Cross-Domain Solutions

What they are, how they work, and where you use them

waterfall

Cross-Domain Solutions

What they are, how they work, and where you use them

Cross Domain Solutions (CDS) are information technologies and architectures used to exchange information between systems and networks at different levels of trust. CDS is used most often to manage interfaces between "business," civilian, or other Internet-connected networks and sensitive or even classified government networks.

The top priority for CDS is almost always to prevent the theft of data, preventing cyber-espionage, and a second important priority is to protect the integrity of the high-trust network, preventing malware and other intrusions from entering that network.

This guide surveys the spectrum of cross-domain solutions and looks at which solutions are most applicable to which kinds of networks.

It is very common to assign colors to different kinds of networks and sensitivities. There is no standard color scheme across all organizations world-wide. But to give an example, here is the scheme used by the US Department of Defense (DoD):

Classified Purple	Top Secret (Sensitive Compartmentalized Information) - Yellow
	Top Secret - Orange
	Secret - Red
	Confidential - Blue
Unclassified Green	Controlled Unclassified - Dark Purple
	Restricted (or private industry)

Table (1) US DoD color scheme

With this background, let's look at some existing technologies used at domain boundaries.

Quick Reference

Air Gaps

Traditionally, an air gapped system or network is one without any online connection to any external system or network – no Ethernet cable, no Wi-Fi, no optical fiber – "no nothing." The only way to move data into or out of an air gapped network is to physically carry the data – including on a USB thumb drive, a laptop computer, or brand-new computer equipment brought into a system.

Traditional air gaps are still used to protect extremely sensitive networks and information assets but are not considered "cross domain solutions" because there is absolutely no online information that crosses between domains in a traditional air gap.

Caution: in modern usage, the term "air gap" can mean a network or system that is connected only indirectly to the Internet and cannot route packets to arbitrary Internet locations. The term can also mean any high-trust assets that are connected to less-trusted networks only through cross-domain systems. When you hear the phrase "air gap" it is very important to ask clarification as to what the speaker means, since traditional air gaps are much more secure than modern "not Internet routable" or systems protected by the simplest of CDS.

Firewalls

Firewalls forward network traffic between systems and networks, and so technically, all firewalls are also routers. Modern firewalls are routers and much more – modern firewalls look at every message that passes through them and ask the question "is this message allowed?" Sometimes the firewalls answer that question based on characteristics of the TCP transport, such as source IP address and destination port. Other times the firewalls look deeper into the payload of the message using so-called "deep packet inspection" to determine whether the payload protocol and content are allowed.

Firewalls are often used to create layers of networks between two security domains. The "in between" networks are called "demilitarized zones" (DMZ's). Best practice is to use firewalls from different vendors in different layers of DMZs, in hopes that if attackers find a software vulnerability that they can exploit in one layer of firewall, then that same vulnerability will not exist in the next layer of firewall from a different vendor.

"Protocol break" is another term used with firewalls and DMZs. A minimal protocol break is when intervening systems, such as firewalls or proxies, throw away one or more layers of communications headers – for example throwing away the IP and TCP header information – and forward packet payloads with new headers. A deep protocol break is where one kind of communication protocol enters a DMZ through one layer of firewall, terminating in servers or other systems in the DMZ. Those systems communicate information through the next layer of firewalls in a different protocol. In principle, this makes reaching through the firewalls into the high-trust network somewhat more difficult, because additional systems must be compromised in additional stages of a multi-stage attack.

Caution: Layers of firewalls, next-gen firewalls, and DMZs are generally not considered secure enough to act as a CDS in and of themselves. One of many reasons for this is that firewalls are software, and all software has vulnerabilities that may be exploited by attackers. Furthermore, firewalls alone – looking at characteristics of communicated messages – are generally alone not powerful enough to prevent the propagation of malware, remote-control attacks, or data exfiltration attacks.

Data Loss Prevention

Data loss prevention systems (DLP's) are used to manage sensitive data. Such systems generally do the following:

- Label data – with machine-readable labels as to how sensitive is the information,
- Control and track access to the data – both across networks and on USB thumb drives,
- Encrypt data - both at rest and in motion,
- Analyze tracking data to look for anomalous patterns of access that might tend to indicate that large amounts of data, or comparatively large amounts of very sensitive data are being stolen, and
- Provide disaster recovery and incident response teams a suite of technologies and procedures designed to minimize data loss during emergency conditions.

Caution: As with firewalls and all other software systems, DLP systems may have software vulnerabilities that adversaries can exploit, and the machine learning systems for anomaly detection may miss thefts of small amounts of data, or large amounts of sensitive data trickling very slowly out of the system over long periods of time. For these reasons and others, firewalls and DLP systems together are not generally considered sufficient to serve as a CDS.

Content Filters

Content filters are systems and features that look at files moving into a high-trust domain. The function of filters is to ensure that malware and other attack tools do not enter a high-trust domain to compromise or confuse systems in that domain. To that end, filters can look at file meta-data, commonly enforcing limits on:

- **File size** – typically configured to allow only small files to move,
- **File suffixes** – configured to allow only documents to move, not executable files, and
- **File names** – allowing only certain file names or patterns of file names to pass, typically excluding compressed, encrypted or 'zip' or other file archives with additional layers of embedded files and folders.

Filters can also open up the files and look inside, commonly checking:

- **Content type** – checking that the type of information in the file matches the file suffix,
- **Character sets** – for example ensuring that a file contains only ASCII characters or other character sets, not arbitrary binary data,
- **Basic schema checking** – for example enforcing maximum numbers of rows and columns in a CSV file, and enforcing maximum content lengths on individual values in individual fields, and/or
- **Advanced schema checking** – for example comparing the contents of an XML or JSON file to a schema that specifies the required layout of the file.

Note: Firewalls, DLP, and basic file filtering are widely considered the minimum acceptable cross-domain solution (CDS) for protecting the least sensitive government networks, domains, and data at their boundaries.

Anti/malware – Filtering "Known Bad" Content

Anti-virus systems and more general anti-malware protection systems look at files passing through a CDS and try to identify malware. There are a variety of approaches used by modern anti-malware systems, including:

- **Scanning library files and executable files** – though arguably few or no such files should be passing at all routinely across domain boundaries,
- Scanning complex files such as spreadsheets, PDF files and other files that can contain embedded code, or are otherwise complex enough to potentially embed compromises and malware,
- **Scanning for "known bads"** - so-called "signatures" are patterns of file content that identify known attacks or known malware, and
- **Sandboxing** – opening files using their usual applications in a Windows or Linux virtual machine and watching for suspicious behavior, such as unexpected communication attempts, or filesystem activity.

Any file that is identified as malware, containing malware, or otherwise suspect is blocked from entering the high-trust domain. Anti-malware scanning is often part of online CDS systems that send content through networks and is often deployed at USB-scanning kiosks at physical security checkpoints as well, in hopes of ferreting out malware that may be embedded in devices carried into a sensitive site by employees, contractors or visitors.

Note: Because different vendors' systems have different success rates for detecting malware, anti-malware systems used in CDS applications are often configured to scan content using more than one vendor's anti-malware product. Four vendors' products scanning content in parallel is widely thought to be the minimum acceptable anti-malware deployment in a CDS for the least sensitive networks, and sensitive sites may scan with up to 16 or 20 anti-malware vendors' products in parallel.

Caution: A limitation of anti-malware systems is that they only reject files that seem suspect for some reason. If an attack in a file or other content is thoroughly disguised, that attack can slip past anti-malware systems.

Talk to a cross-domain security expert today about your specific security needs:
info@waterfall-security.com

Content Disarm & Reconstruction (CDR) Systems

Whereas anti-malware systems seek to identify "bad" or "suspicious" content passing across a domain boundary, content disarm and reconstruction systems (CDR) seek to pick out only the "known good" content in files moving between domains. Basic or first-generation CDR systems tend to render "good" content as images.

For example, a basic CDR might take a PDF file, open the file in a PDF reader in a virtual machine and render each page of the PDF on a virtual screen. The CDR then takes the images of each page and assembles the images into a new PDF, one that contains as many pages as the original file, but each page consists of exactly one image – the image of the original page – and nothing else. The PDF of images is then forwarded across the domain boundary.

The same process works for any file whose primary purpose is to be read by people. Basic CDR can often produce Word files with one image per page, Excel spreadsheets with one large image per worksheet, and so on. The result is reasonably safe, though one must watch that the CDR software itself does not become compromised in a cyber attack.

Caution: The result of this kind of basic CDR can be hard to use. Human users can read the text, but the content is no longer easily searchable, nor can cut-and-paste operations easily access snippets of text to embed in email or other documents.

Advanced CDR – Extracting the Known Good

Advanced CDR systems dig deeper into the meaning of complex files. These systems reach into PDF, Word, Excel, PowerPoint and other complex file types looking for elements of the files that the CDR systems recognize as harmless.

For example, a CDR might open a PDF file and extract text, fonts, formatting commands such as boldfacing and italicization, images, positioning information for text and images and so on. While PDF files can be almost arbitrarily complex, any content an advanced CDR does not recognize as benign, such as malware or in Excel a malicious macro, that CDR simply leaves behind. The recognized known-good content is assembled into a new file of the same type as the original and passed into the high-trust domain.

The advantage in using advanced CDR is that the new "known good" document contains only benign content, and generally can be searched, cut-and-pasted and otherwise used normally in ways familiar to users of such documents.

Caution: The new PDF, Word, or other files may be smaller than the originals because something was left behind. It may be that malware or some other attack was left behind, which is a good thing. Or, it may be that some formatting element, content element or other part of the file was left behind, elements that the CDR simply did not recognize as "good." This may mean the new "clean" file contains less information than the original, or less useful formatting, or other missing features or content.

Note: When selecting among providers of this modern generation of CDR's, it is important to understand what kinds of files the CDR understands, and what kinds of content in each kind of file the CDR extracts. Generally speaking, a greater understanding of a wider array of "known good" content types is valuable, provided the CDR does not extract and propagate any content so complex that it is not truly understood and so might harbor malware.

Cross-Domain Workflow & Orchestration

Workflow and orchestration systems provide different ways to organize and manage the movement of information across domain boundaries. Capabilities often include:

- Orchestration may allow different kinds of content and files through different kinds of cross domain systems. One CDR might have a deep understanding of the structure of PDF files but have a weaker understanding of the structure of Excel files, or vice-versa. Orchestration tools can route different kinds of files through different kinds of filters and CDRs, depending on the file type or source or sensitivity.
- Orchestration can support optional manual intervention during file transfers. For example, a user who is typically allowed to send specific files into classified networks may trigger additional scrutiny if they attempt to transfer an unusual file. In such cases, the file may be redirected to a specialized CDR system or a human analyst for review and approval before the data is allowed to move.
- Orchestration systems may support signing and tracking of file transfers – adding metadata to a file to explain who or where it came from, what processing was done on it, and why it was allowed to be transferred across the domain boundary.

Sophisticated cross-domain systems are not 'one size fits all'—they often integrate a variety of CDS tools from different vendors and may also include manual human oversight and inspection when needed. Sophisticated systems route the right content through the right steps before sending the content, or a variation or subset of the content, across the domain boundary.

Multi-Level CDS

Thus far in our discussion we have assumed that cross-domain systems act at network boundaries – at the connection between networks containing less-sensitive or more-sensitive documents or data. Multi-level CDS are different. These are databases or other repositories of documents or data that can be at a variety of levels of sensitivity.

These systems can have many users, each authorized to access information at different levels. To ensure users access only the information they are authorized to access, multi-level CDS:

- Have detailed profiles for each user as to what they can and cannot access,
- Annotate documents and data as to what kind of data it is and how sensitive it is,
- Manage sometimes detailed permission lists for documents and data as to who and what kind of users can access the data, and
- Enforce these permissions and rules when users attempt to access the data.

Multi-level CDS technology often make heavy use of mandatory access control (MAC) mechanisms. The most widely available such technology is Security-Enhanced Linux (SELinux). MAC is a technology that provides access rules over and above standard computer permissions.

Caution: In a multi-level CDS, the configuration of descriptions, permissions and user authorizations is critical – any mistake in these areas risks exposing sensitive information to unauthorized users. MAC can help in this area – use the mandatory rules to describe in very broad strokes what programs and people can access what kind of data, and the system should enforce these high-level rules uniformly. This way, even if there is a mistake made in detailed permissions lists, the mistakenly permitted user can still not access files out of their clearance scope because the broad-strokes MAC rules overrule normal permissions lists.

Caution: A second concern with these systems arises if they are deployed such that they are exposed to networks where both less-authorized and more-authorized users work. While multi-level systems may be very good at managing who has access to what, these are still software systems with potential vulnerabilities. This means that if we connect a multi-level CDS to a less sensitive, less thoroughly secured network, there is a risk that one or more machines in that network are compromised. Those compromised machines could launch attacks and exploits at the CDS seeking to compromise the CDS itself and so bypass its permission checking systems. When deploying multi-level CDS, the level of protection should be based on the sensitivity of the most classified documents handled by the system. All connected networks must be secured accordingly to prevent potential attacks.

Access CDS

Access cross-domain-systems are essentially keyboard-video-mouse (KVM) systems. These are most often used when a highly-authorized user is physically located in a heavily-secured building or part of a building – an area that contains computers and networks containing highly-classified material. That same user may require access to less-sensitive data on less-sensitive networks. When that user requires such data very frequently, it may be impractical to walk constantly back and forth between classified and unclassified parts of the building. Deploying a less-sensitive computer or workstation in the highly classified area is dangerous, because if the less-classified system is compromised, it might be used by an attacker to eavesdrop or use covert channels to steal classified information.

An access CDS solves this problem. Access CDS are based on keyboard-video-mouse technology – systems that let a single monitor, keyboard and mouse be connected to multiple physical computers. A switch connects the trio to one computer or another. The wiring involved is not network cabling, but rather USB, HDMI, and other types typically used to connect keyboards, mice, and monitors to computers.

An access CDS allows a keyboard, mouse, and monitor in a highly classified area to connect to and operate less-sensitive computers in unclassified areas. The risk of information leakage from the highly classified area is minimal, because the CDS can not send files or arbitrary information from the classified area or network into the unclassified computer. The CDS can send only the keystrokes and mouse movements into the unclassified network that the authorized user enters.

Caution: Using conventional KVM systems to switch between low and high security equipment risks cooperating malware communicating systems between high and low security domains using keyboard, mouse and even video hardware. USB keyboards contain CPUs that can be compromised and used as covert channels. The same is true for USB mice. Monitors often hold a small amount of configuration information that is invisible in normal use and can be used as a low-volume communications channel when switching between low and high security systems. Access CDS have features to reduce these risks.

Data Diodes

The US National Institute of Standards and Technology (NIST) defines a data diode as hardware that can send information in only one direction. In CDS, data diodes are deployed most commonly to send information into high-security networks with no possibility of leaking information back out to a lower-security network through the diode hardware.

Globally, there are over 100 data diode vendors. Most vendors are comparatively small, serving only their national government, intelligence, and military needs. This is because many governments are willing to purchase data diode products for classified networks only from suppliers in their own country, with a certified supply chain. In such suppliers, every person who touches the diode product during development, production and deployment has a national security clearance. Personnel and vendors in one country generally cannot be cleared in any other country, again leading to a proliferation of small vendors with limited product offerings.

The "gold standard" in data diodes is optical signalling. Such diode products consist of at least two circuit boards or network appliances: one board contains a fiber optic transmitter and the other a receiver. This means the transmitting board can send information – photons – to the receiving board, but the receiver can send nothing back. The receiving circuit board does not contain an optical transceiver; it contains only an optical receiver.

Common Criteria (CC) certification is widely seen as a measure of the strength of protection provided by data diodes. This is a military-grade certification of the degree to which vendor claims are accurate. Most vendors of data diodes make only a single claim – that the hardware is truly unidirectional and is therefore physically incapable of leaking data through the device from a high-security network into a lower-security network.

CC certifications can be awarded in any of seven levels of strength, from the weakest Evaluation Assurance Level (EAL) EAL1 through the strongest EAL7. In the CC scheme, an important capability is called "HIGH_ATTACK_POTENTIAL" – certifying the vendor's unidirectional claim even in the face of the most sophisticated imaginable cyber attacks. HIGH_ATTACK_POTENTIAL is required of EAL6 and EAL7 and can be added as an optional criterion when evaluating diodes at lower assurance levels.

Data diodes are widely used for sending open-source data into the most secure, most sensitive networks, most often coupled with filtering, CDR, and other systems. While the highest priority for most sensitive networks is preventing data exfiltration, a second priority is very often preventing malware from entering and impairing operations in those networks. CDR, filtering, and other similar capabilities provide this second capability in conjunction with the strong assurances of preventing data exfiltration that the diodes provide.

Caution: Purchasers of data diodes should be aware that not all diodes follow the NIST standard:

- Some diodes use electrical rather than optical signalling. A concern with electrical signalling is that all electric circuits are circular – we cannot send electrons in only one direction for any length of time without building up a large electrical charge. It can be very difficult to inspect an electrical diode and prove to ourselves that it is truly unidirectional, given the circular nature of electrical currents. While electrical data diodes, with very careful evaluation, may be appropriate for handheld or back-pack style devices carried into battle, where optical signalling hardware is too bulky, optical signalling is preferred whenever practical.

- Some products sold as data diodes are implemented entirely in software, not hardware. This is a misuse of the NIST definition – NIST requires that true data diodes be implemented in hardware that is physically able to send information in only one direction. All software has potential vulnerabilities and a compromised software diode risks being used by an attacker to leak classified information out to less-sensitive networks.

- Some software data diodes appear to be hardware devices, because they are shipped as hardware appliances. An engineering inspection of these devices, however, reveals a CPU and software inside, with no unidirectional hardware. These devices are in fact software data diodes.

- Some products are sold as "unidirectional firewalls". This is a misnomer. Essentially all firewalls permit TCP connections and "unidirectional" refers only to which side of the device is permitted to initiate the TCP connections. Once a TCP connection is established, communications on that connection is bi-directional and can be used by an adversary to leak classified information out of a sensitive network.

- Some practitioners and organizations may be tempted to "roll their own" data diodes, by clipping wires in twisted-pair Ethernet cabling or in serial cabling. This is error prone. Most Ethernet and serial hardware will automatically negotiate a bi-directional half-duplex connection if wires are cut, able to leak information. Even if this risk is somehow addressed, modified wiring is sometimes replaced by conventional wiring by well-meaning technicians during routine troubleshooting, re-introducing bidirectionality by accident.

Want to learn more about Data Diodes?

Download our guide to Data Diodes vs Unidirectional Gateways

Get the Guide

Unidirectional Gateways

The US National Institute of Standards and Technology (NIST) defines a unidirectional gateway as hardware that can send information in only one direction, working with software that makes copies of servers and emulates devices. Unidirectional gateway hardware is thus similar to data diode hardware – the difference lies in the software.

Unidirectional gateway software makes copies of servers. For example, consider an Oracle or Microsoft SQL relational database on a less-secure network that contains information that classified applications on a high-secure network need to access. The unidirectional gateway software might log into the low-security database and ask the database for all of the data and changes that are new – that occurred since the last time the software logged in, say one second ago. That query might return a dozen megabytes of changed data. The unidirectional gateway software sends that snapshot of new information through the one-way hardware into the high-security network. There, unidirectional gateway software inserts or updates the data into a high-security Oracle or SQL database. Users and applications on the high-security database can then access the high-security relational database normally, querying for the data they need.

In addition to the one-way hardware and the convenience of server and database replication, unidirectional gateways can also serve many of the functions of a CDR. For example, if a gateway is configured to make a copy of a half dozen tables in a relational database, and those tables have fixed-length fields containing only textual and numeric data, then how would an adversary embed malware into the unidirectional data stream? Yes, relational databases often support "blobs" – large unstructured data types that can store large files, such as executable files but unidirectional gateways can most often be configured to pull and make a copy of only a subset of relational data or other data.

In the relational example, the gateway should be configurable to pull data from only select tables – not tables containing the "blob" data type. In this case, the gateways are pulling "known good" text and numeric content through the domain boundary, leaving unknown "blob" and other potentially compromised data types behind.
In practice, unidirectional gateways are deployed:

- As hardware-enforced unidirectional protection at cross-domain boundaries,
- To make copies of structured data sources, pulling only "known good" data types across to the high-security domain,
- In series with CDR and filtering tools for less-structured data such as files and electronic mail, and
- In concert with sophisticated data exfiltration prevention and intrusion detection tools – as a second and tertiary lines of defense against malicious activities on high-security domains.
- True hardware-enforced, optically isolated unidirectional gateways are powerful cross-domain solutions, with built-in CDR capabilities in the unidirectional connector software, and the option of using the devices in series with file-oriented CDR systems as well.

Note: Specific unidirectional gateway connectors support specific TCP query/response applications and protocols, again by replicating the underlying servers and/or devices. When selecting a unidirectional gateway model or vendor, the types of connectors available for the gateways are a critical selection criterion – is the device in question able to replicate the systems we need replicated, without costly custom engineering? The "depth" of connector integrations is important as well – do the connectors we need have the features we need – from specific data selection capabilities to built-in content filtering, performance characteristics and even high availability / redundant hardware / no-single-point-of-failure options if we should need them in the future?

Note: Many but not all data diode offerings in the marketplace come with basic file transfer and often Syslog replication capabilities. Technically, these connectors make these devices unidirectional gateways. In practice however, the vendors who call their devices data diodes tend to do so because these very simple connectors are the only ones available off-the-shelf. All other connectors, when needed, are available only via costly custom engineering.

Caution: Many unidirectional vendors include a "TCP connector" whose name suggests that it transfers all TCP protocols, making this one connector a universal tool. This does not work. In practice, almost all TCP-based protocols are query/response at the application level, making it impossible to proxy those protocols through a unidirectional TCP connector.

Covert Channels

For the very most sensitive government and military targets, we should expect adversaries to go to extreme lengths to steal information. In addition to the well-understood potential for insider attacks and errors and omissions removing very sensitive data on removable media, in the most sensitive networks there is also the possibility of covert channel attacks. A covert channel is any means of signalling information from one place or system to another that is not a conventional communications channel.

The obvious example of such information signalling or leakage is electromagnetic emissions. If a long Ethernet cable is laid across a room with electromagnetically porous walls, it may be possible that a large antenna in a building across the street could pick up electromagnetic (EM) emissions from communications passing through the wire. Faraday cages are the obvious way to prevent this kind of information leakage – designing buildings containing classified systems with microwave-scale electric meshes in the walls to block these kinds of emissions. In fact, some military-grade laptops and other equipment have faraday-style protections built into the devices, again, to prevent EM emissions.

There are less-obvious examples of covert channels leaking information, most of which rely on cooperating malware. In this case, malware has somehow been introduced into the high-security domain and that malware is actively trying to communicate with malware in a low-security domain, both to exfiltrate data, and for command and control to permit external attackers to operate the malware to search for specific high-value data to exfiltrate. Some examples of covert channels exploited by cooperating malware include:

- Liquid crystal screens emit characteristic noises that change, depending on the image being displayed. What appears to be a high-security screen saver alternating between two images that make subtly different noises could in fact be communicating in audible binary code with a nearby low-security laptop or other system equipped with a built-in microphone. The reverse works as well, with a low-security device audibly signalling command and control to a high-security device with a microphone.

- Most computer speakers and microphones have frequency ranges that are greater than the human hearing range. Signals of up to 10-20kbps can be encoded into ultrasonic carriers from computer speakers to nearby computer microphones, outside of the range of human hearing.

- Modern CPUs often contain detailed thermometers as part of a toolkit to manage the temperature of the CPUs for maximum performance. When low-security and high-security CPUs are physically near enough to each other, such as in the same rack or server cabinet, it may be possible for malware on set of CPUs to do meaningless work on the CPUs for long enough to generate a material amount of heat, which may be detected by the other CPU as a measurable change in the ambient temperature, or the temperature of the detecting CPU without the detecting CPU's workload having changed measurably.

There are many other examples. The general rule that we can use to defeat these covert channels is physical distance. Essentially all of these channels degrade quickly with distance. Do not host low-security and high-security devices or systems physically close to each other. Wherever practical:

- Host different security levels in different buildings, introducing distance and randomness between potential signal sources and destinations,

- Supply electric power to the different security domains from different parts of the wiring infrastructure, again introducing randomness into electric power supplied to the different domains, and

- Do not permit portable low-security devices to come physically near high-security devices, nor vice-versa, because of the risk that both devices are compromised by cooperating malware that will use the opportunity of proximity to start communicating.

You might wonder why meeting rooms designed for the exchange of classified data have such noisy fans, why they have no windows (faraday cage material in the walls), and why we are forbidden from carrying laptops or cell phones into the rooms? The risk of covert channels is one of the reasons for all these precautions.

Talk to a cross-domain security expert today about your specific security needs:
info@waterfall-security.com

Use Case: Government SOC

One example of how cross-domain systems are used is a central government security operations center (SOC). A government may demand that all of a nation's critical infrastructures, from banks to telecom providers to large power plants, provide security logs, alerts, alarms and other information to a central government SOC. In this example, the government authority may choose to deploy cross-domain solutions such as unidirectional gateways between the SOC and the network connections to the source infrastructure providers.

When the only connections into the SOC are unidirectional "inbound" from external sources into the SOC, the government can be assured with a high degree of confidence that confidential information communicated from one infrastructure provider into the SOC is not leaked back out to another provider on a different network connection.

The government can also be confident, at least in the absence of insider attacks or covert channels, that insights gained about adversaries who may be attempting to attack the nation's infrastructure are not leaked back to the adversaries in time for the adversaries to modify their attacks.

Use Case: Situational Awareness

A national intelligence command and control center may well need to gather large amounts information from a very wide range of public and private sources to evaluate in terms of intuiting an enemy's troop movements, election tampering or other subversive behavior. To this end, an intelligence center may host a large computing, storage, AI and "big data analysis" capability. Feeding this capability may be video feeds from most of the world's news networks, news feeds from most of the world's private and government news networks, and possibly even feeds intercepted voice and other communications. Again, an intelligence agency analyzing these large amounts of data generally does not want their conclusions about their adversaries' activities leaked back into the public domain.

In this example, the intelligence agency may choose to deploy high-volume cross-domain solutions such as unidirectional gateways and data diodes in concert with filters and CDRs, all forwarding video, text, image, email, voice and other communications into the analysis center from external sources. Again, when the only connections into the command-and-control center are unidirectional "inbound" from external sources, the intelligence agency can have a high degree of confidence that their classified conclusions will remain confidential.

Use Case: Immigration Records

Some nations may wish to keep detailed records of who enters the nation from abroad – both citizens and non-citizens. These detailed records are of course highly confidential – travellers may be reluctant to enter a country if they know that every detail about their identity, their passports, and their coming and going is available to the public. The government however, from immigration authorities to intelligence agencies and even policing of organized crime, may well need to record and analyze the movement of individuals long into the past, such as when an individual is arrested or suspected of criminal activity.

Thus, for example, a nation may choose to report detailed information about individuals crossing through all of the nation's land, sea, and air ports to a central authority. To assure confidentiality, cross-domain solutions may be deployed both at the ports of entry and at the central government authority. Unidirectional gateways at the ports of entry assure the ports that only heavily encrypted data leaves the ports, and no cyber attack information can enter back into the ports through the one-way outbound gateway hardware. Unidirectional gateways in the central authority assures the authority that only heavily encrypted data can enter the authority, and no personally identifiable information can leak back out.

What these gateways communicate from the ports of entry to the central authority might include Message Queue Telemetry Transport (MQTT) information from passport readers and other systems at the port – MQTT is the most widely used communications protocol used by "Internet of Things" technologies. The gateways might also communicate video streams encoded as high-volume Kafka feeds – Kafka is a modern, high-volume streaming protocol able to transmit video, audio, and many other kinds of data streams across the Internet.

Use Case: Logistics

Even the world's biggest, most heavily weaponed, most highly trained army is crippled without logistics. Soldiers must be transported to the field of battle and supplied with a ready supply of weapons, food, medicine and other necessities. Logistics are as important to the success of military missions as are soldiers, training, and weapons.

To this end, armies all over the world rely on their own trucks, aircraft, ships and other vehicles for moving supplies and personnel all over the world, and as a rule these armies rely on civilian logistics infrastructure as well. Shipments on civilian transport may be deliberately disguised to avoid drawing enemy attention or becoming targets of attack or sabotage. Civilian trucking, rail, and commercial air services can play a role in military operations, often using falsified, encoded, or generic labels and manifests to conceal their true purpose.

All this means that military planners in high-security and classified command and control centers often need access to shipment data from civilian logistics providers in low-security networks. Different civilian providers might make available information about shipments on websites or through Internet-based application programming interfaces (APIs). Military planners might therefore build custom applications to gather all of these data sources into a single repository – for example a SQL database on a low-security network.

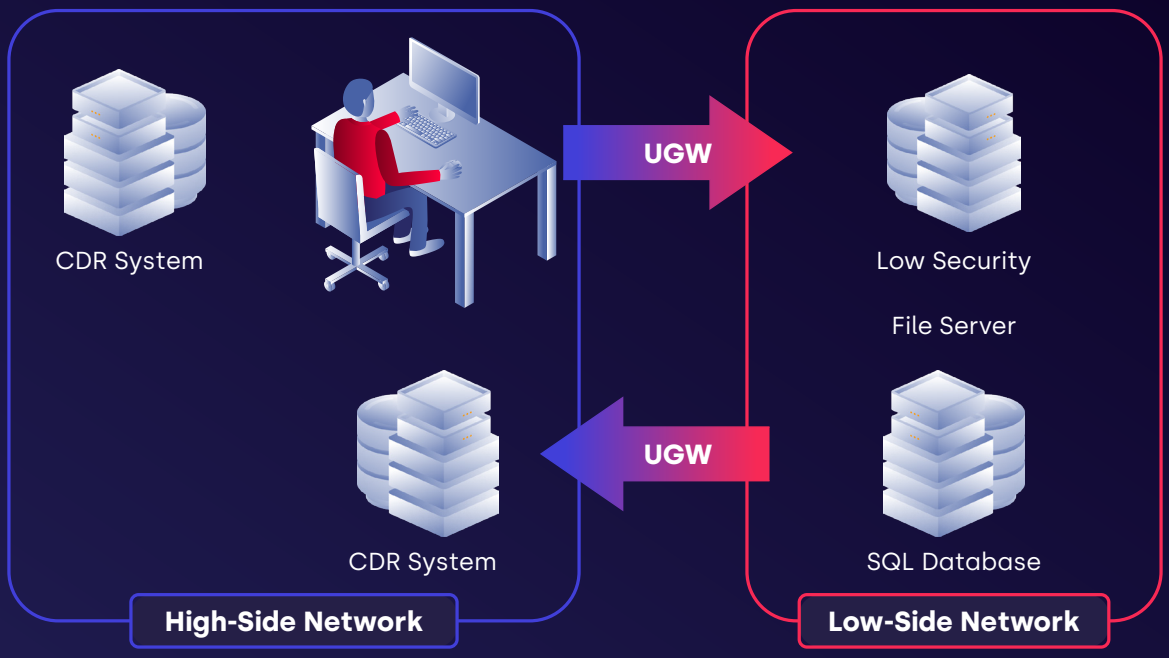

To make this information available to military planners in the command-and-control center, the military might deploy a unidirectional gateway, replicating the low-security database of anonymized shipments-in-progress data into the high-security planning network.

In addition, military planners and automation systems might very often need to request hundreds of new shipments per day in times of crisis. These requests might be generated automatically in the classified command center, but they would need to be communicated to civilian suppliers for those suppliers to act on the requests. This is an example of data LEAVING a high-security network destined for a low-security network. To accomplish this with minimum risk of leaking sensitive information out to the low-security network, a command-and-control center might configure their automated logistics system to produce civilian logistics requests as small XML files.

A CDR and file filtering system inspects the shipment requests, assuring that they are small text-only XML that meet a specific XML schema. The CDR system might then be physically wired into a human inspection station where the XML text of each shipment request is brought up and inspected visually by a human aide to ensure that the request contains nothing that seems in the least confidential or revealing of plans or motives. If the aide approves a request, the inspection workstation is also wired to a unidirectional gateway, which sends that request out to the low-security network, for encoding into the web pages or APIs of the various civilian logistics providers for execution and payment.

Again, when sending information out of a high-security domain, we must take special care to ensure that we are not unintentionally, or under the influence of an adversary, disclosing sensitive information into the low-security domain.

Waterfall's Unidirectional Security Gateways

Waterfall Security provides the world's most powerful, most widely used unidirectional gateway solutions for cross-domain systems. Some of the many advantages of Waterfall's solutions include:

- Waterfall's unidirectional hardware is optically isolated, certified unidirectional to Common Criteria EAL4+ (HIGH_ATTACK_POTENTIAL).
- Waterfall's unidirectional connector library is the most diverse, most feature-full in the world.
- Standard features and options for Waterfall products include high availability (no single point of failure), 10Gbps throughput, detailed selection criteria configuration for extracting known-good content from data sources, optional metadata synchronization, and a simple-to-use web-based GUI for configuration, management, monitoring and even troubleshooting. No command-lines needed.
- Waterfall products are commercial-off-the-shelf (COTS). No custom engineering is needed for the products, and no software or hardware needs to be installed on source or destination networks.

These and many other reasons are why Waterfall's cross-domain unidirectional and other solutions are deployed and proven in classified and sensitive military, intelligence and government networks and cross-domain deployments all over the world.

In addition, Waterfall Security is partnered with and deployed routinely with RESEC – the most advanced "known good" CDR system in the world. With RESEC and Waterfall's local service and support providers, you can be sure of secure, powerful and simple cross-domain integration, in even the most demanding of circumstances.

Wrapping It Up: Securing Our Future

Nation-state, ransomware criminal groups and other adversaries only become more capable over time, as they gain access to or develop more and more powerful attack tools and techniques. Hardware-enforced physical protection for sensitive data and networks is in a real sense future-proof. It does not matter how many passwords our enemies phish nor vulnerabilities our enemies develop exploits for if our defenses are physically incapable of leaking sensitive data. In many government, intelligence, and even military missions, it is vital to the success of the missions that we prevent theft of sensitive data and prevent disruption of important networks.

Many vendors provide CDS solutions of varying cost, capability and security. In this document we have surveyed the most important parts of the CDS landscape, focused on concepts and technologies that help us to select the most appropriate solutions for our unique needs. In particular, understanding what capabilities different vendors offer "out of the box" in COTS offerings is important to minimizing the cost of CDS solutions and avoiding costly, time-consuming and error-prone custom engineering.

Waterfall Security's Unidirectional Gateway and related technologies, partner technologies and partner services are leaders in cross-domain solutions. Waterfall boasts the world's most comprehensive, most sophisticated and most powerful CDS solutions – off the shelf.
For more information on Waterfall's hardware-enforced technologies and capabilities, or for a free consultation with one of Waterfall's cross-domain experts, please contact us at info@waterfall-security.com.

Talk to a cross-domain security expert today about your specific security needs:
info@waterfall-security.com

Glossary

Access CDS – Keyboard-video-mouse systems deployed to provide access to low-sensitivity networks and assets for high-sensitivity users in high-sensitivity areas.

Access Control List (ACL)) – A list of which users or groups of users have what kinds of permissions

Advanced Content Disarm and Reconstruction – Content Disarm and Reconstruction technology that extracts known good content from files and rebuilds mostly-equivalent files for transmission into higher-security domains, rather than converting files into images.

Air Gap – A confusing term that can mean any of:
- No online connection to a system or network,
- No way to route a packet to the Internet from a system or network
- A high-security domain connected to a lower-security domain through a CDS

Anti-Malware – Technology to detect malware embedded in files and other content crossing from low-security to high-security domains.

Content Filters – Systems that carry out basic inspection of files crossing domain boundaries, looking at easily-detected characteristics such as file size, names and types.

Content Disarm and Reconstruction (CDR) – Technology that extracts known good information from complex documents and rebuilds mostly-equivalent documents consisting only of the known-good content, leaving all other unidentified and potentially malicious content behind.

Covert Channel – Any means of signalling information from one place or system to another that is not a conventional communications mechanism.

Cross-Domain System (CDS or sometimes XDS) – Technology and processes at domain boundaries that avoid both exfiltration and accidental disclosure of sensitive data from high-trust domains, prevent infiltration of those domains with malware, and prevent remote control of malware within sensitive networks from adversaries outside those networks.

Cross-Domain Workflow and Orchestration – Systems that route files through a variety of manual and automatic CDS technologies, inspections and other activities, depending on the nature of the files, of the senders of the files, and other characteristics and rules.

Data Diode – Hardware that is physically able to send information in only one direction – usually into a high-security domain from a less-secure one.

Data Loss Prevention – A technology that labels files or data as to their sensitivity, controls and tracks who accesses the files, and raises alarms when anomalous access attempts or patterns are detected.

Domain – A system, network, or network of networks of connected systems, all at the same level of security is called a domain.

Domain Boundary – Any connection between networks or systems at two different security levels. Exfiltration – Theft of information from a protected network.

Firewall – A router that inspects and filters Internet Protocol (IP) traffic between systems and networks.

Mandatory Access Control – Technology that enforces access control and other policies over and above access control lists. Even if an ACL says a user has permission to do something, the user is forbidden unless the MAC system also says they have permission.

Multi-Level CDS – Databases or other repositories of documents or data that can be at a variety of levels of sensitivity.

Unidirectional Gateway – Hardware that is able to send information in only one direction (see "data diode") coupled with software that makes copies of databases or other systems for ease of integration and "known good" data extraction.