



Information Security Manual

DECEMBER 2021

Guidelines for Networking

Network design and configuration

Network documentation

It is important that network documentation accurately depicts the current state of a network. This typically includes network devices such as firewalls, data diodes, intrusion detection and prevention systems, routers, switches, and critical servers and services. Furthermore, as this documentation could be used by an adversary to assist in compromising a network, it is important that it is appropriately protected.

Security Control: 0516; Revision: 4; Updated: Sep-18; Applicability: All

Network documentation includes a high-level network diagram showing all connections into the network; a logical network diagram showing all network devices, critical servers and services; and the configuration of all network devices.

Security Control: 0518; Revision: 4; Updated: Sep-18; Applicability: All

Network documentation is updated as network configuration changes are made and includes a 'current as at [date]' or equivalent statement.

Security Control: 1178; Revision: 3; Updated: Sep-18; Applicability: All

Network documentation provided to a third party, or published in public tender documentation, only contains details necessary for other parties to undertake contractual services.

Network segmentation and segregation

Network segmentation and segregation is one of the most effective security controls to prevent an adversary from propagating through a network and accessing target data after they have gained initial access. Technologies to enforce network segmentation and segregation also contain logging functionality that can be valuable in detecting an intrusion and, in the event of a compromise, isolating compromised devices from the rest of a network.

Network segmentation and segregation involves separating a network into multiple functional network zones with a view to protecting important data and critical services. For example, one network zone may contain user workstations while another network zone contains authentication servers. Network segmentation and segregation also assists in the creation and maintenance of network access control lists.

Security Control: 1181; Revision: 4; Updated: Jun-21; Applicability: All

Networks are divided into multiple functional network zones according to the sensitivity or criticality of data or services.

Security Control: 1577; Revision: 0; Updated: Jul-20; Applicability: All

Organisation networks are segregated from service provider networks.

Using Virtual Local Area Networks

Virtual Local Area Networks (VLANs) can be used to implement network segmentation and segregation as long as the networks belong to the same security domain. In such cases, if a data spill occurs the impact will be lesser than if a data spill occurred between two networks of different classifications or between an organisation's network and public network infrastructure. Should organisations choose to risk manage implementing VLANs between networks belonging to different security domains, such as at the same classification, additional security controls relating to not using VLAN trunking and terminating VLANs on separate physical network interfaces of networks devices will apply.

For the purposes of this section, Multiprotocol Label Switching is considered to be equivalent to VLANs and is subject to the same controls.

Security Control: 1532; Revision: 2; Updated: Dec-21; Applicability: All

VLANs are not used to separate network traffic between organisations' networks and public network infrastructure.

Security Control: 0529; Revision: 6; Updated: Dec-21; Applicability: All

VLANs are not used to separate network traffic between networks belonging to different security domains.

Security Control: 1364; Revision: 3; Updated: Dec-21; Applicability: All

Network devices managing VLANs terminate VLANs belonging to different security domains on separate physical network interfaces.

Security Control: 0535; Revision: 6; Updated: Dec-21; Applicability: All

Network devices managing VLANs belonging to different security domains do not share VLAN trunks.

Security Control: 0530; Revision: 6; Updated: Dec-21; Applicability: All

Network devices managing VLANs are administered from the most trusted security domain.

Using Internet Protocol version 6

Internet Protocol version 6 (IPv6) functionality can introduce additional security risks to a network. As such, organisations using Internet Protocol version 4 (IPv4)-based networks should disable IPv6 functionality until it is intended to be used to assist in minimising the attack surface of the network and ensure that any IPv6 functionality that is not intended to be used cannot be exploited.

To aid in the transition from IPv4 to IPv6, numerous tunnelling protocols have been developed that are designed to allow interoperability between the protocols. Disabling IPv6 tunnelling protocols on network devices and ICT equipment that do not explicitly require such functionality will prevent an adversary bypassing traditional network defences by encapsulating IPv6 data inside IPv4 packets.

Stateless Address Autoconfiguration is a method of stateless Internet Protocol (IP) address configuration in IPv6 networks. Notably, it reduces the ability of an organisation to maintain effective logs of IP address assignment on a network. For this reason, stateless IP addressing should be avoided.

Security Control: 0521; Revision: 5; Updated: Sep-18; Applicability: All

IPv6 functionality is disabled in dual-stack network devices and ICT equipment unless it is being used.

Security Control: 1186; Revision: 3; Updated: Sep-18; Applicability: All

IPv6 capable network security devices are used on IPv6 and dual-stack networks.

Security Control: 1428; Revision: 1; Updated: Sep-18; Applicability: All

Unless explicitly required, IPv6 tunnelling is disabled on all network devices and ICT equipment.

Security Control: 1429; Revision: 2; Updated: Jan-20; Applicability: All

IPv6 tunnelling is blocked by network security devices at externally-connected network boundaries.

Security Control: 1430; Revision: 2; Updated: Jun-21; Applicability: All

Dynamically assigned IPv6 addresses are configured with Dynamic Host Configuration Protocol version 6 in a stateful manner with lease data stored in a centralised logging facility.

Network access controls

If an adversary has limited opportunities to connect to a network, they have limited opportunities to compromise that network. Network access controls not only prevent unauthorised access to a network but also prevent users carelessly connecting a network to another network.

Network access controls are also useful in segregating data for specific users with a need-to-know or limiting the flow of data between network segments. For example, computer management traffic can be permitted between workstations and systems used for administration purposes but not permitted between standard user workstations.

Security Control: 0520; Revision: 6; Updated: Sep-18; Applicability: All

Network access controls are implemented on networks to prevent the connection of unauthorised network devices.

Security Control: 1182; Revision: 3; Updated: Sep-18; Applicability: All

Network access controls are implemented to limit traffic within and between network segments to only those that are required for business purposes.

Network device register

Maintaining and regularly auditing a register of authorised network devices can assist in determining whether devices such as switches, routers, wireless access points and internet dongles on a network or connected directly to workstations are rogue or not. The use of automated discovery and mapping tools can assist in this process.

Security Control: 1301; Revision: 2; Updated: Aug-19; Applicability: All

A network device register is maintained and regularly audited.

Default accounts for network devices

Network devices can come pre-configured with default credentials. For example, wireless access points with an account named 'admin' and a passphrase of 'admin' or 'password'. Ensuring default accounts are disabled, renamed or have their passphrase changed can assist in reducing the likelihood of their exploitation by an adversary.

Security Control: 1304; Revision: 2; Updated: Sep-18; Applicability: All

Default accounts for network devices are disabled, renamed or have their passphrase changed.

Disabling unused physical ports on network devices

Disabling unused physical ports on network devices such as switches, routers and wireless access points reduces the opportunity for an adversary to connect to a network if they can gain physical access to network devices.

Security Control: 0534; Revision: 2; Updated: Sep-18; Applicability: All

Unused physical ports on network devices are disabled.

Functional separation between servers

Implementing functional separation between servers can reduce the security risk that a server compromised by an adversary will pose an increased security risk to other servers.

Security Control: 0385; Revision: 6; Updated: Sep-18; Applicability: All

Servers maintain effective functional separation with other servers allowing them to operate independently.

Security Control: 1479; Revision: 0; Updated: Sep-18; Applicability: All

Servers minimise communications with other servers at both the network and file system level.

Management traffic

Implementing security measures specifically for management traffic provides another layer of defence on a network should an adversary find an opportunity to connect to that network. This also makes it more difficult for an adversary to enumerate a network.

Security Control: 1006; Revision: 6; Updated: Sep-18; Applicability: All

Security measures are implemented to prevent unauthorised access to network management traffic.

Use of Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) can be used to monitor the status of network devices such as switches, routers and wireless access points. The first two iterations of SNMP were inherently insecure as they used trivial authentication methods. Furthermore, changing all default SNMP community strings on network devices and limiting access to read-only access is strongly encouraged.

Security Control: 1311; Revision: 2; Updated: Sep-18; Applicability: All

SNMP version 1 and 2 are not used on networks.

Security Control: 1312; Revision: 2; Updated: Sep-18; Applicability: All

All default SNMP community strings on network devices are changed and have write access disabled.

Using Network-based Intrusion Detection and Prevention Systems

A Network-based Intrusion Detection System (NIDS) or Network-based Intrusion Prevention System (NIPS), when configured correctly and supported by suitable processes and resources, can be an effective way of identifying and responding to known intrusion profiles.

In addition, generating alerts for data flows that contravene any rule in a firewall rule set can help security personnel respond to suspicious or malicious traffic entering a network due to a failure or configuration change to firewalls.

Security Control: 1028; Revision: 7; Updated: Aug-20; Applicability: All

NIDS or NIPS are deployed in all gateways between an organisation's networks and other networks they do not manage.

Security Control: 1030; Revision: 7; Updated: Jun-21; Applicability: All

NIDS or NIPS in gateways are located immediately inside the outermost firewall and configured to generate a log entry, and an alert, for any data flows that contravene any rule in firewall rule sets.

Security Control: 1185; Revision: 3; Updated: Sep-18; Applicability: All

When deploying NIDS or NIPS in non-internet gateways, they are configured to monitor unusual patterns of behaviour or traffic flows rather than internet-based communication protocol signatures.

Blocking anonymity network traffic

Inbound network connections from anonymity networks (such as Tor, Tor2web and I2P) to an organisation's internet-facing services can be used by adversaries for reconnaissance and malware delivery purposes with minimal risk of detection and attribution. As such, this traffic should be blocked provided it will not meaningfully impact accessibility for legitimate users. For example, some organisations might choose to support anonymous connections to their websites to cater for individuals who want to remain anonymous for privacy reasons. In such cases, it is suggested that traffic from anonymity networks be logged and monitored instead. Additionally, outbound network connections to anonymity networks can be used by malware for command and control or data exfiltration and should be blocked given they rarely have legitimate business uses.

Security Control: 1627; Revision: 0; Updated: Nov-20; Applicability: All

Inbound network connections from anonymity networks to internet-facing services are blocked.

Security Control: 1628; Revision: 0; Updated: Nov-20; Applicability: All

Outbound network connections to anonymity networks are blocked.

Further information

Further information on wireless networks can be found in the wireless networks section of these guidelines.

Further information on functional separation of servers using virtualisation can be found in the virtualisation hardening section of the [Guidelines for System Hardening](#).

Further information on implementing network segmentation and segregation for administration purposes can be found in the system administration section of the [Guidelines for System Management](#).

Further information on event logging and auditing can be found in the event logging and auditing section of the [Guidelines for System Monitoring](#).

Further information on gateways can be found in the [Guidelines for Gateways](#).

Further information on network segmentation and segregation can be found in the Australian Cyber Security Centre (ACSC)'s [Implementing Network Segmentation and Segregation](#) publication.

Further information on network plans can be found in the United States' National Security Agency's [Manageable Network Plan Guide \(version 4.0\)](#) publication.

Further information on blocking anonymity network traffic can be found in the ACSC's [Defending Against the Malicious Use of the Tor Network](#) publication.

Further information on Domain Name Systems can be found in the ACSC's [Domain Name System Security for Domain Owners](#) and [Domain Name System Security for Domain Resolvers](#) publications.

Wireless networks

Choosing wireless devices

Wireless devices that have been certified against a Wi-Fi Alliance certification program provide an organisation with the assurance that they conform to wireless standards. Deploying wireless devices that are guaranteed to be interoperable with other wireless devices will prevent problems on a wireless network.

Security Control: 1314; Revision: 2; Updated: Sep-21; Applicability: All

All wireless devices are Wi-Fi Alliance certified.

Wireless networks for public access

When an organisation provides a wireless network for the general public, connecting such a wireless network to, or sharing infrastructure with, any other network creates an additional entry point for an adversary to target connected networks to steal data or disrupt services.

Security Control: 0536; Revision: 6; Updated: Sep-18; Applicability: All

Wireless networks provided for the general public to access are segregated from all other networks.

Administrative interfaces for wireless access points

Administrative interfaces allow users to modify the configuration and security settings of wireless access points. Often wireless access points, by default, allow users to access the administrative interface over methods such as fixed network connections, wireless network connections and serial connections. Disabling the administrative interface for wireless network connections on wireless access points will assist in preventing unauthorised connections.

Security Control: 1315; Revision: 2; Updated: Sep-18; Applicability: All

The administrative interface on wireless access points is disabled for wireless network connections.

Default settings

Some wireless access points and wireless devices come with a default Service Set Identifier (SSID) and/or weak default configuration settings. As the default SSIDs of wireless access points are often documented in internet forums, along with default accounts and passphrases, it is important to change the default SSID of wireless access points along with default passphrases and weak configuration settings for all wireless devices.

When changing the default SSID, it is important that the new SSID does not bring undue attention to an organisation's wireless network. In doing so, the SSID of a wireless network should not be readily associated with an organisation, the location of their premises or the functionality of the wireless network.

A method commonly recommended to lower the profile of a wireless network is disabling SSID broadcasting. While this ensures that the existence of the wireless networks is not broadcast overtly using beacon frames, the SSID is still broadcast in probe requests, probe responses, association requests and re-association requests. As such, it is easy to determine the SSID of the wireless network by capturing these requests and responses. By disabling SSID broadcasting, organisations will make it more difficult for users to connect to a wireless network. Furthermore, an adversary could configure a malicious wireless access point to broadcast the same SSID as the hidden SSID used by a legitimate wireless network, thereby fooling users or devices into automatically connecting to the adversary's malicious wireless access point instead. In doing so, the adversary could steal authentication credentials in order to gain access to the legitimate wireless network. For these reasons, it is recommended organisations enable SSID broadcasting.

Security Control: 1316; Revision: 2; Updated: Sep-18; Applicability: All

The default SSID of wireless access points is changed.

Security Control: 1317; Revision: 2; Updated: Sep-18; Applicability: All

The SSID of a non-public wireless network is not readily associated with an organisation, the location of their premises or the functionality of the wireless network.

Security Control: 1318; Revision: 2; Updated: Sep-18; Applicability: All

SSID broadcasting is enabled on wireless networks.

Security Control: 1709; Revision: 0; Updated: Sep-21; Applicability: All

Default accounts and passphrases of wireless devices are changed.

Security Control: 1710; Revision: 0; Updated: Sep-21; Applicability: All

Configuration settings for wireless devices are hardened.

Static addressing

Assigning static IP addresses for devices accessing wireless networks can prevent a rogue device when connecting to a wireless network from being assigned a routable IP address. However, some adversaries will be able to determine IP addresses of legitimate users and use this information to guess or spoof valid IP address ranges for wireless networks. Configuring devices to use static IP addresses introduces a management overhead without any tangible security benefit.

Security Control: 1319; Revision: 2; Updated: Sep-18; Applicability: All

Static addressing is not used for assigning IP addresses on wireless networks.

Media Access Control address filtering

Devices that connect to wireless networks generally have a unique Media Access Control (MAC) address. As such, it is possible to use MAC address filtering on wireless access points to restrict which devices can connect to a wireless network. While this approach will introduce a management overhead, it can prevent rogue devices from connecting to a wireless network. However, some adversaries will be able to determine valid MAC addresses of legitimate users already on a wireless network. Adversaries can then use this information to spoof valid MAC addresses and gain access

to the wireless network. MAC address filtering introduces a management overhead without any tangible security benefit.

Security Control: 1320; Revision: 2; Updated: Sep-18; Applicability: All

MAC address filtering is not used to restrict which devices can connect to wireless networks.

Confidentiality and integrity of wireless network traffic

As wireless networks are often capable of being accessed from outside the perimeter of secured spaces, all wireless network traffic requires suitable cryptographic protection. For this purpose it is recommended that Wi-Fi Protected Access 3 (WPA3) be used as it provides equivalent or greater security than its predecessor Wi-Fi Protected Access 2 (WPA2). WPA3 has also prohibited the use of various outdated and insecure cipher suites.

WPA3-Enterprise supports three enterprise modes of operation: enterprise only mode, transition mode and 192-bit mode. Preference is given to WPA3-Enterprise 192-bit mode as this mode incorporates changes that satisfy Commercial National Security Algorithm Suite requirements and ensures no algorithms with known weaknesses are used. However, if any other WPA3-Enterprise modes are used then Authentication and Key Management suite 00-0F-AC:1 should be disabled (if this option is available).

Security Control: 1332; Revision: 3; Updated: Sep-21; Applicability: All

WPA3-Enterprise 192-bit mode is used to protect the confidentiality and integrity of all wireless network traffic.

802.1X authentication

WPA3-Enterprise uses 802.1X authentication which requires the use of an Extensible Authentication Protocol (EAP). A number of EAP methods supported by both WPA2 and WPA3 are available.

Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) is considered one of the most secure EAP methods and is widely supported. It uses a Public Key Infrastructure to secure communications between devices and a Remote Access Dial-In User Service (RADIUS) server through the use of X.509 certificates. While EAP-TLS provides strong mutual authentication, it requires an organisation to have established a Public Key Infrastructure. This involves deploying their own certificate authority and issuing certificates, or purchasing certificates from a commercial certificate authority, for every device that accesses the wireless network. While this introduces additional costs and management overheads, the security advantages are significant.

Security Control: 1321; Revision: 2; Updated: Sep-21; Applicability: All

802.1X authentication with EAP-TLS, using X.509 certificates, is used for mutual authentication; with all other EAP methods disabled on supplicants and authentication servers.

Security Control: 1711; Revision: 0; Updated: Sep-21; Applicability: All

User identity confidentiality is used if available with EAP-TLS implementations.

Evaluation of 802.1X authentication implementation

The security of 802.1X authentication is dependent on four main elements and how they interact with each other. These four elements include supplicants, authenticators, wireless access points and authentication servers. To provide assurance that these elements have been implemented correctly, they should have completed an evaluation.

Security Control: 1322; Revision: 4; Updated: Sep-21; Applicability: All

Evaluated supplicants, authenticators, wireless access points and authentication servers are used in wireless networks.

Generating and issuing certificates for authentication

When issuing a certificate to a device in order to access a wireless network, organisations should be aware that it could be stolen by malicious code. Once compromised, the certificate could be used on other devices to gain unauthorised

access to the wireless network it was issued for. Organisations should also be aware that in only issuing a certificate to a device, any actions taken by a user will only be attributable to a device and not a specific user.

When issuing a certificate to a user in order to access a wireless network, it can be in the form of a certificate that is stored on a device or a certificate that is stored within a smart card. Issuing certificates on smart cards provides increased security, but at a higher cost. Specifically, a user is more likely to notice a missing smart card and alert their security team, who are then able to revoke the credentials on the RADIUS server, which can minimise the time an adversary has access to the wireless network. In addition, to reduce the likelihood of a stolen smart card from being used to gain unauthorised access to a wireless network, multi-factor authentication can be implemented through the use of personal identification numbers on smart cards. This is particularly important when a smart card grants a user any form of administrative access.

Security Control: 1324; Revision: 3; Updated: Aug-19; Applicability: All

Certificates are generated using an evaluated certificate authority solution or hardware security module.

Security Control: 1323; Revision: 3; Updated: Sep-21; Applicability: All

Certificates are required for both devices and users accessing wireless networks.

Security Control: 1327; Revision: 2; Updated: Sep-21; Applicability: All

Certificates are protected by encryption, user authentication, and both logical and physical access controls.

Caching 802.1X authentication outcomes

When 802.1X authentication is used, a shared secret key known as the Pairwise Master Key (PMK) is generated upon successful authentication of a device. This PMK is then capable of being cached to assist with fast roaming between wireless access points. When a device roams away from a wireless access point that it has authenticated to, it will not need to perform a full re-authentication should it roam back while the cached PMK remains valid. To further assist with roaming, wireless access points can be configured to pre-authenticate a device to other neighbouring wireless access points that the device might roam to. Although requiring full authentication for a device each time it roams between wireless access points is ideal, organisations can choose to use PMK caching and pre-authentication if they have a business requirement for fast roaming. If PMK caching is used, the PMK caching period should not be set to greater than 1440 minutes (24 hours).

Security Control: 1330; Revision: 1; Updated: Sep-18; Applicability: All

The PMK caching period is not set to greater than 1440 minutes (24 hours).

Fast Basic Service Set Transition

The WPA3 standard specifies support for Fast Basic Service Set Transition (FT) (802.11r). FT is a feature designed to improve user mobility and combat lag introduced by the need to authenticate to each wireless access point. However, FT requires authenticators to request and send keys to other authenticators within a security domain. If any of these keys are intercepted, all security properties are lost. Therefore, it is imperative that communications are appropriately secured. As such, FT should be disabled unless it can be confirmed that authenticator-to-authenticator communications are secured by a suitable ASD Approved Cryptographic Protocol that provides confidentiality, integrity and mutual authentication.

Security Control: 1712; Revision: 0; Updated: Sep-21; Applicability: All

The use of FT (802.11r) is disabled unless authenticator-to-authenticator communications are secured by an ASD Approved Cryptographic Protocol.

Remote Authentication Dial-In User Service authentication

Separate to the 802.1X authentication process is the RADIUS authentication process that occurs between authenticators and a RADIUS server. RADIUS is what is known as an authentication, authorisation and accounting protocol, and is intended to mediate network access. However, RADIUS is not secure enough to be used without

protection. To protect credentials communicated between authenticators and a RADIUS server, communications should be encapsulated with an additional layer of encryption, such as RADIUS over Internet Protocol Security or RADIUS over Transport Layer Security.

Security Control: 1454; Revision: 2; Updated: Sep-21; Applicability: All

Communications between authenticators and a RADIUS server are encapsulated with an additional layer of encryption using RADIUS over Internet Protocol Security or RADIUS over Transport Layer Security.

Interference between wireless networks

Where multiple wireless networks are deployed in close proximity, there is the potential for interference to impact the availability of a wireless network, especially when operating on commonly used 802.11b/g (2.4 GHz) default channels of 1 and 11. Sufficiently separating wireless networks through the use of frequency separation can help reduce this security risk. This can be achieved by using wireless networks that are configured to operate on channels that minimise overlapping frequencies or by using both 802.11b/g (2.4 GHz) channels and 802.11n (5 GHz) channels. It is important to note though, if implementing a mix of 2.4 GHz and 5 GHz channels, not all devices may be compatible with 802.11n and able to connect to 5 GHz channels.

Security Control: 1334; Revision: 2; Updated: Sep-18; Applicability: All

Wireless networks implement sufficient frequency separation from other wireless networks.

Protecting management frames on wireless networks

An effective denial of service can be performed by exploiting unprotected management frames using inexpensive commercial hardware. The 802.11 standard provides no protection for management frames and therefore does not prevent spoofing or denial of service activities. However, the 802.11w amendment specifically addresses the protection of management frames on wireless networks and should be enabled for WPA2, in WPA3 this feature is built into the standard.

Security Control: 1335; Revision: 1; Updated: Sep-18; Applicability: All

Wireless access points enable the use of the 802.11w amendment to protect management frames.

Wireless network footprint

Instead of deploying a small number of wireless access points that broadcast on high power, a greater number of wireless access points that use less broadcast power can be deployed to achieve the desired footprint. This has the benefit of providing service continuity should a wireless access point become unserviceable. In such a case, the output power of nearby wireless access points can be increased to cover the footprint gap until the unserviceable wireless access point can be replaced.

In addition to minimising the output power of wireless access points to reduce the footprint of a wireless network, the use of Radio Frequency (RF) shielding can be used for an organisation's facilities. While expensive, this will limit the wireless communications to areas under the control of an organisation. RF shielding on an organisation's facilities has the added benefit of preventing the jamming of wireless networks from outside of the facilities in which wireless networks are operating.

Security Control: 1338; Revision: 1; Updated: Sep-18; Applicability: All

Instead of deploying a small number of wireless access points that broadcast on high power, a greater number of wireless access points that use less broadcast power are deployed to achieve the desired footprint.

Security Control: 1013; Revision: 6; Updated: Dec-21; Applicability: S, TS

The effective range of wireless communications outside an organisation's area of control is limited by implementing RF shielding on facilities in which SECRET or TOP SECRET wireless networks are used.

Further information

Further information on implementing segregation using VLANs can be found in the network design and configuration section of these guidelines.

Further information on selecting evaluated products can be found in the evaluated product acquisition section of the [Guidelines for Evaluated Products](#).

Further information on encryption for wireless networks can be found in the [Guidelines for Cryptography](#).

Further information on [Wi-Fi technologies](#) and associated [certification programs](#) are available from the Wi-Fi Alliance.

Service continuity for online services

Cloud-based hosting of online services

Using a cloud service provider can allow an organisation to build highly resilient online services due to the increased computing resources, bandwidth and multiple separate physical sites made available by the cloud provider. Organisations can achieve the same results using their own infrastructure; however, this may require significant upfront costs and may still result in a limited capability to scale dynamically to meet increased demand. In case of a denial-of-service attack, cloud-based hosting can also provide segregation from self-hosted or other cloud hosted services ensuring that other systems, such as email services, are not affected.

Security Control: 1437; Revision: 4; Updated: Dec-21; Applicability: All

A cloud service provider is used for hosting online services.

Location policies for online services

When using cloud service providers, organisations will need to consider whether they should lock their data to specific regions or availability zones. In doing so, organisations that specify locking policies will have an expectation that their data won't be relocated to different regions or availability zones by the cloud service provider.

Security Control: 1578; Revision: 1; Updated: Dec-21; Applicability: All

Organisations are notified by cloud service providers of any change to configured regions or availability zones for online services.

Availability planning and monitoring for online services

It is important that the connectivity between organisations and their cloud service providers meets organisational requirements for bandwidth, latency and reliability. To support this, organisations and cloud service providers should discuss and document any specific network requirements, performance characteristics or planned responses to availability failures, especially when requirements for high availability exist. This includes whether network connections between organisations and cloud service providers will use dedicated communication links, or connect over the internet, and whether any secondary communications links will provide sufficient capacity to maintain operational requirements should the primary communication link become unavailable.

Furthermore, capacity monitoring should be performed in order to manage workloads and monitor the health of online services. This can be achieved through continuous and real-time monitoring of metrics such as latency, jitter, packet loss, throughput and availability. In addition, feedback should be provided to cloud service providers when performance does not meet service level agreement targets. To assist with this, anomaly detection can be performed through network telemetry that is integrated into security monitoring tools.

Security Control: 1579; Revision: 1; Updated: Dec-21; Applicability: All

Cloud service providers' ability to dynamically scale resources due to a genuine spike in demand or a denial-of-service attack is tested as part of capacity planning processes for online services.

Security Control: 1580; Revision: 1; Updated: Dec-21; Applicability: All

Where a high availability requirement exists for online services, the services are architected to automatically transition between availability zones.

Security Control: 1441; Revision: 3; Updated: Dec-21; Applicability: All

Where a requirement for high availability exists for online services, a denial of service mitigation service is used.

Security Control: 1581; Revision: 1; Updated: Dec-21; Applicability: All

Organisations perform continuous real-time monitoring of the availability of online services.

Using content delivery networks

Similar to cloud-based hosting, the use of content delivery networks (CDNs) and denial of service mitigation services can allow an organisation to create highly resilient online services by leveraging the large bandwidth, geographically dispersed hosting locations, traffic scrubbing and other security controls offered by CDN and denial of service mitigation service providers.

The use of CDNs is particularly effective when serving static, bandwidth intensive media such as images, sound or video files. However, the services offered by a CDN can include more than basic content hosting such as web response caching, load balancing, web application security controls or denial of service mitigations.

Care should be taken when configuring the use of a CDN or denial of service mitigation service to ensure that the IP address of the organisation's web server is not identifiable by an adversary as this could allow for protections to be bypassed. Additionally, appropriate network security controls should be applied to only allow communication between an organisation's server, the CDN or denial of service mitigation service provider and the authorised management environment.

Security Control: 1438; Revision: 2; Updated: Dec-21; Applicability: All

Where a high availability requirement exists for website hosting, CDNs that cache websites are used.

Security Control: 1439; Revision: 2; Updated: Dec-21; Applicability: All

If using a CDN, disclosing the IP address of the web server under the organisation's control (referred to as the origin server) is avoided and access to the origin server is restricted to the CDN and an authorised management network.

Denial of service strategies

Denial-of-service attacks are designed to disrupt or degrade online services such as website, email and Domain Name System services. To achieve this goal, adversaries may use a number of approaches to deny access to legitimate users of online services:

- using multiple computers to direct a large volume of unwanted network traffic at online services in an attempt to consume all available network bandwidth
- using multiple computers to direct tailored traffic at online services in an attempt to consume the processing resources of online services
- hijacking online services in an attempt to redirect legitimate users away from those services to other services that the adversary controls.

Although an organisation cannot avoid being targeted by denial-of-service attacks, there are a number of measures they can implement to prepare for and potentially reduce the impact if targeted. This includes engaging with their cloud service providers to identify the denial of service detection technologies that may be available for use. For example, real-time capacity reporting dashboards, that provide out-of-band and real-time alerts based on organisation-defined thresholds, can assist with the rapid identification of denial-of-service attacks. In addition, not all online services or functionality offered by an organisation may be business critical. Understanding what services can be offered with reduced functionality, deprioritised, disabled or lived without can help an organisation reduce or eliminate the impact on other more essential services or free up resources to respond to more critical services first.

Overall, preparing for denial-of-service attacks before they occur is by far the best strategy as it is very difficult to respond once they begin and efforts at this stage are unlikely to be effective.

Security Control: 1431; Revision: 3; Updated: Dec-21; Applicability: All

Denial-of-service attack prevention and mitigation strategies are discussed with cloud service providers, specifically:

- *their capacity to withstand denial-of-service attacks*
- *any costs likely to be incurred as a result of denial-of-service attacks*
- *thresholds for notification of denial-of-service attacks*
- *thresholds for turning off online services during denial-of-service attacks*
- *pre-approved actions that can be undertaken during denial-of-service attacks*
- *denial-of-service attack prevention arrangements with upstream service providers to block malicious traffic as far upstream as possible.*

Security Control: 1458; Revision: 2; Updated: Dec-21; Applicability: All

The functionality and quality of online services, how to maintain such functionality, and what functionality can be lived without during a denial-of-service attack, are determined and documented.

Domain name registrar locking

The use of domain name registrar locking can prevent a denial of service caused by unauthorised deletion or transferal of a domain, or other unauthorised modification of a domain's registration details.

Security Control: 1432; Revision: 2; Updated: Dec-21; Applicability: All

Domain names for online services are protected via registrar locking and confirming domain registration details are correct.

Monitoring with real-time alerting for online services

Organisations should perform automated monitoring of online services with real-time alerting to ensure that a denial-of-service attack is detected and responded to as soon as possible.

Security Control: 1435; Revision: 2; Updated: Dec-21; Applicability: All

Availability monitoring with real-time alerting is implemented for online services to detect denial-of-service attacks and measure their impact.

Segregation of critical online services

Denial-of-service attacks are typically focused on highly visible online services, such as an organisation's core website, in order to have a publicly noticeable impact. By segregating online services (e.g. having one internet connection for email and internet access and a separate connection for web hosting services) the impact of a denial-of-service attack can be limited to just a targeted service.

Security Control: 1436; Revision: 2; Updated: Dec-21; Applicability: All

Critical online services are segregated from other online services that are more likely to be targeted.

Preparing for service continuity

Depending on the nature of a denial-of-service attack, replacing a full-featured website with a minimal impact static version can help provide a level of service which would otherwise not be possible.

An organisation's standard full-featured website may have higher processing or resource demands due to database integration or the presence of large media files such as high-resolution images or videos. These additional resource requirements may make the website more susceptible to denial-of-service attacks.

Security Control: 1518; Revision: 1; Updated: Dec-21; Applicability: All

A static version of a website is pre-prepared that requires minimal processing and bandwidth in order to facilitate at least a basic level of service when under a denial-of-service attack.

Further information

Further information on mitigating denial-of-service attacks can be found in the ACSC's [Preparing for and Responding to Denial-of-Service Attacks](#) publication.