Network segmentation and segregation

1. Network segmentation and segregation is, along with multi-factor authentication, one of the most effective controls an agency can implement to mitigate the second stage of a network intrusion, propagation or lateral movement. If implemented correctly, it can make it significantly more difficult for a malicious cyber adversary to locate and gain access to your agency's most sensitive information.

2. When implementing network segmentation and segregation, the aim is to minimise the methods and level of access to sensitive information for those systems and people who don't need it, whilst ensuring your agency can continue to operate effectively. This can be achieved using a number of techniques and technologies depending on your network's architecture and configuration.

3. Traditionally, network segmentation and segregation has been implemented at the Internet gateway. With the malicious cyber adversary's intrusion methods evolving to target your internal network directly using socially engineering, and the increasing use of mobility and remote working, it is becoming increasingly important to segment and segregate your sensitive information from the environment within which your users access the web, email and other external services.

What is network segmentation and segregation?

4. Network segmentation involves partitioning the network into smaller networks. Network segregation involves developing and enforcing a ruleset controlling which computing devices are permitted to communicate with which other computing devices.

5. When implementing network segmentation and segregation correctly you are minimising the method and level of access to sensitive information, whilst not stopping your agency from operating effectively.

6. This can be achieved using a variety of technologies and methods. Depending on the architecture and configuration of your network, some of the common technologies and methods used include:

   a. Implementing demilitarised zones (DMZ) and gateways between systems or networks with different security requirements (security domains) utilising technologies at various layers such as:

      i. Separate physical links and systems;
      ii. Traffic flow filters;
      iii. Virtual Local Area Networks (VLANs);
      iv. Network and host-based Firewalls;

v. Network Access Control;
vi. Data Diodes;
vii. Application Firewalls;
viii. Application and service proxies;
ix. User and service authentication and authorisation; and
x. Content based filtering.

b. Implementing server and domain isolation using Internet Protocol Security (IPsec)[1].

c. Implementing storage based segmentation and filtering using technologies such as:

i. Encryption; and
ii. Logical Unit Number (LUN) masking.

d. Implementing DSD-evaluated cross-domain solutions (CDS) where necessary[2].

7. Network segmentation and segregation is not just important for protecting systems or networks of different security classifications, but also for protecting systems of the same classification but with different security requirements. A classic example of this is providing segmentation between the systems where your users browse the web and access email and your most sensitive information. You may be prepared to let someone else's unauthenticated code from the web run on your user's workstation (i.e. JavaScript, Java Applets and Flash content), but you should not allow the same unauthenticated code to execute on, or access, your database server containing your organisation's most sensitive information.

Why is network segmentation and segregation important?

8. Once a malicious cyber adversary compromises your network, usually through the compromise of a system under the control of a legitimate user by means of social engineering, they will attempt to move around your network to locate and access the information they are targeting. This is known as propagation or lateral movement.

9. In order to minimise the impact of such a compromise, it should be as hard as possible for the malicious cyber adversary to find and access the information they seek and move undetected around a system or network, and remove the information from the network once they locate it.

10. The malicious cyber adversary may attempt to make connections directly from the compromised system(s) to the more sensitive system(s) using tools and techniques they have at their disposal. This is typically executed using their own or enhanced versions of legitimate network administration tools. For example, if the malicious cyber adversary has initially compromised a workstation, they may seek

[1] http://technet.microsoft.com/en-us/library/cc756066(v=WS.10).aspx
[2] Additional information on implementing a cross domain solution can be accessed from the OnSecure website at https://members.onsecure.gov.au/ in the *Guide to Secure Configuration of Cross Domain Solutions* publication.

to create a remote connection to a sensitive server, map a network resource or use installed legitimate network administration tools in order to access information on that server, or execute software remotely on the server. This is particularly common when the adversary targets an organisation's authentication server. Properly planned and implemented network segmentation and segregation is a key control to stop such activities from occurring. You may be able to explicitly disallow remote desktop connections or the use of common network administration tools from end-user workstations on sensitive servers (as most users do not require such functionality) and/or configure sensitive servers to prohibit the sharing of files and restrict their ability to communicate via remote connections.

11. Network segmentation and segregation also assists an organisation to both detect and respond to an intrusion. The technologies implemented to enforce segmentation and separation will:

a. contain audit and alerting capabilities that may prove critical in identifying an intrusion;

b. allow an organisation to better focus their auditing and alerting tools to a limited subset of attacks based on the approved access methods; and

c. provide a ready way to isolate a compromised device from the rest of your network in the event of an intrusion.

12. Network segmentation and segregation is a key enabler for the implementation of workforce mobility and a secure bring your own device (BYOD) strategy as it allows you to better isolate a compromised or potentially compromised device from the key information on your network.

Best practice in implementing network segmentation and segregation

13. Regardless of the technology chosen to implement network segmentation and segregation, there are five common themes for good network segmentation and segregation, including:

a. Apply technologies at more than just the network layer. Each system and network should be segmented and segregated, where possible, from the data link layer up to and including the application layer. It is not sufficient to implement a hardware-based firewall as the only protective security measure.

b. Use the principles of least privilege and need-to-know. If a system doesn't need to communicate with another system on the network, it should not be allowed to. If a system only needs to talk to another system on a specific port or protocol and nothing else, it should be restricted as such.

c. Separate information and infrastructure based on your security requirements. This may include using different hardware or platforms based on security classifications or different threat and risk environments in which each system or network segment operates.

d. Identify, authenticate and authorise access for entities based on your security requirements. This includes users, systems and services that should have their access restricted to that required to perform their intended function.

e. Implement whitelisting instead of blacklisting. That is, grant access to the known good, rather then denying access to the known bad. This will also improve an organisation's capacity to analyse log files.

14. The following types of filtering should be considered when implementing segmentation and segregation. Further, these filtering techniques should be implemented using a whitelisting approach:

a. Logical access restrictions of network traffic, including:

i. Network layer filtering that restricts which systems are able to communicate with others on the network based on IP and route information.
ii. State-based filtering that restricts which systems are able to communicate with others on the network based on their intended function or current state of operation.
iii. Port and/or protocol level filtering that restricts the number and type of services that each system can use to communicate with others on the network.

b. Authentication filtering to restrict access to systems and services based on strong authentication, commonly implemented using public key cryptography, such as certificate based IPsec.

c. Application filtering that commonly filters the content of communications between systems at the application layer. Common implementations include email and web content filtering, intrusion prevention systems and web application or eXtensible Markup Language (XML) firewalls.

Network segmentation and segregation – two common approaches

Segmenting your network so key systems are protected from the corporate network

15. To effectively implement this type of segmentation and segregation you need to consider:

a. Minimising logical network connectivity to sensitive servers and/or applications to only those hosts and ports/protocols that are essential (whitelisting).

b. Where possible, only allow connections to be established from more trusted to less trusted zones and not vice versa (it is accepted that for user access to application interfaces this can be very difficult).

c. Where possible whitelist the application layer content so that only required content is able to flow between the zones.

d. Using a separate set of credentials for users or services if their function is more sensitive than other users or services sharing use of a system or network. Where possible, use multi-factor authentication for the most sensitive users and services on a system or network.

e. Minimising the use of implicit trust relationships between systems in the same and different zones. Trust relationships defined across different zones should be implemented so that each

side of the trust authenticates the other. Further, data exchanges across different zones should be explicitly documented and managed according to your security requirements.

f. Implementing content filtering, anti-virus and intrusion prevention to block the known bad (blacklisting).

g. Implementing logging, alerting, monitoring and auditing capabilities. Network segmentation and segregation technologies often provide a great opportunity for additional logging and alerting, however these features need to be enabled and actively managed.

16. This list is not exhaustive, however the key thing to note is that segmentation and segregation must be considered at all layers, it is not as simple as implementing a firewall with restrictive access control lists.

Segregating high risk services from the corporate network

17. In this approach the organisation has identified that much of their internal network contains sensitive information and segmenting the network or segregating all of that information is not cost effective. Instead, the organisation chooses to isolate or segregate their highest threat applications, i.e. web browsing and email from the Internet from the rest of their corporate network.

18. This approach is generally implemented by:

a. Remotely accessing a separate user environment (such as another desktop stored in a less trusted or more isolated area of the network) from which users can access the web. This is commonly implemented using Citrix or similar remote desktop application software to directly access separate user environments; or

b. Remotely accessing an application directly from the internal network. The remote application runs in the less trusted environment, however the user accesses it from their normal corporate desktop.

19. The key point with both options is that users do not store or process potentially malicious information directly from their corporate desktop. Each user provides input to the remote application or desktop and, if required, output is sent back to the user through a sufficiently structured and limited capability that prevents malware or potentially malicious content from executing or propagating throughout the corporate network.

20. One of the important controls when implementing this type of segmentation and segregation is to ensure that untrusted web browsing environments are non-persistent and regularly patched. That is, if the web browsing environment becomes compromised with malware, the infection is quickly removed when the user completes their web browsing session. Common examples of systems that enable this type of functionality include :

a. For separate desktop – your favourite virtualisation product.

b. For application virtualisation - Citrix XenApp, VMware ThinApp and Microsoft App-V.

Conclusion

21. Network segmentation and segregation, along with multi-factor authentication, is one of the most effective controls an agency can implement to mitigate the second stage of a network intrusion, known as propagation or lateral movement. When implemented correctly, it can make the challenge of locating and gaining access to your agency's most sensitive information, and removing that information significantly more difficult for a malicious cyber adversary. It also provides an excellent opportunity to isolate compromised systems when they are discovered (therefore assisting incident response) and implement more effective logging, alerting and auditing for systems on your internal network.

22. The key considerations for implementing good network segmentation and segregation are:

a. Understand your network, including how users interact with systems and how systems communicate with each other.

b. Plan your implementation carefully.

c. Use least privilege and need-to-know as guiding principles.

d. Whitelisting is always better than blacklisting.

e. Logging, alerting, monitoring and auditing are essential for managing your implementation and ensuring its effectiveness.

f. It is not just about logical network level segregation, it must be considered at all levels up to, and including the application layer.

23. DSD strongly recommends the implementation of network segmentation and segregation for separating sensitive information and systems from high risk environments used to access external services, such as the Internet. In addition, improving the security of systems and their associated networks through segmentation and segregation reduces the overall risk that an organisation must address. Implementing the advice presented by this publication provides a sound foundation for the steps required in successfully integrating a CDS.

Contact details

Australian government customers with questions regarding this advice should contact the DSD Advice and Assistance Line on 1300 CYBER1 (1300 292 371) or dsd.assist@defence.gov.au.

Australian businesses or other private sector organisations seeking further information should contact CERT Australia at info@cert.gov.au or by calling 1300 172 499.