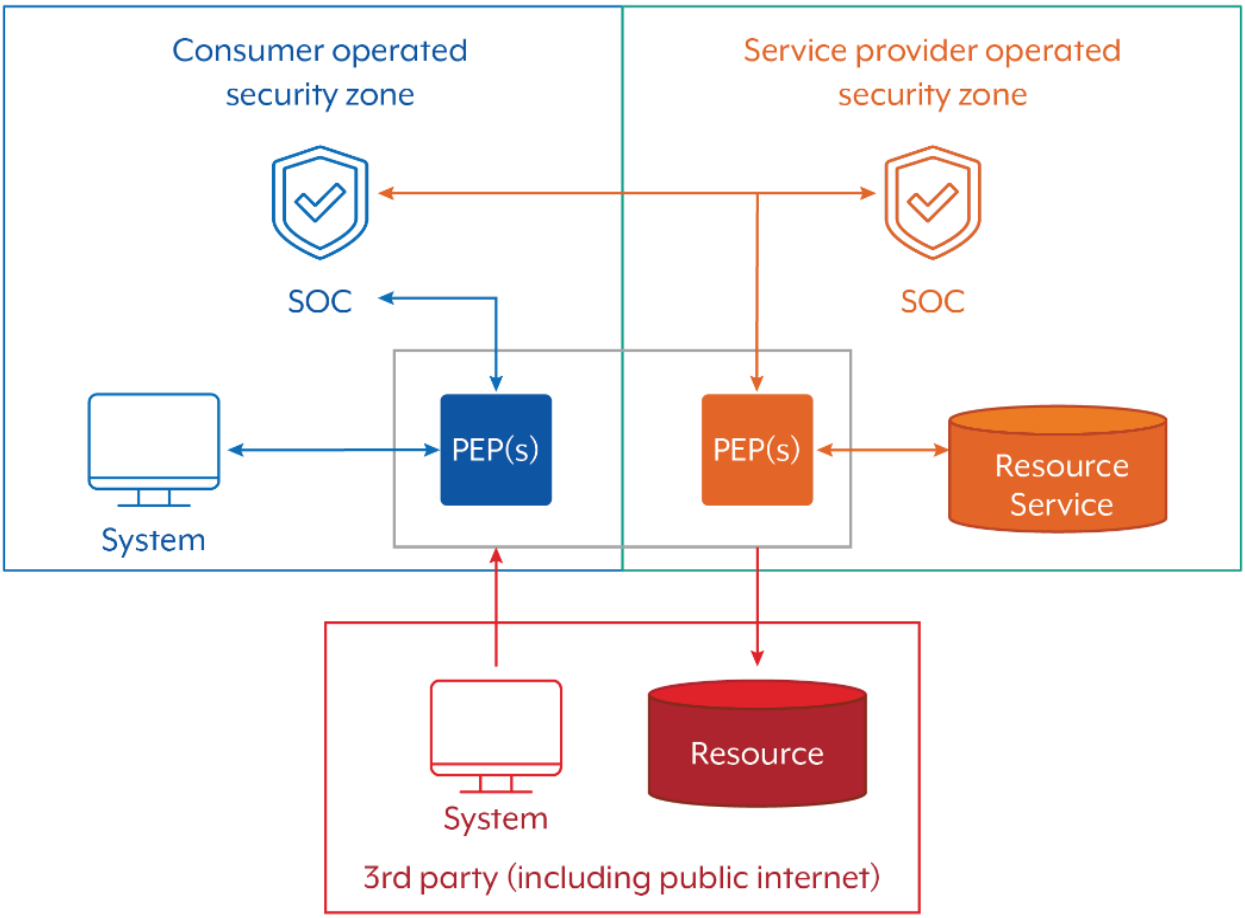

Figure 6: PEPs and security domains

Risk considerations for outsourcing gateway services

An organisation's risk management process should consider and balance the benefits of an in-house gateway with the cyber security risks associated with outsourcing management of a gateway to a service provider. A risk assessment should consider whether the organisation, as a gateway consumer, is willing to trust its reputation, business continuity and data to a gateway service provider. Consideration should also be given to recovering the consumer's data should it be insecurely transmitted, stored or processed.

The risks should be weighed against the security capabilities available in-house compared to an outsourced arrangement. In some cases, it may not be possible to independently verify whether a gateway provider is adhering to contractual terms, which leaves the consumer having to rely on third-party audits to find out. Consumers should consider which documents to request from a gateway service provider in order to assess risks, and determine whether the contents of the documents provide the appropriate information to satisfy their assurance requirements or to identify gaps in their visibility.

Roles and responsibilities

Under the PGPA Act and the PSPF, an NCE's Accountable Officer is the position responsible for understanding and managing risk within their organisation, and for granting systems an ATO. An organisation's architecture, engineering and operations teams all play an important part in

providing information that helps the organisation develop an understanding of its risk profile and the operating threat landscape. NCEs should review both the PSPF and the ISM for more information on determining formal organisational roles and responsibilities.

As part of using a gateway provider's services, gateway consumers need to understand their own responsibilities, as well as the responsibilities of the other parties involved in delivering the complete gateway solution. This includes understanding each party's responsibilities for securing the gateway. For example:

- responsibilities for cyber security policy
- incident detection and response
- data retention and backup
- monitoring
- system hardening
- patching and encryption.

In some of these examples, one party may be entirely responsible, or different aspects may be shared between parties.

As part of the IRAP security assessment report for a gateway, IRAP assessors are to document or otherwise identify which party is responsible for managing risks associated with key aspects of each gateway solution in scope of the assessment. This provides gateway consumers with a clear understanding of the different responsibilities that each party has for securing the gateway solution, including their own.

Regardless of the shared responsibility model, gateway consumers remain accountable for their data, including taking steps to ensure the data is appropriately secured. Organisations should verify that gateway controls are in place, operating effectively and providing the required visibility and capabilities. Access requirements should be proportional to the gateway service, and organisations should only provide MSPs and CSPs with the access required to operate the gateway environments.

Tenancy considerations

It is important that services exposing interfaces for clients to use also provide mechanisms for clients' PEPs to appropriately restrict traffic to only the expected tenant or instance. For example, by providing either:

- a unique tenant domain or URL path
- a combination of IP address and port, with a unique HTTP header on all requests
- an SDN construct including unique virtual network interface
- a virtual network route that can be enforced by a PEP on the client side.

If any system component can make an outbound request to an unauthorised or uncontrolled tenancy (even if the client has a valid signed payload), this could enable an outbound command-and-control path that the consumer would be unable to identify or constrain.

Inbound and outbound data flows need to be protected by PEP capabilities. CSPs need to consider, particularly for Software-as-a-Service and Platform-as-a-Service with multi-tenant services, how they will enable each consumer's own security policy applied across both outbound and inbound data flows. A CSP should support one or more methods. As an example, if the responsibility for some of the security controls falls on the consumer to implement, then the service they are receiving should support a consumer-provided proxy via either a SDN construct (including a gateway load balancer construct as referred to by some CSPs) or application construct (including API-based event bus for policy enforcement).

Cyber security incident reporting

When using a service provider, cyber security incident reporting processes should be formalised through a shared responsibility model. For example, if a security policy violation was discovered through a gateway system, both parties should clearly understand how this will be reported to the appropriate authorities. Organisations should consider contractual obligations to report any cyber security incident or breach to the gateway consumer.

Early detection of a cyber security incident is critical to expediting containment and recovery, including timely reporting to an organisation's Chief Security Officer or Chief Information Security Officer. The PSPF outlines the requirements and obligations of NCEs for the reporting of security incidents, including cyber security incidents. Refer to the ISM Guidelines for cyber security incidents.

CSPs, MSPs and their customers will benefit from contract arrangements that clearly define responsibilities.

- CSPs and MSPs, when negotiating the terms of a contract with their customer, should provide clear explanations of the services that the customer is purchasing, services that the customer is not purchasing, and all contingencies for cyber security incident response and recovery.
- Customers should ensure that they have a thorough understanding of the security services that their service provider is providing and address any security requirements that fall outside the scope of the contract. If contracting to an MSP, contracts should detail how and when MSP notifies the customer of a cyber security incident that affects the customer's environment or data.
- Customers should ensure that they gain trust in the service delivery models provided to them by MSPs and CSPs, particularly with respect to the countries or jurisdictions in which their support service teams may be based.
- Contracts should clearly define sanctions, penalties and exit clauses for not fulfilling contract terms, noting that penalties rarely compensate an organisation for the losses incurred as a result of a cyber security incident.

The Cloud Security Alliance describes a shared responsibility model as follows:[2]

> *In a traditional data centre model, you are responsible for security across your entire operating environment, including your applications, physical servers, user controls, and even physical building security. In a cloud environment, your provider offers valuable relief to your teams by taking on a share of many operational burdens, including security. In this shared responsibility model, security ownership must be clearly defined, with each party maintaining complete control over those assets, processes, and functions they own. By working together with your cloud provider and sharing portions of the security responsibilities, you can maintain a secure environment with less operational overhead.*

Gateway visibility and telemetry

Event logs and system telemetry support the continued delivery of operations and improve the security and resilience of critical systems by enabling network visibility. Best practices for event logging and threat detection outlines best practice for event logging and threat detection for cloud services, enterprise information technology (IT) networks, enterprise mobility and operational technology (OT) networks. It also provides recommendations to improve an organisation's resilience in the current cyber threat environment, with regard for resourcing constraints.

Gateway logs and telemetry may come from a variety of sources:

- systems and applications (e.g. authentications, operational logs, headers and geo-location)
- network infrastructure telemetry (e.g. IPFIX/NetFlow/JFLOW/SYSFLOW)
- traffic payload artefacts (e.g. packet captures of decrypted content and remote object inspection or 'ICAP').
- health and performance monitoring tools.

An organisation's Security Operations Centre (SOC) needs to have access to logs, telemetry and other artefacts produced by their gateways to ensure effective incident response.

High-value logs from a gateway include traffic to and from:

- identity and authentication systems
- credential and access management systems
- DNS servers
- web proxies
- WAFs

[2] https://cloudsecurityalliance.org/blog/2020/08/26/shared-responsibility-model-explained

- mail relays
- load balancers
- IPS/IDS
- remote access solutions
- sandbox and ICAP services
- other security services.

Additionally, network flow telemetry data from gateways can enhance security capabilities.

While full packet capture and storage of all gateway traffic is often impractical, organisations should build the capability to decrypt and store network traffic and test the capability to conduct targeted data capture of higher risk sessions and services. This could be achieved either through random sampling or conducted in response to a cyber security incident or investigation. Processes should be developed to provide captured data to the organisation's SOC, or typically a cyber security incident response team, in a format that facilitates forensic analysis. This data could be in the form of a packet capture file or streamed data.

During a cyber security incident, the availability of logs becomes more important to cyber security incident response teams. Gateways typically have the functionality for logging to be tuned to increase verbosity, which provides more visibility and insights. An organisation should develop and test procedures to increase and decrease log verbosity in response to different scenarios.

An organisation's SOC will be interested in more than just gateway events. An organisation should also collect and analyse internal network flow telemetry and endpoint event logs (servers and workstations), behavioural analytics of processes, endpoint security logs, and internal security capability events. **Figure 7** describes a mechanism where logs and telemetry, generated by multiple gateway services, can be forwarded to several separate stakeholders.

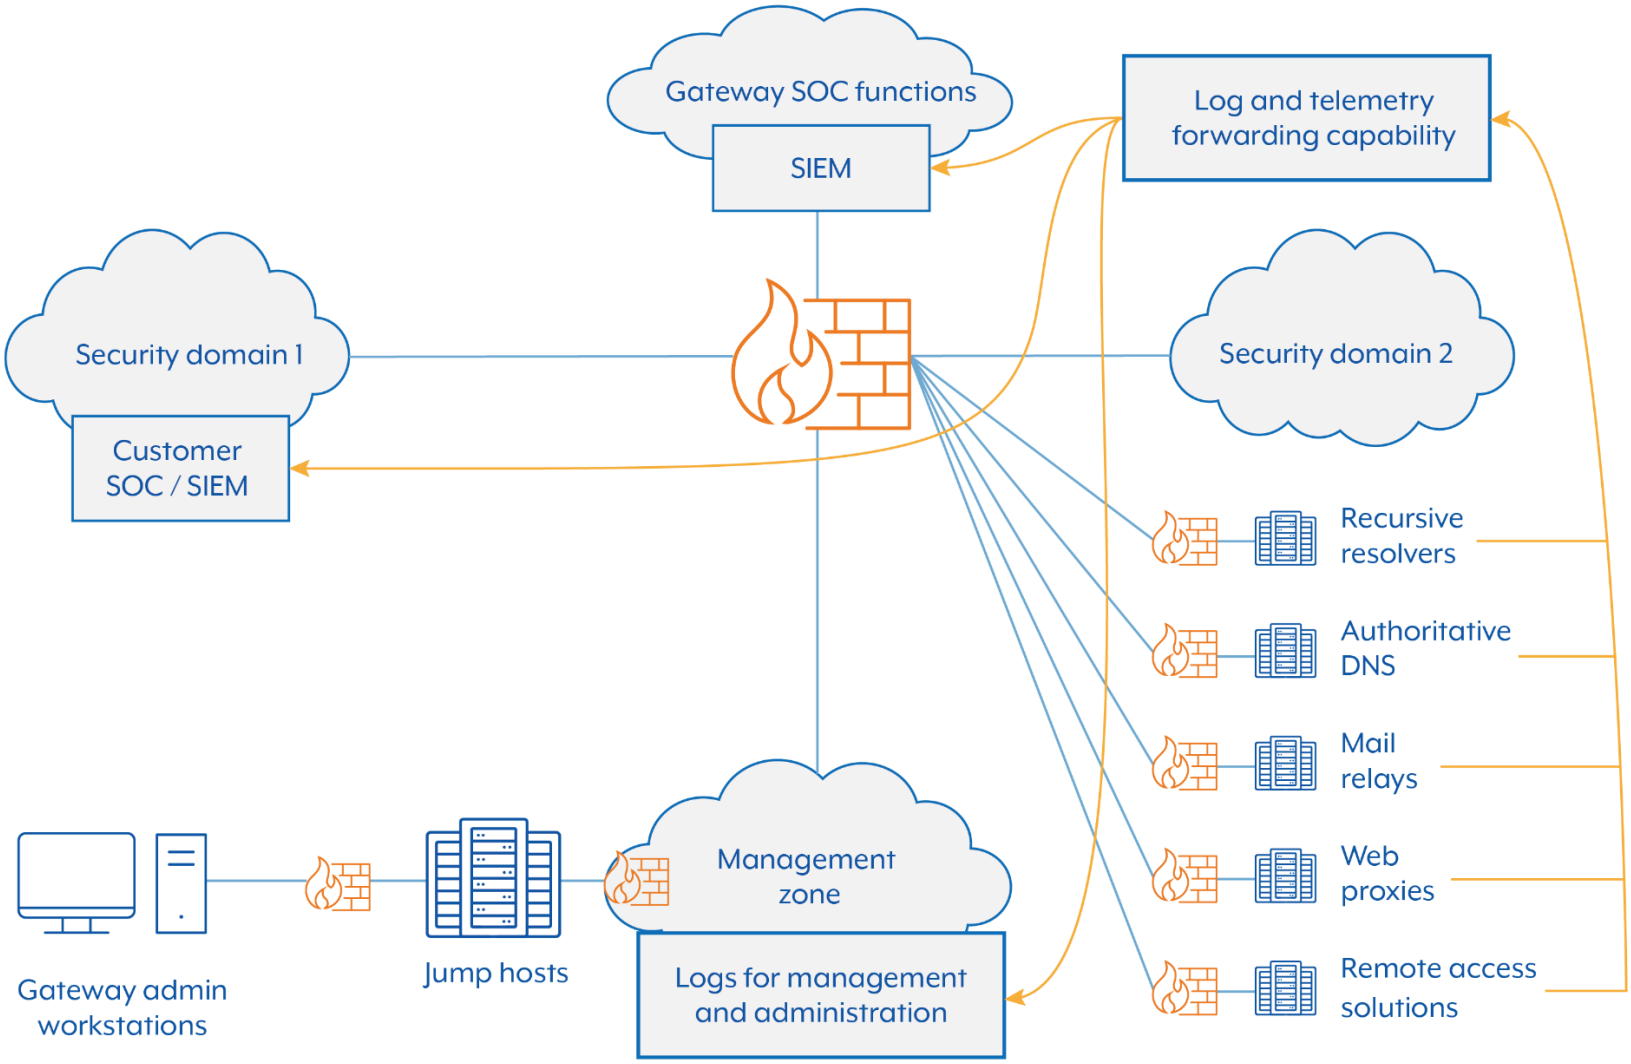

Figure 7: Logs and telemetry transfer from gateway systems

Cyber security incident response

Managing responses to cyber security incidents is the responsibility of affected organisations. As such, all organisations should have a CIRP to ensure an effective response and prompt recovery if system controls do not prevent a cyber security incident from occurring. This plan should be regularly tested and reviewed. Organisations can refer to the following guidance when developing their CIRP:

- Cyber security incident response planning: Executive guidance
- Cyber security incident response planning: Practitioner guidance.

Gateways are critical control points between internal systems and the internet and their compromise can disrupt business operations or expose sensitive data. As such, an organisation's CIRP should include scenarios such as a misconfigured firewall, unpatched mail relay, or compromised VPN concentrator being used to pivot to internal environments.

Gateways can help an organisation respond to vulnerabilities. For example, as a temporary measure, a network intrusion prevention system or WAF may be able to actively block known attacks against vulnerable systems that require patching (also called 'virtual patching'), and firewalls can ingest cyber threat intelligence (CTI) to automatically block sources of malicious traffic.

A gateway service provider's response to the following questions can indicate their ability to handle cyber security incidents.

Timely service provider support:

- Is the service provider readily contactable and responsive to requests for support during an incident, and is the maximum acceptable response time captured in the service-level agreement (SLA)?
- Is the support provided locally, from a foreign country, or from several foreign countries?
- Which mechanism does the service provider use to obtain a real-time understanding of the security posture and configuration of the service provider's services?

Service provider's CIRP:

- Does the service provider have a CIRP that specifies how to detect and respond to cyber security incidents, that impact on the gateway provider or their customers, in a way that is consistent with the ISM's cyber security incident handling procedures?
- Can the organisation review the service provider's CIRP?

Training of service provider's employees:

- Which qualifications, certifications and regular information security awareness training do the service provider's employees receive to know how to use the service provider's systems in a secure manner, and to identify potential cyber security incidents?

Notification of cyber security incidents:

- Will the service provider notify the organisation through secure communications of cyber security incidents that are more serious than an agreed threshold, especially in cases where the service provider might be liable?
- Will the service provider automatically notify law enforcement or other authorities, who may confiscate computing equipment used to store or process an organisation's data?

Extent of service provider support:

- To what extent will the service provider assist the organisation with investigations if there is a security breach, such as an unauthorised disclosure of the organisation's data, or if there is a need to perform legal electronic discovery of evidence?

Access to logs:

- How does an organisation obtain access to time-synchronised audit logs and other logs to perform a forensic investigation, and how are the logs created and stored to be suitable evidence in a court of law?

Cyber security incident compensation:

- How will the service provider adequately compensate a consumer if the service provider's actions or inaction, faulty software or hardware contributed to a security breach?

Data spills:

- If data that an organisation considers too sensitive to be stored in a location is accidentally placed in that location (referred to as a data spill), how can the spilled data be deleted using forensic sanitisation techniques?
- Is the relevant portion of physical storage media 'zeroed' whenever data is deleted? If not, how long does it take for deleted data to be overwritten by consumers as part of normal operation, noting that some service providers have significant spare unused storage capacity?
- Can spilled data be forensically deleted from the service provider's backup media?
- Where else is the spilled data stored, and can it be forensically deleted?

For more information refer to the Cyber Security Incident Response Planning: Executive Guidance, ASD, Cyber Security Incident Response Planning: Practitioner Guidance, and ASD, Cyber Incident Management Arrangements for Australian Governments.

Architect for maintenance

It is common for vulnerabilities in edge devices to be exploited within hours of a security patch becoming available. An organisation should adopt modern defensible architecture and business processes to ensure patches and mitigations can be deployed with very short notice. Organisations should generally prioritise unplanned patching outages over the risk of compromise. Where patching is not feasible, gateways may provide organisations with options for compensating controls for vulnerable systems (e.g. WAF or intrusion prevention solution signatures).

Organisations should align change management processes to support unplanned emergency patching scenarios. As SecDevOps processes are more widely adopted in supply chains, organisations should consider if their existing change management processes are introducing barriers for operational teams. For example, CSPs do not have obligation to a consumer's change management processes, but they do proactively implement processes to patch infrastructure to protect their consumers, which is consistent with the service providers' shared responsibilities model. For more information refer to the Patching Applications and Operating Systems.

Cyber threat intelligence

A gateway should use CTI to proactively identify and respond to threats and vulnerabilities that are relevant to the consumer(s) of the gateway. A gateway should allow an organisation to both derive and consume CTI from its operation and generate data that allows security analysts to produce CTI. For more information on how to use a gateway to generate and use CTI, refer to the *Gateway security guidance package: Gateway technology guides*.

Gateways should be architected to generate and use high-confidence CTI. Mature gateways may also be able to take automated responses to security risks based on high-confidence CTI it receives. This can include automatically checking if an Indicator of Compromise (IOC) has been observed or taking action to block an attacker's IP address or email address.

Some CTI, especially IOCs, can be highly perishable and will only be relevant and actionable for a short time (less than a day, sometimes only hours). CTI should be able to inform organisations of what they should be responding to as a priority. As gateway teams, SOC, and incident response teams have limited resources, organisations should prioritise and consider environmental contexts to assist with triaging CTI. Organisations should make decisions based on CTI (e.g. by changing security settings, adding new security capabilities, changing business processes, adjusting training and policy, or making architectural changes).

Priority services for security visibility

NCEs may choose to consume gateway services through an existing commercial or government gateway provider. While there are many services that gateway providers supply, ASD's ACSC recommends NCEs prioritise uplifting the security capabilities in the following five services:

- recursive resolvers

- web proxies
- mail relays
- reverse proxies
- remote access.

The Gateway security guidance package: Gateway technology guides contain service-specific advice on different types of gateway services (e.g. mail relays, DNS, web proxies and remote access). They also include why organisations should focus on uplifting the security of the above five services.

More information

For more information on topics covered in this guidance, refer to the following ASD's ACSC publications:

- Information security manual
- Patching applications and operating systems
- Cyber security incident response planning: Executive guidance
- Cyber security incident response planning: Practitioner guidance
- Cyber incident management arrangements for Australian governments
- Foundations for modern defensible architecture

Contact us

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

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

Disclaimer

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

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

Copyright

Use of the Coat of Arms

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

Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre