



Implementing Network Segmentation and Segregation

APRIL 2019

Introduction

This document intends to assist staff responsible for an organisation's network architecture and design to increase the security posture of their networks by applying network segmentation and segregation strategies.

Network segmentation and segregation are highly effective strategies an organisation can implement to limit the impact of a network intrusion. If implemented correctly, these strategies can make it significantly more difficult for an adversary to locate and gain access to an organisation's most sensitive information; and increase the likelihood of detecting an adversary's activity in a timely manner.

Historically, network segmentation and segregation has been implemented with a perimeter firewall at an organisation's internet (or inter-organisation) gateway; or a pair of firewalls surrounding a demilitarised zone which provides a segregated environment between completely trusted and untrusted zones. However, simply zoning an entire network as 'trusted' and treating it as 'flat' (i.e. the entire network as a single segment) creates an environment that requires only a single network intrusion for an adversary to gain widespread access. A flat network also allows an adversary to pivot between hosts and services with minimal obstruction and limited chance of detection.

With adversary tradecraft targeting internal networks directly using techniques such as spear-phishing and social engineering, along with the increasing use of mobile and remote working, these common flat network architectures do not protect organisations from contemporary cyber threats. As a result, it is important for organisations to segment networks and segregate sensitive information, hosts and services from the environment in which users access external resources; in particular the Web, email and other internet services.

What is network segmentation and segregation?

Network segmentation involves partitioning a network into smaller networks; while network segregation involves developing and enforcing a ruleset for controlling the communications between specific hosts and services.

When implementing network segmentation and segregation, the aim is to restrict the level of access to sensitive information, hosts and services while ensuring an organisation can continue to operate effectively. To be effective, network segmentation and segregation measures must be carefully planned, robustly enforced, closely monitored and be unable to be bypassed.

Implementing network segmentation and segregation can be achieved using a number of techniques and technologies, including:

- Implementing demilitarised zones and gateways between networks with different security requirements (security domains) utilising technologies at various layers such as:
 - routers or layer 3 switches to divide a large network into separate smaller networks to restrict traffic flow using measures such as access control lists

- virtualised networking and routing protocols, including Virtual Local Area Networks and Virtual Routing and Forwarding to segment the network
 - virtual machines, containers and virtual functions to isolate activities of different trust or threat levels (such as accessing the Internet or email, or performing privileged administrative tasks)
 - virtual hosts, virtual switching, cloud tenancies and managed security groups to segregate and segment applications, data and other services
 - host-based security and firewall software to filter network traffic at the host level
 - network firewalls and security appliances between networks to filter network traffic
 - network access controls to control the devices which can access networks
 - application and service firewalls and proxies (or service brokers) to permit only authorised communications between applications and services in different networks
 - user and service authentication and authorisation, including multi-factor authentication and policy-based access controls to enforce least privilege
 - data diodes and one-way transfer devices to enforce the directionality of data flows between networks
 - content filtering techniques including recursive decomposition, validation, verification and sanitisation to comprehensively assure network and application traffic flows.
- Implementing server and domain isolation using Internet Protocol Security (IPsec).
 - Implementing storage based segmentation and filtering using technologies such as disk and volume encryption and Logical Unit Number masking.
 - For extremely sensitive network connections, implementing Cross Domain Solutions or other technologies recommended by the Australian Signals Directorate’s Australian Cyber Security Centre (ACSC).

To be successful, implementation of these techniques and technologies must be driven by a network architecture based on achieving organisational business and security requirements. It is vital that network, system and security architects work together with business analysts and customers to ensure that an accurate and considered strategy is adopted.

Why is network segmentation and segregation important?

Once an adversary compromises a network, usually through the compromise of a host under the control of a legitimate user by means of social engineering, they will attempt to move around the network to locate and access sensitive information, hosts and services. In order to minimise the impact of such a network intrusion, it should be as hard as possible for the adversary to find and access such information, move undetected around the network and remove information from the network.

An adversary may attempt to make connections directly from a compromised host to a more sensitive host using tools and techniques at their disposal. For example, if an adversary has initially compromised a workstation, they may seek to create a remote connection to a server, map a network resource or use legitimate network administration tools in order to access sensitive information or execute malicious code on that server. This is particularly common when an adversary targets an organisation’s authentication server. Properly planned and implemented network segmentation and segregation is a key security measure to assist in preventing such activities from occurring. Example mitigations include explicitly disallowing remote desktop connections or the use of common network administration tools from user workstations (as most users do not require such functionality), configuring servers to limit the sharing of files, and restricting servers’ ability to communicate via remote connections.

Network segmentation and segregation can also assist security personnel in their duties. Organisations implementing these measures should consider the audit and alerting capabilities of candidate technologies, as these features may prove critical in identifying a network intrusion and ensuring timely incident response activities. A mature segmented and segregated environment will also allow an organisation to better focus their auditing and alerting strategies on a prioritised subset of attack vectors informed by the approved network access methods. Additionally, it can provide a way to isolate compromised hosts or networks in a timely manner following a network intrusion.

How can network segmentation and segregation be implemented?

Regardless of the technologies chosen for network segmentation and segregation, there are five common themes for best practice implementations:

- Apply technologies at more than just the network layer. Each host and network should be segmented and segregated, where possible, at the lowest level that can be practically managed. In most cases, this applies from the data link layer up to and including the application layer; however, in particularly sensitive environments, physical isolation may be appropriate. Host-based and network-wide measures should be deployed in a complementary manner and be centrally monitored. It is not sufficient to simply implement a firewall or security appliance as the only security measure.
- Use the principles of least privilege and need-to-know. If a host, service or network doesn't need to communicate with another host, service or network, it should not be allowed to. If a host, service or network only needs to talk to another host, service or network on a specific port or protocol, and nothing else, it should be restricted to this. Adopting these principles across a network will complement the minimisation of user privileges and significantly increase the overall security posture of the environment.
- Separate hosts and networks based on their sensitivity or criticality to business operations. This may include using different hardware or platforms depending on different security classifications, security domains or availability/integrity requirements for certain hosts or networks. In particular, separate management networks and consider physically isolating out-of-band management networks for sensitive environments.
- Identify, authenticate and authorise access by all entities to all other entities. All users, hosts and services should have their access to all other users, hosts and services restricted to only those required to perform their designated duties or functions. All legacy or local services which bypass or downgrade the strength of identification, authentication and authorisation services should be disabled wherever possible and have their use closely monitored.
- Implement whitelisting of network traffic instead of blacklisting. Only permit access for known good network traffic (i.e. that which is identified, authenticated and authorised), rather than denying access to known bad network traffic (e.g. blocking a specific address or service). Not only will whitelisting result in a superior security policy to blacklisting, it will also significantly improve an organisation's capacity to detect and assess potential network intrusions.

The following types of traffic flow filtering techniques should be considered when implementing network segmentation and segregation. As stated above, these filtering techniques will be significantly more effective if implemented using a whitelisting approach:

- Logical access restrictions of network traffic such as:
 - network layer filtering that restricts which hosts are able to communicate with other hosts based on Internet Protocol and route information

- state-based filtering that restricts which hosts are able to communicate with other hosts based on their intended function or current state of operation
- port and/or protocol level filtering that restricts the number and type of services that each host can use to communicate with other hosts.
- Authentication filtering to restrict access to hosts, services and networks based on strong authentication, commonly implemented using public key cryptography, such as certificate-based IPsec.
- Application filtering to filter the content of communications between hosts and networks at the application layer, commonly implemented using email and web content filtering, intrusion prevention systems, and web application or XML firewalls.
- For particularly sensitive environments, it may also be appropriate to implement physical isolation between networks. Where limited interaction between physically-isolated networks is necessary, one or more of the following may be required:
 - bespoke or tailored security device
 - High Assurance product, or
 - Cross Domain Solution.

As with any security strategy, network segmentation and segregation implementations must be adapted when significant changes occur to an organisation’s network. Additionally, environmental changes such as new business functions and evolving cyber threats will necessitate reviews of an organisation’s security posture and network architecture to ensure appropriate mitigation strategies are applied.

Example implementations of network segmentation and segregation

Segmenting a network to protect key hosts

In this scenario an organisation had decided to segment their network to protect key hosts from a network intrusion. In doing so they implemented the following security measures:

- compiled an inventory of key hosts documenting their sensitivity and any necessary communications with such hosts
- planned the introduction of security measures in a schedule that was achievable with the resources allocated ensuring sufficient testing prior to deployment
- restricted logical network connectivity to key hosts to only those ports and protocols that were essential
- only allowed connections to be established from more trusted to less trusted zones and not vice versa (with the exception of necessary user access to application interfaces)
- whitelisted application layer content so that only required content was able to flow between different trust zones
- implemented multi-factor authentication in addition to using a separate set of credentials for users and services if their function was more sensitive than other users or services sharing the same host or network
- minimised the use of implicit trust relationships between hosts in the same and different trust zones (the trust relationships defined across different trust zones were implemented such that each side of the trust relationship authenticated and authorised the other)

- implemented Web, email and file content filtering for connections to external organisations and the Internet to detect and sanitise potentially malicious content
- applied intrusion prevention and host-based antivirus to detect and quarantine identified malicious content
- implemented centralised logging, alerting, monitoring and auditing capabilities which were the responsibility of a dedicated security operations team.

The above list is not an exhaustive set of security measures; however, it is a realistic overview which demonstrates that network segmentation and segregation must be considered at all layers to be effective. Implementing a secure network architecture is never as simple as implementing a gateway firewall with restrictive access control lists.

Segregating high-risk applications from a network

In this scenario an organisation had identified that most of their network contained sensitive information and segmenting the network or segregating all of that information was not cost-effective. Instead, the organisation chose to segregate high-risk applications (i.e. web browsers, email clients and content management systems) from the rest of the network. In doing so, they implemented the following security measures to maintain business requirements while reducing the risk of a successful network intrusion:

- Users requiring internet access launched a remote desktop application on their corporate workstation to access a virtual desktop and authenticated with a user account used only for that purpose. This virtual desktop was served from a dedicated server hosted in a different network segment within a different authentication domain. This dedicated remote desktop allowed users to conduct high-risk activities such as web browsing and reading emails while limiting the utility of a single compromised application to an adversary.
- Users requiring access to high-risk applications launched a local virtualisation application to run a hardened virtual host which connected to a less-trusted remote environment which was protected by a layered security gateway that broke apart and abstracted all necessary communications protocols between high-risk applications and the organisation's corporate network.

Summary of example implementations

The key takeaway from both approaches was that users did not store or process potentially malicious data directly on their corporate workstation or use the corporate servers which were relied upon for sensitive and business-critical functions. Each user's interaction was with a remote desktop or application and, if required, output was sent back to the user through a sufficiently structured and limited capability that prevented malicious code from executing or propagating throughout the corporate network.

It is important to remember that when implementing security measures an organisation will incur a resource cost to ensure that the additional systems are appropriately maintained. As with other technology assets, these security measures should be managed and monitored, with security patches applied as soon as possible after release.

Finally, it is recommended that all web browsing environments should be non-persistent, rigorously hardened and subject to regular technical security assessments. Therefore, if the web browsing environment does become compromised with malicious code, the infection is quickly removed when the user completes their web browsing session.

Further information

The **Australian Government Information Security Manual (ISM)** assists in the protection of information that is processed, stored or communicated by organisations' systems. It can be found at <https://www.cyber.gov.au/ism>.

The ***Strategies to Mitigate Cyber Security Incidents*** complements the advice in the ISM. The complete list of strategies can be found at <https://www.cyber.gov.au/publications/strategies-to-mitigate-cyber-security-incidents>.

Contact details

Organisations or individuals with questions regarding this advice can email asd.assist@defence.gov.au or call 1300 CYBER1 (1300 292 371).