



Australian Government
Australian Signals Directorate

Fundamentals of Cross Domain Solutions

First published: December 2019
Last updated: October 2021

Table of contents

Executive summary	1
Background	1
About this guidance	1
Introduction	2
Summary	2
Why a CDS is special	2
Common use cases for a CDS	2
Common technologies in a CDS	3
Cross domain security	3
Common misconceptions about a CDS	7
Avoiding the need for a CDS	8
Security domains explained	9
Security domains	9
Security domain enclave	9
High side and low side	9
Domain boundary	10
Air gap	10
Security zone, trust zone and demilitarised zone	11
Connections between security domains	11
Domain security models	12
CDS explained	13
CDS overview	13
CDS functions	13
Comparison with network gateways	14
Types of CDS	15
Cross domain security principles	19

Context-appropriate security-enforcing mechanisms	19
Secure architecture and design	21
System assurance and secure operation	22
Security assurance	26
Controls and mitigation strategies	26
Risk acceptance	26
The next step	27
Further information	28
Contact details	29

Executive summary

Background

Complex interconnected systems are now ubiquitous in the modern global information environment. Meanwhile, the need to share information across traditionally isolated boundaries continues to increase as numerous threats to Australia's economic and national security continue to evolve and become more sophisticated.

Organisations seeking to update or expand information sharing capabilities must do so without introducing vulnerabilities to their most sensitive systems. Meanwhile, established and emerging threats will continue to challenge the strategic, business and technology transformations necessary to achieve these capabilities. Implementing secure connectivity between systems of different sensitivity, trust or security classifications is a cornerstone of successful transformation.

Cross Domain Solution (CDS) technologies address this requirement by enabling organisations to share information across physically, logically and administratively separated networks (known as security domains) in a reliable, secure and interoperable manner.

About this guidance

This guidance introduces technical and non-technical audiences to cross domain security principles for securely connecting security domains. It explains the purpose of a CDS and promotes a data-centric approach to a CDS system implementation based on architectural principles and risk management. This guidance also covers a broad range of fundamental concepts relating to a CDS, which should be accessible to readers who have some familiarity with the field of cyber security. Organisations with complex information sharing requirements are encouraged to refer to this guidance in the planning, analysis, design and implementation of CDS systems.

This guidance is also intended to support cyber security guidance contained within the [Information Security Manual](#) (ISM).

Introduction

Summary

A CDS is a capability that can be used to securely connect discrete systems or networks. These separate systems or networks (known as security domains) may have different security policies to address their exposure to different types of threats and levels of risk, and therefore hold differing levels of trust.

The ISM defines a CDS as ‘a system capable of implementing comprehensive data flow security policies with a high level of trust between two or more differing security domains’. Put simply, a CDS is an assured system which perform the security functions necessary to control the flow of information and data between security domains. In this way, a CDS enables controlled connectivity (or hardens existing connectivity) to otherwise isolated networks, such as across an air gap. Whilst a CDS may occasionally be described as a X-Domain Solution or XDS, this is a stylistic decision and CDS requirements still apply.

Enabling the secure interconnectivity of systems will help organisations to:

- become more efficient and agile by presenting information and data to users where and when it is needed, including via automated processes
- increase accountability and compliance with governance and security requirements
- address credible threats from motivated and capable threat actors, such as foreign intelligence services.

Whilst a CDS is the key element to obtaining effective cross domain security, a thorough approach should consider more than any single CDS product in isolation. Additionally, not all cross domain issues can be solved with technical solutions alone. Human factors, information management practices and other business processes must be comprehensively understood to ensure the appropriate solution is implemented. A secure capability is only achievable if these aspects are considered throughout the capability’s lifecycle.

Why a CDS is special

CDS systems are an example of technological convergence where discrete security concepts and technologies are carefully combined within a comprehensive system architecture to protect sensitive, critical and classified networks.


Unlike a commercial off-the-shelf (COTS) network gateway, such as those used in lower risk applications, a CDS is an advanced gateway systems that employs a wide range of overlapping controls to provide assured data flows and defence-in-depth. A CDS used to protect classified systems is specifically designed to address a threat model encompassing national security threats, rather than just consumer-type threats. However, implementing controls alone will not be sufficient to make a CDS safe to operate or to perform their core functions adequately.

The Australian Signals Directorate (ASD) advocates that organisations perform rigorous analysis of the security, financial and sustainment risks when considering the business benefits of any proposed CDS in their environment. In order to gain assurance in the effectiveness of security functions performed by a CDS, systematic risk-based analysis and thorough technical assessments are strongly recommended.

Common use cases for a CDS

The security functions of a CDS are relied upon to protect isolated networks (typically classified systems or sensitive critical infrastructure). They must also protect information and data objects that pass between these networks while delivering on business requirements. A CDS enforces security policies through controls developed to address a given use case and threat environment. Security policies for a CDS are generally traceable to information security objectives such as preventing the introduction of malware and limiting the loss of an organisation’s important data.

A CDS provides a controlled interface between networks of differing trust, as illustrated by the following diagram.



Indicative use cases for a CDS include:

- transferring fixed-format messages between government and military systems
- transferring business files into or out of isolated systems
- ingesting bulk operational data from a sensitive network into a classified network
- aggregating input from multiple differently-classified environments in a central audit or monitoring system
- accessing multiple workspaces, with different sensitivities, from a single client device on the most trusted network.

Common technologies in a CDS

A CDS may be an integrated appliance or, more commonly, be composed of a combination of general purpose security components and specialist cross domain security-enforcing mechanisms. Common technologies in CDS systems include:

- CDS appliances or other cross domain guards
- data diodes for one-way data flow control and network protocol termination
- content filtering software and appliances
- filter orchestration and workflow management software
- anti-malware software and detonation chambers
- data loss prevention and insider threat detection software
- identity and access management infrastructure
- logging and audit systems
- application layer firewalls.

It should be noted that the CDS technology landscape continues to mature and technology advancements can be expected to yield significant improvements to both security and business functionality.

Cross domain security

Cross domain security involves more than just focusing on an appliance or implementing a data diode with no additional security enforcement. It is a comprehensive approach to defending against both known and unknown threats to data connections at the boundaries of sensitive or classified networks and enclaves. In practice, cross domain security is achieved through a combination of robust cyber security protections tailored for high risk use cases.

High level cross domain security considerations include:

- determining the most critical security properties and security functionality for a CDS and connected systems, prior to commencing design or acquisition
- understanding the risks inherent to the data that will be transferred, and ensuring security policy enforcement meets these risks
- recognising that security functionality may exist in a single appliance or be distributed across multiple components in a CDS or adjacent systems
- ensuring design and implementation assurance requirements, including physical and personnel controls, are proportionate to the criticality of the system
- tailoring a CDS to address the unique security and business environment
- blocking all data flows by default and only allowing known good data to pass based on predefined rules
- considering each layer of the Open System Interconnect (OSI) model of networking (as well as the human factor).

A CDS must balance the goals of upholding the security policies of connected networks and security domains (to a level of assurance appropriate to address identified risks) and enabling information sharing outcomes for the business. However, these goals are often in competition. Balance can be achieved by adopting a risk-based approach to cross

domain security, but the actual risks need to be considered and managed, with decisions revisited throughout the life of the system.

For these reasons, it can be said that secure CDS technology is only a part of cross domain security. A comprehensive approach to security will also encompass organisational policies, processes and procedures and be complemented by additional controls in the surrounding environment to provide layered defence.

Business objectives and threats

As we have seen, a CDS provides numerous benefits for organisations and their users, including:

- increased business interconnectivity (e.g. getting the data you need, where you need it)
- increased efficiency through automation of processes
- improved accountability and compliance with governance and security requirements.

If a CDS is poorly implemented, threats to business objectives may include:

- loss of efficiency
- disproportionate security imposts (e.g. increased cost, time and business process)
- missed opportunities.

Information security objectives

As the purpose of a CDS is to uphold security policies, it is important that information security objectives are well understood prior to commencing design or implementation of a CDS. These include:

- Confidentiality: the assurance that information is disclosed only to authorised entities.
- Integrity: the assurance that information and systems are not modified without authorisation.
- Availability: the assurance that systems are accessible and useable by authorised entities when required.
- Authenticity: the assurance that the identity of a user, process or device is known and verified, as a prerequisite to allowing authorised access to resources in a system.
- Accountability: the assurance that the origin and integrity of data has been proven, or that an authentication can be asserted to be genuine and any subsequent actions cannot be denied (also known as non-repudiation).

Secure design and operation of a CDS is only achievable if the outputs of each stage of the system lifecycle are traceable to security requirements derived from these information security objectives. As such, systems must:

- protect the security of information (e.g. avoid import and execution of undesired and malicious content, and mitigate the uncontrolled export of sensitive or classified data)
- protect the security of the systems that store and process that information, in their respective security domains and within a CDS itself (e.g. defend against network attack and monitor for attempted breaches).

The following diagram illustrates the information security objectives for securing a cross domain connection.



Common cross domain threats

Threat actors

External threat actors, typically attempting to compromise the high side (i.e. more sensitive or trusted security domain) via the low side (i.e. less sensitive or trusted security domain), might include:

- issue-motivated groups seeking to disrupt or embarrass governments, international organisations or multinational corporations
- state-sponsored groups conducting espionage on behalf of foreign entities
- cyber criminals using the compromise of legitimate businesses for illegal benefit
- legitimate low side users attempting unauthorised access (these could also be considered internal threat actors if they are part of the same organisation).

Internal threat actors, typically enabled via high side access, might include:

- authorised users, either non-privileged or privileged, enabling a data spill or system compromise by accident
- authorised administrators misconfiguring or failing to patch a CDS or other security system, which weakens security and enables a data spill or the provision of malware
- malicious or deceived insiders conducting human-enabled espionage or interference.

Other threat actors, enabled through access to the cyber supply chain, might include:

- system developers making a poor security decision or mistake during system design, development or integration
- external threat actors attempting to interfere in critical security functionality during system design, manufacture or delivery.

Ultimately, the source of a threat is unlikely to matter. Due to the sensitivity and therefore desirability of information protected by a CDS, they should be designed to be resilient against all possible threat actors.

Threats

First-order threats to security domains needing protection by a CDS may include:

- threat actors with low side access deliberately or accidentally attempting to enable malicious content to pass to the high side, thereby threatening high side integrity or availability via a low side to high side data connection
- threat actors with high side access deliberately or accidentally attempting to enable data to spill or leak onto the low side, thereby threatening high side confidentiality via a high side to low side data connection
- threat actors with low side access attempting to pass malicious software to the high side in order to enable the leak of data to the low side, thereby threatening high side confidentiality via the high side to low side return path.

Second-order threats to a CDS and their security-enforcing mechanisms may include:

- threat actors with low side access attempting to disrupt or degrade security functionality, or cause an insecure failure condition, and then enabling any of the first-order threats, thereby threatening the ability of a CDS to protect high side confidentiality, integrity and availability
- threat actors attempting to disrupt the operation of a CDS, thereby threatening availability of the CDS and impacting cross domain business operations
- threat actors interfering with security functionality of a CDS within the cyber supply chain, thereby threatening the assurance in the CDS to perform their function and high side confidentiality, integrity and availability.

Third-order threats impacting support systems used by a CDS may include:

- threat actors attempting to disrupt or compromise peripheral systems such as audit, monitoring, timekeeping or identity and access management to indirectly disrupt or compromise secure operation of a CDS, thereby enabling any of the second-order threats.

Note though, the threats and corresponding mitigations for a CDS may be better captured through threat modelling.

Failure scenarios

To support a threat and risk assessment, tailored failure analysis exercises should be undertaken for each CDS. Issues that could lead to some degree of failure might include architectural vulnerabilities, platform vulnerabilities, software vulnerabilities or security management issues. In any case, it is important to consider:

- What can fail?

- What are the consequences of failure?
- Where are the single points of failure?
- What are the redundancies or safeguards against failure?
- Will a failure affect the organisation's reputation or ability to conduct business?
- Will a failure get the attention of security bodies?
- Will a failure get the attention of the organisation's executive?

A CDS should be designed to fail securely in a controlled manner. The disabling, compromise or failure of a single or chosen number of components should not lead to the compromise of a CDS or high side security domain.

Common cross domain risks

When controls for a CDS are inadequately enforced, connections between different security domains may allow threat actors to:

- gain unauthorised access to steal, copy or interfere with information
- establish covert channels (i.e. data paths not intended in the original design of the system or product) into, or out of, systems
- compromise the integrity of trusted systems or data (e.g. audit logs)
- bypass security-enforcing mechanisms
- interrupt the availability of critical systems or services
- propagate by pivoting through less-protected networks to access more sensitive systems.

Note that this list of risks is not complete, and risks specific to each CDS will need to be considered. Tailored threat and risk assessment activities will assist in identifying such risks, along with potential mitigation strategies and controls to address them.

Cross domain security principles

The cross domain security principles are a set of ideals that inform best practice for the secure implementation of a CDS, which in turn help determine the security functionality, controls and other mitigation strategies required to meet a given use case and threat environment. The mapping of cross domain security principles to the functions and components in a system design is often illustrated, for the purpose of discussion, in a security-focused logical decomposition or 'security at a glance' diagram.

To address the threats and risks described earlier, a CDS must contain context-appropriate security functions that are applied in a considered pattern and configuration, and implemented to a sufficient level of assurance. These security functions can be explained as:

- A CDS contains context-appropriate security functionality including controls to implement security-enforcing mechanisms and other security functions required for the specific cross domain application, as well as to protect the CDS system itself. These security-enforcing mechanisms must be securely configured to address content-based, protocol-based and trusted insider threats.
- Security functions and other components of a CDS are applied in a considered pattern and configuration implemented according to a sensible system architecture that ensures complete coverage and reduces the risk of any bypass of controls.
- CDS systems and components are implemented to a sufficient level of assurance with enough trust that the system will operate consistently as expected, is hardened against unauthorised modification and will fail only in a predictable and secure manner.

In the context of risk management, these requirements can be translated into the cross domain security principles of:

- context-appropriate security-enforcing mechanisms
- secure architecture and design
- system assurance and secure operation.

These cross domain security principles must all be achieved and maintained in order to address the information security objectives of a CDS system. Conversely, if any of these cross domain security principles were to fail then the overall security of a CDS would likely be compromised – which would impact on the security of the connected high side security domains.

Risk management of a CDS

A number of factors complicate the risk management of a CDS. By definition, a CDS connects discrete security domains which may be operating under separate administrative or organisational authorities. Additionally, technical risks will exist in any CDS. Outside of technical mitigation strategies, a significant number of security policy and procedural measures will likely be required to ensure access to, and use of, any CDS are adequately controlled.

Risk managed approach to the implementation of a CDS

Involvement of security authorities throughout the development of CDS capabilities is crucial to ensuring that decisions are well understood in the context of an organisation's risk management framework and expected business outcomes. The implementation of a CDS should solve security problems by design rather than blindly apply controls to an existing system.

Projects involving a CDS should engage with their organisation's security teams and CDS advisory bodies early and often to ensure that security problems and risks are comprehensively understood and managed. An informed risk acceptance decision is only achievable as a result of a security assessment of system architectures, controls and documented risk mitigation strategies. Critically, security should be baked in throughout the development and implementation process and not bolted on at the end.

Common misconceptions about a CDS

Through extensive involvement in a variety of projects implementing a CDS, ASD has identified misconceptions surrounding the use of CDS products and systems. The most common of these are addressed below and should be factored into capability planning.

'A CDS is one size fits all.'

Although CDS components may be re-usable, the risk assessment and risk acceptance activities relevant to one CDS are not automatically transferable to another. A tailored and considered risk assessment is necessary when expanding CDS capabilities beyond the specific well-understood use cases and threat environments which formed the basis of the original risk acceptance activities. Additionally, new use cases may also inadvertently reduce the security of the initial implementation. Therefore, changes in the security posture of a CDS must be thoroughly assessed against business requirements to ensure residual risks are fully understood and acceptable.

'A CDS is sufficient protection in isolation.'

A CDS is only capable of a finite set of functions, which are based upon the sum of components used to satisfy security requirements. The effectiveness of a CDS can be undermined by weak security policies or poor security practices within connected security domains. Management and monitoring data from a CDS should be analysed and correlated with other security events to ensure a comprehensive understanding of the security environment.

Due to the effectively unlimited methods for encoding hidden data among valid data (i.e. steganography), a CDS handling complex data types is unlikely to eliminate the risk of data exfiltration entirely. A CDS can only reduce the bandwidth of covert data channels available to any given threat actor. Reliance on dirty word filtering alone will not stop malicious attempts at data exfiltration.

'A CDS is secure out of the box.'

Connecting different security domains is inherently risky. A CDS will not achieve intended business and information security objectives if deployed without an accurate understanding of business and security requirements. A measure of a system's security is impossible without the context of the intended deployment.

Furthermore, a CDS is typically not 'plug and play'. A solid understanding of the information and corresponding data structures that are to be shared, and the associated security policies to enforce that data flow, is necessary to ensure security measures are appropriate and comprehensive.

'A CDS is straightforward and inexpensive to deploy successfully.'

A CDS is a complex integrated system of hardware, software, business processes and operational policies. As a result, the financial and resource costs to organisations and integrators in developing, deploying and supporting a CDS is likely to be substantial. Operating costs for securely managing a CDS is likely to significantly exceed the budget allocated to other systems of a similar scale. Costly requirements that are often overlooked include a secure cyber supply chain, secure transport and disposal, cleared personnel, and specialist practitioners skilled in administering a CDS.

'A CDS is the same as high assurance IT equipment.'

To be effective, a CDS must be built and tested to demonstrate an appropriate level of assurance in order to uphold the policies of the connected security domains. A high level of trust is placed in CDS components but, as the complexity of these components increases, the determination of a high level of assurance becomes increasingly difficult. Whilst an entire CDS capability is unlikely to be evaluated through ASD's High Assurance evaluation program, the use of high assurance IT equipment as components (e.g. data diodes) is strongly encouraged where appropriate.

For complex systems and platforms that implement a CDS as an enabling or supporting function, the CDS and its security functions should be considered as critical to the security of the overall capability (and vice-versa) during risk assessment and risk acceptance activities.

Avoiding the need for a CDS

In many business cases that propose introducing an additional CDS, the optimal solution for addressing business requirements may be to leverage a pre-existing CDS or even to avoid using a CDS altogether. Alternative approaches to implementing an additional CDS can include:

- using existing enterprise CDS (eCDS) capabilities where available
- using an existing CDS where available and appropriate, noting that a non-enterprise CDS is highly tailored to both the environment and data requirements
- changing business practices, noting the potential for increased business risk (i.e. working within existing security domains and reducing data transfer requirements)
- using existing data transfer procedures, noting a potential for increased risk compared to a CDS if insufficient security enforcement, such as content filtering, is applied to data transfers.

For use cases involving SECRET or TOP SECRET networks, connecting these security domains without use of an assured CDS is in violation of the guidance within the ISM and an organisation would be doing so at their own risk.

Security domains explained

The following content introduces security domains, a foundational concept necessary to contextualise a CDS. It also introduces a number of sub-classes of security domains and the relationships between them.

Security domains

Generically, an information domain is 'a system, or collection of systems in a network, within a defined administrative, functional or security boundary'. More specifically, a security domain is 'a system, or collection of systems in a network, operating under a consistent security policy and under the ownership or control of one organisation'.

In the case of security domains, the security policy will define the security classification, releasability caveats, community of interest and any other special handling caveats for information stored and processed within it. The security policy will also inform the physical and personnel controls needed to protect the security domain. The concept of security domains is central to a CDS and the discipline of cross domain security.

In practice, security domains can be implemented as an organisation's default network domain at a given security classification with a specific security profile and policy, or as an enclave within a network. Security domain boundaries may differ where multiple organisations access a common system under different security policies that define their administrative or ownership responsibilities.

Examples of security domains include:

- a system at a single security classification (e.g. PROTECTED, SECRET or TOP SECRET)
- a community of interest with releasability restrictions or special handling caveats within a single security classification (e.g. Australian Eyes Only, Australian Government Access Only, Releasable To, Originator Controlled or No Contractor)
- a development or test network with a more permissive security policy, but a more restricted user access policy, when compared with a production network
- any external network of which there is no defined security policy and unbounded risk (e.g. the internet).

A security domain, being a type of network, can be depicted in a network diagram as an unspecified collection of systems using a cloud-shaped image.



Security domain enclave

A collection of systems connected by one or more internal networks under the control of a single authority and security policy, within a wider security domain, is considered to be an enclave of that security domain. In short, an enclave is a security domain within a security domain.



As a result of an enclave's reliance on the surrounding security domain to provide key security functionality, there is a risk that any compromise of that security domain will lead to a compromise of the enclave within it. An example enclave might include a SECRET Australian Eyes Only enclave within a broader SECRET network.

High side and low side

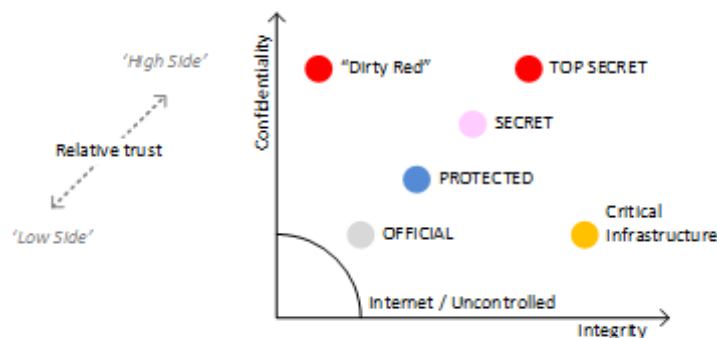
When two security domains are considered together, one security domain will be more trusted than the other. The more trusted security domain, typically with a higher security classification or caveat, is known as the high side, or occasionally the red side or domain high. The high side security domain is likely to have more restrictive access requirements, handling requirements and security policy. It should maintain higher levels of integrity and confidentiality, as well as higher levels of availability where this doesn't compromise other security properties.

The less trusted security domain, with a lower security classification, caveat or security policy, is known as the low side, or occasionally domain low. It will be considered to hold lower levels of integrity and confidentiality, particularly when this security domain has no defined security policy (e.g. the internet, which is an environment uncontrolled by the organisation).



Sometimes the distinction between high side and low side security domains will be difficult to identify as the security classification of each security domain will be the same. If this is the case, CDS system owners assume responsibility for the high side security domain, while all other connected security domains are considered to be the low side.

The following diagram demonstrates relative trust between example security domains through their respective confidentiality and integrity requirements.



The high side and low side terminology is often more relevant in scenarios where confidentiality is considered to be the most important security property. The relationship between high trust and high sensitivity may be different for systems where integrity is the most important security property, such as a CDS designed to protect critical infrastructure.

Domain boundary

A domain boundary is a logical delineation around a security domain wherever systems that are contained within that security domain can interface externally, such as with other security domains.

If there is no permanent interface to any external systems, the domain boundary will exist as an air gap since the necessary logical separation is also supported by physical separation.

In a typical cross domain environment, the domain boundary (also considered as a trust boundary) will be located within a CDS. More specifically, the actual domain boundary within the architecture of a CDS will likely be located wherever a network bridge, protocol break and/or one-way data flow control (e.g. data diode) occurs. Some system architectures may otherwise define a staged boundary by introducing an intermediary zone or security domain separated from other security domains between a pair of data diodes or similar controls. This is known as a diode sandwich design pattern.

Air gap

A high side security domain that is separated or segregated from lesser security domains through physical separation is considered to be protected by an air gap. Notably, a stand-alone network or system will maintain a high level of security, possibly at the expense of utility. Data can only enter or leave air-gapped environments via removable media or a secure CDS capability.



There is a risk of compromising an air-gapped network when performing routine system maintenance and patching. Users may also attempt to transfer data across the air gap in an uncontrolled fashion. Note that as soon as you pass removable media between systems, you have effectively removed or bridged the air gap. There are public examples of threat actors and malicious insiders successfully compromising an air gap via removable media.


Physical isolation via an air gap is likely to be maintained, or at least very carefully controlled, wherever extraordinary security is essential. Examples include military and intelligence networks, industrial control systems, financial services, and life-critical systems such as medical equipment and avionics.

Security zone, trust zone and demilitarised zone

In networking, a security zone is a logical area in a network with a defined level of security. This security is typically defined relative to the security of another security zone such as the internet (where there is zero trust). Security zones layered in this way are also known as trust zones, which are used as a mechanism for establishing defence-in-depth in a network design.

Unlike a proper security domain, a security zone or trust zone may not have its own defined or enforced security policy. Rather, security enforcement will typically be limited to logical network separation and access control rules.

A demilitarised zone (DMZ) is a specialised form of security zone with one or more servers that are kept separate from the core network, either on the outside of the core network firewall or as a separate network protected by the firewall.



DMZs are used to prevent direct access to information and services on internal networks. DMZs usually provide a selection of information and services to less trusted networks, such as the internet. A DMZ does not necessarily provide additional security functionality, it simply provides a level of network separation for servers that must be made available externally. Therefore, a DMZ may be seen as both a type of security domain and type of trust zone.

If the DMZ and internal network are protected by a single firewall, also known as a three-legged firewall, this is likely to present a single point of failure between public and internal zones. For example, a public web server located within a DMZ will allow access to that server from both the internet (low side) and the internal network (high side) whilst denying untrusted users any access to the internal network.

Connections between security domains

Temporary connections between security domains

Removable media is often used to facilitate data transfers between security domains that are separated via an air gap. Note that a logical connection is still made between systems when removable media is used, thereby introducing associated cross domain risks (e.g. air-gap-jumping malware) that will need to be acknowledged and addressed. However, this connection is of a temporary or ephemeral nature only.

Benefits of creating a temporary connection include a lack of persistence and a reduction in the threat or attack surface of the high side system. However, unless additional controls are introduced to supplement removable media processes, the lack of comprehensive data protections (e.g. content filtering, detailed audit and provenance checking) may allow a threat actor to gain a foothold on the more sensitive security domain. Existing removable media policies for sensitive or classified networks are still a cross domain problem, even if a CDS is not used.

Permanent connections between security domains


A data diode, network gateway and CDS are examples of permanent connections between security domains. Without a comprehensive application of security policy enforcement and other controls, such as an assured and risk accepted CDS system, a permanent connection into an otherwise air-gapped security domain is likely to undermine security benefits afforded by the air gap.

Domain security models

Multiple Independent Levels of Security

The standard domain security model is known as Multiple Independent Levels of Security (MILS). Under the MILS model, there is typically one security domain per security classification or sensitivity level per organisation. For example, a government department may have an internet-connected OFFICIAL or PROTECTED network and a SECRET network. The concept of a high side and a low side is an artefact of the segregated nature of the MILS model. Due to its simplicity and the ease of achieving segregation, this model is used extensively in government contexts.

The following diagram illustrates how a user cleared to level 'A' can access systems and objects within 'Domain A', and a user cleared to 'B' can access systems and objects within 'Domain B', including any objects transferred from 'Domain A' via a CDS or other approved data transfer process.




Multi-Level Security

An alternative domain security model is known as Multi-Level Security (MLS). Rather than focusing on connecting independent environments, MLS implementations consist of a single system that holds information tagged at all security levels and enforces access rules accordingly.

Under the MLS model, all information attributes are to be tagged, including security level and access control list attributes. Whilst MLS capabilities can introduce many business benefits, the complexity of managing these systems means it can be difficult and expensive to achieve true security, particularly as the integrity and provenance of information tagging must be maintained for the life of the information.

A practical example of the MLS model is in the SELinux operating system. When it is operating in MLS mode, the secure kernel is able to broker communications between different parts of the operating system that have been tagged with different security levels. Systems built using SELinux, which follows the MLS model internally, can still be incorporated within a MILS environment.

The following diagram illustrates how a user cleared for level 'A' can access information tagged with 'A', and a user cleared for both levels 'A' and 'B' can access information tagged 'A' or 'B'.



CDS explained

This content introduces a CDS in greater detail, explaining the essential functions of a CDS and categorising a number of CDS sub-classes and related systems. It builds on the knowledge contained in earlier content.

CDS overview

A CDS is a security solution that provides the ability to access and/or transfer information between two or more differing security domains, and is capable of implementing comprehensive data flow security policies with a high level of trust. An assured CDS can be described as ‘a robust gateway environment that is designed, built and tested to a specific use case and high level of trust that it will operate as expected’.

A CDS provides information flow control mechanisms at each layer of the OSI model with a higher level of assurance than network gateways. Unlike a network gateway, a CDS employs a wide range of controls to provide comprehensive content filtering and defence-in-depth.



As the complexity of data that passes through a CDS increases, the more difficult it is to gain a high degree of confidence in the verification process handling that data. Subsequently, the risk of leaking data or importing malware increases with this complexity. Constraining the number of allowed data types forms an important part of the specification of CDS capabilities.

Depending on the deployment, the use of firewalls controlling access to a CDS (from both low and high side security domains) may be optional but recommended. Sometimes it will be appropriate to include firewalls in the scope or boundary of a CDS and sometimes they will be included in the scope or boundary of the connected security domains.

A CDS is intended to securely solve two general business requirements: allowing data transfers into and/or out of a sensitive or classified security domain; or allowing simultaneous access to information from multiple sources at multiple security classifications.

CDS functions

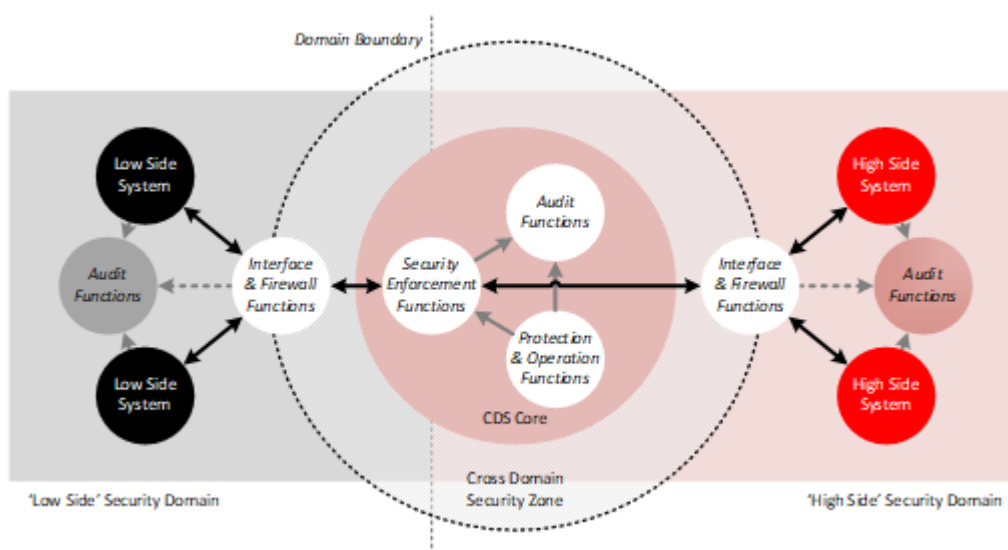
A well designed and assured CDS prevents the flow of information between different security domains by default. A CDS will only permit selected information to pass security-enforcement points based on whether the data conforms to a pre-determined security policy.

To uphold the effectiveness of a CDS, each connected security domain must also implement appropriate controls to protect its core systems and boundary connections. Note that some security functions, such as approval to release data, may actually be located within their source systems rather than within a CDS.

The essential functions of a CDS can be summarised as follows:

- Providing an interface between the CDS core and each of the connected security domains, as well as any relevant management network planes, by:
 - accepting data from authenticated and authorised users and/or source systems for inspection and treatment
 - presenting filtered data in a usable form to approved destination systems and/or users
 - maintaining logical separation between connected security domains
 - reducing the attack surface of the CDS core and the high side security domain.
- Enforcing a security policy for data flowing between security domains, including:
 - blocking all traffic between security domains by default and enforcing the paths that data can transit via the CDS, which typically includes the use of one-way controls
 - performing filtering, data normalisation, transformation and/or sanitisation on traffic to eliminate malicious content and prevent loss of sensitive or classified data

- permitting selected data to pass based on pre-determined security policy rulesets and release approval processes
- operating as a proxy between networks, rather than routing original network traffic.
- Protecting the operation of the CDS, including:
 - providing secure functionality for configuration and management
 - maintaining a secure, verifiable and patched state
 - allowing system monitoring and alerting
 - catching and handling operational errors, and ensuring data channels fail secure by default in the event of any error or failure in a subsystem.
- Maintaining a forensic audit trail, including:
 - maintaining a security audit of the access control mechanisms for all system elements
 - maintaining a security audit of the data channels and security-enforcement decisions
 - maintaining a security audit of the system state and any changes to configuration.



Comparison with network gateways

A CDS could also be described as a robust network gateway that is designed, built and tested to a high level of assurance, and tailored to a specific use case or set of use cases. Unlike a network gateway or network firewall appliance, a CDS employs a wide range of controls to provide defence-in-depth.

A network gateway is considered much less robust when compared to a thoroughly architected and assured CDS, as many important controls are not present in a typical gateway implementation. For this reason, network gateways are restricted to lower risk connections between security domains, while a CDS comprising security functions with higher levels of assurance is used for connections that are higher risk.

The below table offers a side by side comparison between a CDS and a network gateway.

Cross Domain Solution	Network Gateway
Connects security domains across multiple trust levels	Generally connects security domains at the same trust level (although network gateways are also used between OFFICIAL or PROTECTED networks and the internet)
Robust content filtering at the application layer	Protocol filtering at the network and possibly application layer

Controls application transactions	Controls network connections
Few network services permitted	Many network services permitted
Breaks data transport protocols	Generally no protocol break
Uses trusted platforms	Generally no support for trusted platforms
Uses multiple trusted subsystems	Generally a single device or appliance
Typically employs tailor-made solutions, or government-off-the-shelf (GOTS) or military-off-the-shelf (MOTS) products, in addition to COTS products	COTS product
Higher levels of security assurance	Generally lower levels of security assurance

Types of CDS


Functionally, a CDS needs to explicitly control the flow of data between security domains, enforcing a defined security policy through technical means. However, depending on the use case, a CDS may be categorised into one of the following types:

- Transfer: to facilitate information and/or files moving between security domains.
- Access: to serve multiple user desktops or applications that may be hosted in different security domains.
- Multi Level Security: introducing a separate security domain that provides granular access controls for each data object.

Transfer CDS (uni-directional)

This is a security solution that facilitates the transfer of information between security domains in one direction only, such as low side to high side or high side to low side. A uni-directional transfer CDS will typically consist of a one-way flow control (e.g. data diode) surrounded by specialised data processing components.

A CDS transfer from the low side to the high side may otherwise be described from the perspective of the high side as data import, whilst a CDS transfer from the high side to low side may be described as data export. The following diagrams illustrate the data import and data export use cases respectively.




Note that the grey boxes in the above (and following) diagrams are intended to represent a generic collection of security-enforcing mechanisms within a CDS. The internal architectures of a real CDS may differ from this abstraction.

Transfer CDS (bi-directional)

This is a security solution that facilitates the transfer of information between security domains in both directions, with separate low side to high side and high side to low side data flows. A bi-directional transfer CDS should consist of a physically separate upwards and downward path to enforce separation of data flows through one-way flow control

mechanisms. A common implementation is to use a pair of data diodes surrounded by servers running specialised data processing components including independent filters tailored to the data flows expected in each direction.



Transfer CDS (enterprise)

An eCDS is a large-scale transfer systems, or array of systems, designed and deployed to service the general cross domain transfer needs of a whole organisation. They will usually connect multiple security domains with a common framework for operation, management, security policy, audit and monitoring.

The business requirements driving adoption of an eCDS can introduce significant challenges to achieving cross domain security. An eCDS must be designed with great care to ensure risks are appropriately managed. Risks specific to an eCDS include:

- attempting to meet all transfer requirements of an organisation in a single solution may increase design and assurance complexity
- providing additional security domain connections may introduce a higher attack surface compared to point-to-point solutions
- enforcing and separating data flows may become more difficult when multiple, and multi-directional, flows are allowed.

Note, some business requirements may be better met by a point-to-point transfer CDS deployed alongside an eCDS. Further, noting the risks involved with implementing new eCDS capabilities, an alternative approach may be to rely on a common architecture or design pattern to deliver multiple discrete solutions for different use cases or connection requirements.

Transfer CDS (specialised implementations)

Other CDS implementations exist as COTS, GOTS, MOTS or custom capabilities to meet specific business requirements. These include a CDS for voice and video teleconference gateways, mail gateways, database replication, audit collection and streaming video.

Regardless of the business requirements being met, a specialised CDS will ultimately be an implementation of one of the basic types of CDS already described. Specialisation is derived from performance optimisations and the addition of a tightly-constrained security policy that is tailored to address threats and risks applicable to the niche data requirement.

Cross domain guards

A network proxy that also applies some form of security filtering of data, typically when data is to be released from a network, is often referred to as a guard. Some CDS documentation may reference cross domain guards, particularly to describe a core security-enforcing mechanism such as a software appliance intended to provide some cross domain security functionality. Whilst this terminology is sometimes used interchangeably with the more specific description of a CDS, a cross domain guard may not contain the full suite of security enforcement necessary for a best practice CDS.


Access CDS

This is a security solution permitting access to multiple security domains from a single client device. Although user interface data does traverse between systems, these systems do not allow file-based or user-initiated data transfers. Information will remain within their respective security domains, but the user will need to be cleared to at least the level of the highest security domain.

In practice, an access CDS is a specialised form of a transfer CDS that has been designed to only allow user input and system output to be transferred between a CDS core and any connected security domains.

Current access CDS products rely heavily on secure operating system kernels with separation via mandatory access controls and a specialised hypervisor. As virtualisation is used, known threats to virtualisation such as hypervisor escapes should be considered as part of any risk assessment.


A further risk with an access CDS is that users may maliciously or accidentally input information in the wrong security domain as they are both displayed on a common monitor. This may occur directly through data input devices, such as keyboards, or through copy and paste functionality if made available to users. Note that the risk of compromise of client devices will also need to be considered when designing an access CDS.



There are additional approaches to achieving an access CDS that use existing devices within one security domain and provide a controlled tunnel into a different security domain. These are also known as browse-down or browse-up solutions and are explained below.


Access CDS (browse-down)

The browse-down model is an alternative form of an access CDS that allows users to access a representation of a low side desktop from their regular high side client device. Benefits of the browse down model may include avoiding the creation of an all-new client device, along with associated scalability and space, weight and power improvements. A browse-down CDS may also be used to browse-across to systems with a similar trust level.



Access CDS (browse-up)

The reverse of the browse-down model, which might be described as browse-up, introduces a high level of risk to the high side as the trust model is also reversed. This model should never be used to access classified information or systems from an untrusted client device. There is a risk that an untrusted low side client device could intercept classified information from the display and inject unauthorised commands for execution on the high side system.




MLS systems

The MLS model allows a MLS system to process information with different security classifications by tagging information with security classification markers. This in turn permits simultaneous access by users with different security clearances

or permissions. The intention of this approach is to promote collaboration whilst preventing users from obtaining access to information for which they lack authorisation.

MLS systems permit access to less classified information by personnel holding a higher security clearance, while also enabling them to share sanitised information with personnel holding a lower security clearance within the same system. In this situation, sanitisation of information refers to the removal or modification of content so that the information no longer requires a higher security clearance to access. As a result, MLS systems remove the need to duplicate information across different security domains.




One key security feature of MLS systems is that users cannot directly access the connected security domains. Therefore, a MLS system could be desirable when users aren't cleared to the same level but need to work together with varying information within the same environment.

Systems employing MLS are necessarily complex with numerous unbounded risks. As such, it is very difficult to provide assurance. ASD should be contacted for further advice on the use of MLS systems as no such authorised implementations exist in Australia at this time.

Combined access and transfer CDS

In some instances the aforementioned types of CDS may be combined to permit users to access and transfer information in one or multiple directions between security domains from a single client device. This architecture is effectively a combination of one or more access CDS and transfer CDS between two or more security domains. In such cases, information will not be transferred between respective security domains, except where expressly nominated and authorised for transfer. As with an access CDS, users will be cleared to at least the level of the highest security domain. A combined access and transfer solution is likely to require integration between features to provide users with a seamless experience.

These types of CDS have previously been described as a multilevel CDS due to their use of a combination of CDS technology to provide some outward functionality of MLS systems. The following diagram depicts a possible implementation that combines access and transfer CDS functionality.



Cross domain security principles

As noted earlier, a CDS must address the cross domain security principles of:

- context-appropriate security-enforcing mechanisms
- secure architecture and design
- system assurance and secure operation.

These cross domain security principles must all be achieved and maintained in order to address the information security objectives of a CDS capability. Conversely, if any of these cross domain security principles were to fail, the overall security of a CDS would likely be compromised – which would impact the security of connected high side security domains.

Detailed descriptions of how a CDS can achieve these cross domain security principles are described below.

Context-appropriate security-enforcing mechanisms

Security policy enforcement

A CDS will enforce a security policy, developed to meet an organisation's information sharing requirements, whilst upholding the security and risk acceptance assumptions of the high side security domain.

The security policy must consider the information being shared (i.e. the semantic content) and the data format used to convey that information (i.e. the syntax or syntactic structure). A CDS should validate all file content, and any other data, against a schema of known and expected traffic. Furthermore, a CDS will ensure all information, including data type, content, source and originator, is approved for release via conformance to the security policy or a securely-implemented manual intervention process that reveals any hidden content to a knowledgeable reviewer.

Further information on security policy enforcement can be found in the ISM's *Guidelines for Gateways* and *Guidelines for Data Transfers*.

Block first, ask questions later

A CDS will block all network traffic by default, maintaining the air gap and only allow information to transit once it passes security enforcement.

Blocking all information transiting from one security domain to another, except where explicitly approved according to a defined security policy, can increase the level of assurance that information transiting a security domain is legitimate and benign. A CDS should also incorporate firewall and gateway best practice and drop all unexpected protocols.

Malware prevention

A CDS will filter data and information moving from low side to high side, to protect the integrity of the high side security domain from malicious content.

This will include employing traditional signature-based antivirus techniques to identify known bad content, as well as behavioural analysis and advanced content filtering where appropriate. Where possible, passive neutralisation of malicious content is more effective than traditional malware detection and prevention techniques. This can be achieved through lossy transformation of original data formats into alternative formats that ensure business information remains present (i.e. the data is semantically equivalent) but the binary content and metadata has been modified (i.e. the data is syntactically different).

Data loss protection

A CDS will filter data and information to mitigate its inadvertent or deliberate transfer onto a system not approved to handle it.

For fixed-format messages, a data field should be dedicated to the sensitivity or security classification of messages. A CDS will then check that field to ensure the destination security domain can handle it.

For systems that employ a Reliable Human Review process, system developers should note that humans are unreliable at identifying hidden content without technical assistance, this can decrease further when data is more verbose or more frequent. Release approvers must be knowledgeable about the information for which they are attesting, including its sensitivity and appropriateness for release.

Access authorisation

A CDS will ensure all users, including non-human system user accounts, are authenticated and authorised to access the CDS and related security functions.

If it is necessary to receive information from unauthenticated sources (e.g. web or email servers), a CDS should implement cyber security best practice to confirm the authenticity of each data source, such as through verification of Transport Layer Security certificates. This measure would also support provenance checking, as described below.

Data provenance

A CDS will check where data has originated to ensure it aligns with expectations and that provenance is maintained through any links between source and destination systems.

A chain of provenance may be based on the physical design of CDS systems, cryptographic techniques or other approved means. A CDS will also ensure data is not accessed or modified, other than by trusted and approved users or processes. A CDS should support best practice data-in-transit protection for all connections (e.g. point-to-point encryption).

Transformation and normalisation

A CDS, where possible, will convert file types into a standard and normalised format.

Content transformation processes have the potential to remove malicious or hidden content within files or to make malicious content inert. For this reason, transformation is a critical security function for complex document types, particularly those being transferred from low side to high side.

Advanced content filtering, also known as content disarm and reconstruction, is a process for deconstructing a file, removing all elements that do not conform to the file's type specification, and then rebuilding it into a clean known good version suitable for transfer across a CDS. Files embedded within archives or other documents should be recursively extracted and processed individually.

Following transformation and normalisation, data can be validated against a schema of known and expected traffic. Ideally, data will only be accepted in specific file types that can be analysed against the schema.

Protocol break

A CDS will act as a network proxy and terminate transport protocols at multiple layers of the OSI model for all network traffic.

This protocol break helps a CDS protect the high side from malicious network traffic and application layer protocols. For example, at the transport layer (OSI model Layer 4), Transmission Control Protocol connections may be broken and retransmitted using the User Datagram Protocol.

One-way flow

A CDS will enforce uni-directional data flows using one-way control components and one-way inter-process communication paths.

The direction of data flows will not be modified during the operational life of a CDS without approval. Data transport across a one-way control, such as a data diode, typically uses a stateless protocol facilitated by a protocol break.

Flow control

A CDS will monitor the size, volume, quantity and types of traffic and take action when this traffic exceeds defined thresholds.

Unexpected traffic, such as that arriving from an unknown source, on an unused network port or at an unusual time, should trigger a security action (e.g. alert or block) as this may indicate a misconfiguration or compromise of upstream services.

Secure architecture and design

Domain isolation

Organisations will have an architecture that separates information and systems into stand-alone security domains with a common security policy.

This will most likely be done according to sensitivity or security classification, and is a precondition for the most effective use of a CDS. These isolated security domains, with physical, logical and/or cryptographic separation, are said to be protected by an air gap reflecting the deliberate lack of connectivity to external environments.

A security domain adheres to this cross domain security principle if the only way to transfer or access information from an isolated security domain is via an assured CDS or other approved process that maintains logical network segmentation and segregation.

Further information can be found in the [Implementing Network Segmentation and Segregation](#) publication.

Pre-approved patterns

A CDS will follow architecture patterns and/or design patterns, where approved and available to system developers, to ensure risks related to common security problems are addressed consistently and securely.

Architecture patterns are being developed for a number of common business requirements and can be requested from ASD.

Tailored solutions

A CDS will be tailored to the specific business, operational and security requirements of CDS system owners.

Although CDS components may be re-usable, particularly where pre-approved architecture or design patterns have been used, the risk assessment and risk acceptance activities relevant to one CDS are not automatically transferable to another.

A tailored and considered risk assessment is necessary when expanding capabilities of a CDS beyond the specific well-understood use cases and threat environments which formed the basis of the original risk acceptance activities. Additionally, new use cases may also inadvertently reduce the security of the initial implementation. Therefore, changes in the security posture of a CDS must be thoroughly assessed against business requirements to ensure residual risks are fully understood and acceptable.

Defence-in-depth

A CDS will employ layered controls such that a single failure will not compromise security.

A CDS will contain redundant components that eliminate single points of failure at all levels of the system stack. For larger and more complex systems, there should be diversity in platforms, operating systems and software components (e.g. by employing diverse implementations of the network stack across system components along the data path).

Secure-by-design

A CDS will be designed to address security concerns first and foremost with functionality added as necessary to meet business requirements.

CDS system designs should solve security problems rather than blindly apply controls to an existing system. Critically, security should be baked in throughout the development and implementation process and not bolted on at the end.

Simplicity

A CDS will be designed to be as simple as possible while still meeting security and business requirements.

Complexity in CDS system design will complicate system assurance, operation and maintenance, and possibly increase the attack surface. Furthermore, overly complex systems may introduce additional costs to organisations throughout development, implementation and operational phases.

Software packages and physical interfaces not vital to the day-to-day operation of a CDS should be removed.

Note, when following a risk managed approach, a better security outcome may sometimes be achieved if some controls are relaxed. As long as it can be demonstrated that the overall intent of the control is still met and the decision reduces the overheads in managing an otherwise complex system.

No bypass

A CDS will ensure that critical security-enforcing mechanisms are not able to be bypassed.

Security-enforcing mechanisms are actively coordinated by a filter orchestration process and/or passively by the system's architecture. A CDS will also break and inspect any encrypted data in transit to ensure enforcement can occur.

It should not be possible to bypass a CDS or its security-enforcing mechanisms. For example, physical and personnel controls should protect a CDS against tampering.

Separation

A CDS will separate data flow paths to ensure that appropriate enforcement is applied on a data path basis.

Data paths are optimised to address the specific threats to either high side confidentiality or integrity. For example, data moving from low side to high side security domains will typically be channelled through a different arrangement of security-enforcing mechanisms compared to data moving from high side to low side.

CDS system architecture will also employ network segmentation and segregation internally, creating separation between security zones within CDS components to limit bypass of critical security-enforcement points. Movement between these security zones will be enforced by software and hardware firewalls as well as one way controls (e.g. one-way inter-process communication paths and hardware data diodes).

Attack surface reduction

A CDS will present the minimum functional system interface to connected security domains.

Software platforms should be hardened by removing any unnecessary interfaces and functionality. Consider the employment of access and authorisation controls and the use of firewalls to create protected CDS zones. Like a CDS core, these firewalls should also block all traffic by default allowing only the minimum traffic necessary for system operation. Alternatively, consider security through obscurity (although this is an insufficient security measure alone) such as using a network sinkhole to direct traffic into a logically separated CDS zone. This cross domain security principle also applies to the internal components of a CDS.

Cascaded connections mitigation

The risk of cascading connections to lower trust environments will be identified and mitigated.

Consider any additional connectivity into and out of security domains connected by a CDS and how this might affect the threat model. For example, a data spill might not be contained to just one security domain if a cascaded connection to a further security domain exists and is not accounted for.

System assurance and secure operation

Formal system assurance

CDS system owners will ensure CDS capabilities undergo a formal system assessment process and any recommendations are adopted.

In order to ensure projects that involve a CDS are not stalled by formal system assurance activities, CDS system owners will account for these activities in project schedules and plan for any potential failures or findings of unacceptable risk. To reduce this risk, CDS system owners should engage with relevant security authorities early and follow any advice given.

Trusted platforms and components

A CDS will use platforms and components designed and assured for security, especially for critical security-enforcement mechanisms.

A CDS should use components that have been evaluated under ASD's High Assurance evaluation program, or mutually recognised by ASD, where they are available and suitable. This ensures that technical, physical and procedural protections for components are sufficient to address national security threats. System developers and integrators should also incorporate a secure cyber supply chain for software and hardware components to ensure they are not compromised during development and manufacture.

Platforms and components should also be hardened by following the [Strategies to Mitigate Cyber Security Incidents](#) and other system hardening guides such as the Security Technical Implementation Guides developed by the US Defense Information Systems Agency. Such hardening activities should include the removal or neutering of unnecessary users (e.g. system users, particularly 'root') and functionality (e.g. kernel modules, developer tools and any unused third-party software).

Secure administration

A CDS will employ secure administration practices.

Secure administration practices include, but are not limited to:

- managing configuration out-of-band or from the high side only
- using data-in-transit encryption for administration sessions
- authenticating all administrators and management interfaces
- following general information security objectives of role separation with role-based access control and least-privilege
- treating administrative actions as privileged access and auditing appropriately
- implementing separation of duties such that no single person can alter the operation of components of a CDS from end to end (e.g. firewalls, CDS core components and auditing are administered by different people).

Accountability and detection

A CDS will employ ASD's best practice audit advice to generate and store meaningful audit and logging messages.

A CDS should employ meaningful system monitoring and alerting to aid system integrity and availability, and not be reliant on human observation. This functionality should incorporate monitoring tools that are configured to automatically profile normal business behaviour and alert to significant variations and anomalous behaviour.

Further information can be found in the ISM's *Guidelines for System Monitoring* and in the [Windows Event Logging and Forwarding](#) publication.

Self-protection

A CDS will employ passive and active self-protection measures.

Secure CDS implementations should protect themselves against direct attacks, as well as attacks through data channels, that aim to impact connected security domains. A CDS should also include measures to prevent any compromise that gains a foothold from achieving persistence within the system. Further, CDS components should employ a trusted boot process and only operate in clearly-defined states. For example, system configuration and operation should never occur simultaneously. Finally, security-enforcing mechanisms that are intended to process complex data types should be reset between each session to provide a clean execution environment for filtering software.

A CDS should also implement functionality for identifying emerging vulnerabilities across its application stack in an automated and time sensitive way, especially for any public or internet-facing components. The underlying hardware of security-enforcing mechanisms should also employ physical tamper detection and prevention measures. More complex systems may also incorporate additional intrusion detection and prevention functionality within CDS zones.

Finally, any release approval processes should check the internal provenance of data and verify the integrity of security-enforcing mechanisms that data has passed through (such as being signed by the content filter orchestrator or similar process). In addition to passive protections, CDS command and control metadata should be validated and correlated with other CDS logs (e.g. release authorisations).

Secure failure

A CDS will fail securely, such that the disabling, compromise or failure of a single component or a chosen number of components should not lead to compromise of the security of the CDS or the high side security domain.

The definition of a secure failure depends on the information security objectives that a CDS is intending to achieve. Typically, the objective is to ensure the integrity and/or confidentiality of information, in which case a secure failure will see a CDS block all information until secure operations can be resumed. However, this may be undesirable if availability of information is determined to be the most critical business requirement. In this case, a critical alert should be generated so that an alternative system or method can be used while the failure is investigated as a matter of urgency.

CDS system owners should have a plan in place to manage any component or system failures. Similarly, a CDS should not be allowed to operate in a degraded state where the secure operation of security-enforcing mechanisms and audit functions cannot be guaranteed.

Opaque operation

CDS will operate opaquely and not leak information about the high side or themselves to low side systems or users.

Unsuccessful operations shouldn't provide security-related feedback or otherwise present new information to users.

Security maintenance and operational support

CDS system owners will maintain the security level of their CDS and security-relevant systems.

Administrators will apply patches and system updates to operating systems, software libraries, anti-malware signature lists and other components as soon as possible. In doing so, any updates or modifications should follow approved change control procedures and be reflected in formal system documentation. Furthermore, patches and system updates should have an assured provenance and be sourced from the highest, and therefore most trusted, security domain where possible.

Security review

CDS system owners will monitor and review their security posture regularly, as negotiated with the applicable authorising officer.

CDS system owners should recognise that threat environments and technologies are not static. CDS technologies have matured considerably and continue to evolve since many legacy products and capabilities were first designed.

Note that the risk management process defined in the ISM features a continuous monitoring phase that ensures that the impact of changes to the threat environment and any security-relevant components are considered and risk managed in a timely manner.

User education

CDS system owners will ensure users are trained in the secure use, operation, administration and maintenance of their CDS.

Noting that human behaviours are often as important to security outcomes as controls themselves, a CDS should be designed in such a way that users will be willing to use them while following all associated processes and procedures. Regardless, users should still be informed of the consequences of misuse or attempted bypass of any CDS or other security-enforcing mechanisms.

Users of transfer CDS should be informed when security-enforcing mechanisms will modify their data, and the possible business impacts of that transformation. Users of relevant source systems that are connected to a CDS, particularly those in high side security domains, should be informed when their data might be replicated and presented within another security domain. Similarly, users of relevant high side security domains should be informed when data has originated

from outside of that particular security domain and may be unsafe. Note that informing low side users of the existence of a CDS, or that their data may be replicated, may sometimes run counter to business and security outcomes.

Any access CDS, or other specialist implementations that perform a similar function such as voice and video teleconference gateways, should clearly indicate to users which security domain is currently active.

Security assurance

A CDS mitigates the risks of data connections into and out of networks to a high level of assurance. That is, with a high level of confidence that security functionality operates as intended to protect connected security domains. To facilitate this, understanding the risks specific to each cross domain use case is essential, as a CDS is designed to satisfy an organisation's specific business requirements and risk environment. Furthermore, the planning, design, installation and maintenance of any CDS must be based on the careful management of risk.

Compared to normal cyber security practices, the application of cross domain controls are designed to protect an organisation's most sensitive or classified systems and information, therefore higher levels of assurance are necessary. General cyber security best practice is still applicable to a CDS, however, a well-designed CDS will ensure cross domain security principles are applied and layered such that the failure or compromise of a single component or even a number of components cannot lead to a compromise of any connected security domains.

Some organisations house dedicated advisory bodies who are CDS subject matter experts, and operate on behalf of ASD to provide tailored advice and assistance. ASD recommends that projects involving a CDS engage with their organisation's security team and any CDS advisory bodies early and often to ensure that risks are comprehensively understood and managed. Similarly, if any questions haven't been met by this guidance, please contact your organisation's security team, CDS advisory body or ASD.

Controls and mitigation strategies

Controls expand upon the cross domain security principles in order to provide practical mitigation strategies that can be implemented using a risk management approach to deliver a secure configuration for a CDS. ASD can provide a range of such controls for CDS implementations on request in order to supplement content within the ISM.

Note that not all controls will be relevant to all CDS systems. Similarly, it is not intended that supplementary controls be treated as a compliance list. Rigorous analysis is required to ensure technical, physical, personnel, policy and procedural controls are applied to sufficiently address the cross domain security principles described in this guidance.

As with any other system, the [Strategies to Mitigate Cyber Security Incidents](#), including its Essential Eight, are also applicable to CDS implementations.

Architecture patterns

System architecture patterns, design patterns and other reference architectures for a CDS are intended to assist system developers and integrators undertaking work on a CDS by:

- raising awareness of efficient and secure solutions to common business requirements
- building an understanding of the capabilities and limitations of the system architecture pattern in the context of a wider system
- identifying the role of, and requirements placed on, each component of the system architecture pattern.

Architecture patterns are being developed for a number of common business requirements and can be requested from ASD.

Risk acceptance

Risk acceptance of a CDS should be formalised through a risk management process, which is owned by the CDS system owner and applicable authorising officer. This process will typically follow the below sequence, noting that sub-processes and the process as a whole will be cyclical in nature.

- Conduct risk assessments at key stages of the capability development life cycle:
 - Conduct a security design review, including:
 - developing and reviewing the threat model
 - reviewing the security risk management approach, including security policy, and reporting findings
 - reviewing the system architecture and design, including functional components, and reporting findings.

- Conduct a security assessment, including:
 - developing and reviewing the test plan
 - performing a vulnerability assessment to validate the security of the system's design, and reporting findings
 - performing hands-on testing as necessary (typically comprising but not limited to penetration testing) to identify technical vulnerabilities, and reporting findings.
- Conduct risk mitigation activities, including:
 - implementing additional controls as appropriate
 - updating system design and security documentation as appropriate.
- Formally accept the residual risk and authorise the system to operate:
 - Determine residual risk, considering the outcome of prior risk assessments and mitigation activities, and report findings.
 - Formally accept residual risk, thereby authorising the system to operate.
 - Update details for connected security domains as appropriate.
- Continually monitor system status and repeat the cycle as necessary:
 - Monitor for any major cyber security incidents involving the system.
 - Monitor for new or emerging threats to the system or its operating environment.
 - Monitor for the discovery that controls for the system are not as effective as planned.
 - Monitor for changes in security policies relating to the system.
 - Monitor for security-relevant architectural changes to the system.

Note that the level of involvement of ASD's CDS Advice and Assessment team during these activities will be managed on a case-by-case basis.

The next step

In addition to publications such as this guidance, ASD provides advice and assistance to CDS advisory bodies and organisations' security teams seeking further guidance on the threats and risks to security-relevant components of a CDS or the strength of security-relevant components of a CDS.

Further information

The [*Information Security Manual*](#) is a cyber security framework that organisations can apply to protect their systems and data from cyber threats. The advice in the [*Strategies to Mitigate Cyber Security Incidents*](#), along with its [*Essential Eight*](#), complements this framework.

Contact details

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