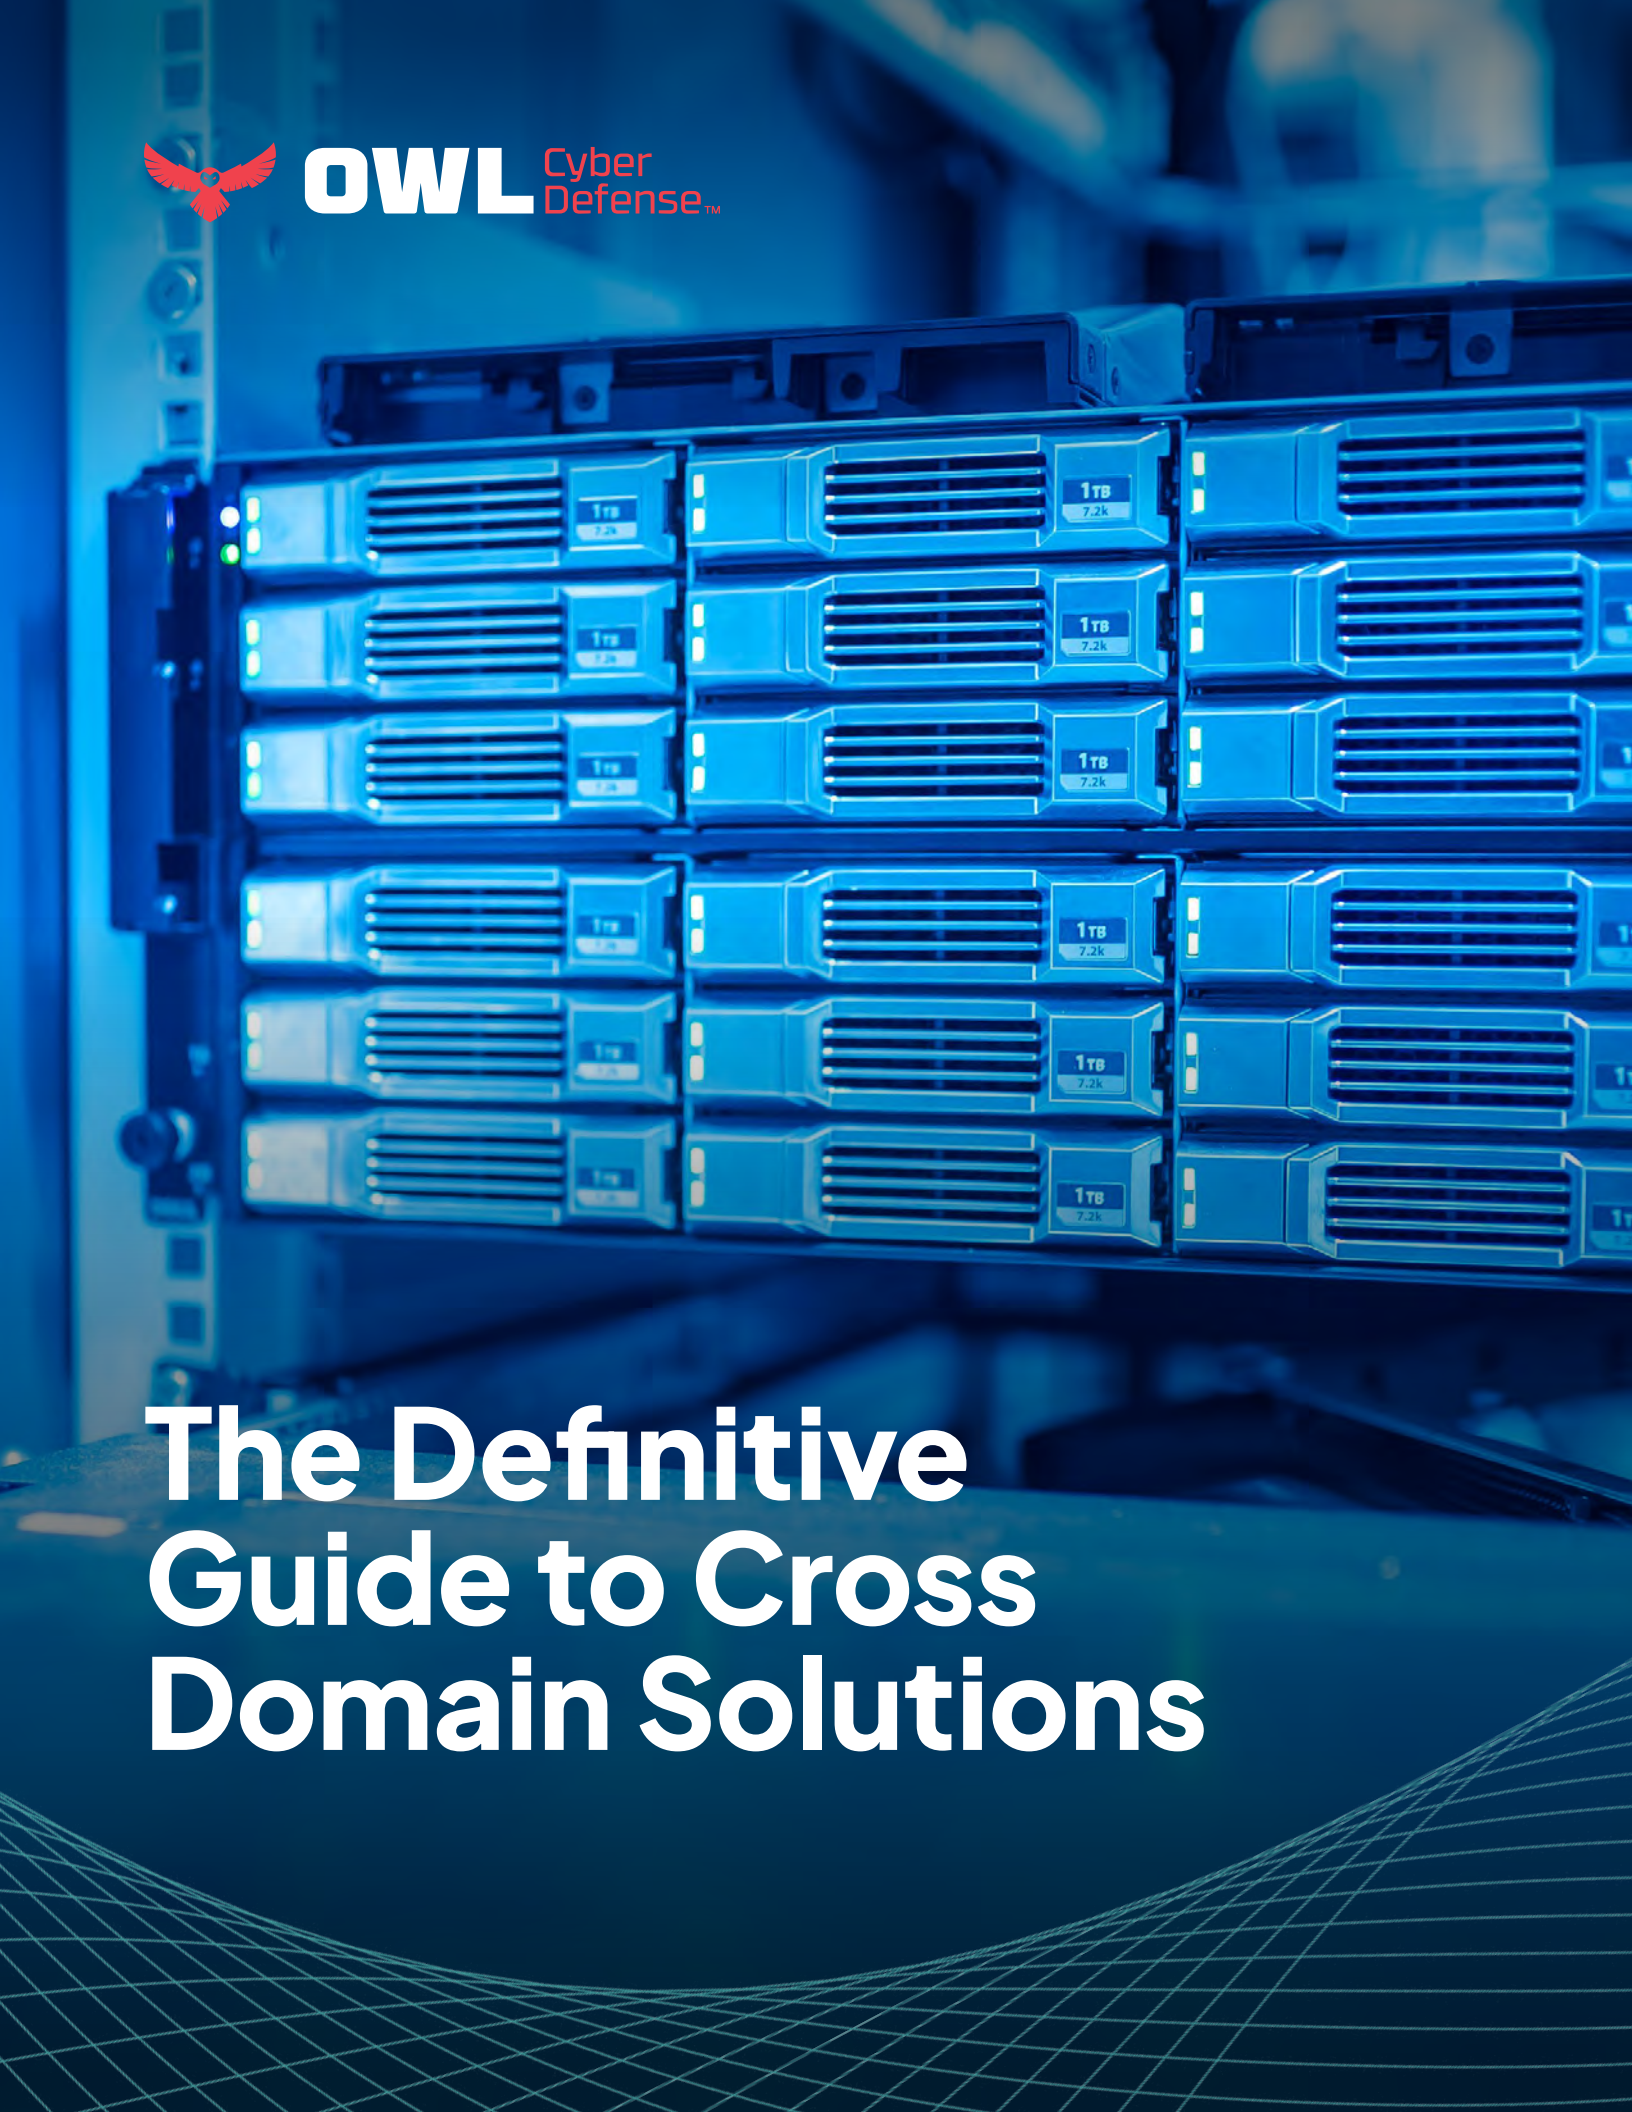

OWL Cyber Defense™

The Definitive Guide to Cross Domain Solutions

Table of Contents

5 Introduction

7 What is a CDS?

9 Background & Principles of a CDS

17 What Makes a “Cross Domain Solution”?

25 ‘Raise the Bar’ (RTB)

27 Common CDS Technologies

35 Security Principles of Cross Domain Solutions

43 How are Cross Domain Solutions Used?

47 Summary

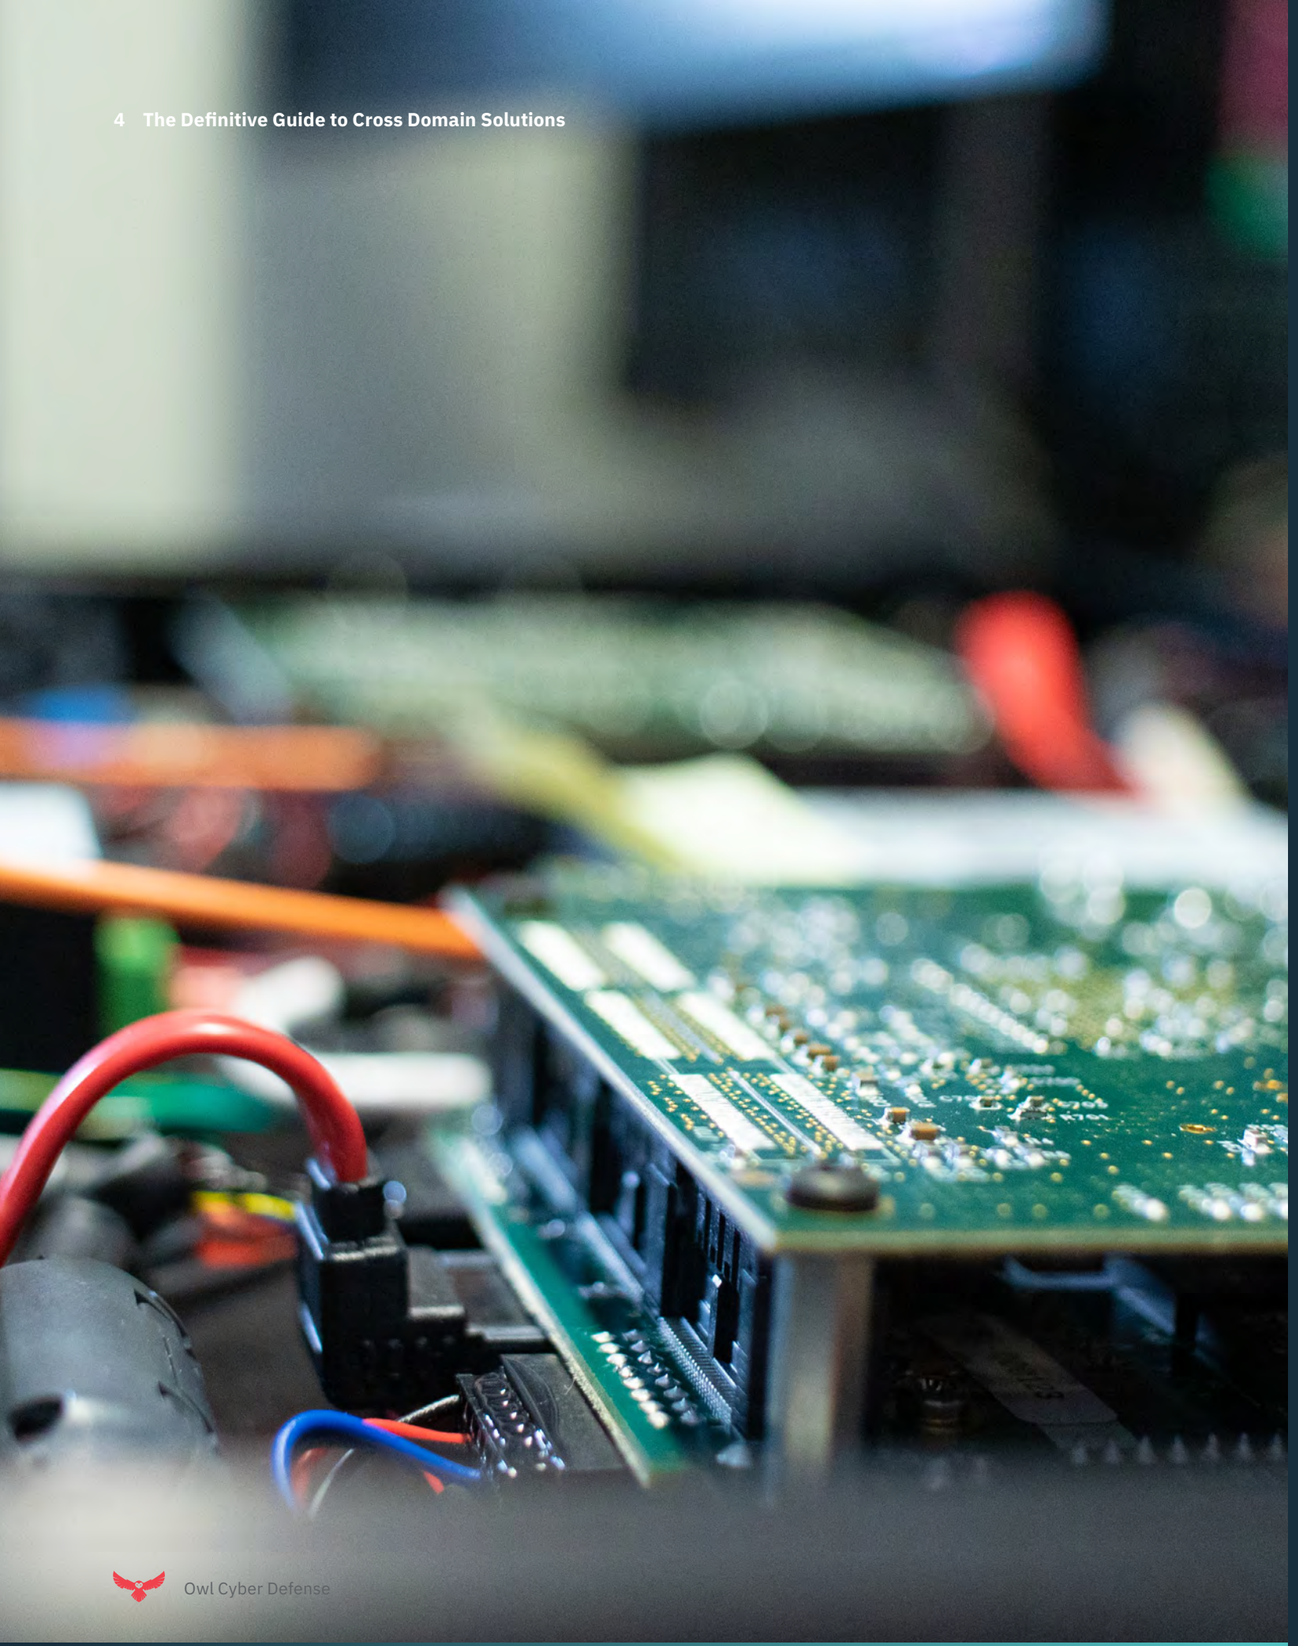

Introduction

Network security is not about technology – it's about trust.

Trust in your data, network, applications, and systems is essential, and keeping malicious infections and other bad actors from corrupting these trusted domains is the heart of network security. Organizations must be able to trust the enforcement of their information sharing policies such that they can be assured that only authorized data is provided to less trusted environments and that only clean or authorized data is consumed from less trusted or untrusted environments.

Imagine a bank's financial records have been manipulated through a malware infection, an entire hospital system's medical records encrypted with ransomware, or a government's intelligence data has been stolen by an adversary. Without trust, you are no longer able to rely or act on your data, eroding the bedrock of your organization until the entire system crumbles.

However, information sharing with external partners and environments is fundamental to the success of business and mission operations. Very few networks can truly operate disconnected from the outside world. Meanwhile, the connectivity required for information sharing creates attack vectors for infiltration of malicious data or exfiltration of sensitive data – the more connections you make, the more potential threat vectors. So how do you maintain the integrity of your trusted domain while also maintaining connectivity with untrusted domains?

Most organizations protect their infrastructure through insurance-based security in implementing industry-acceptable, "good enough" due diligence by constraining information flows. However, "good enough" is not acceptable for organizations that operate in high risk environments (Defense, Intelligence, Financial, Medical) that must have assurance-based security to ensure just not secure connectivity, but more importantly, secure information sharing. **Cross domain solutions** (CDSs) were developed specifically to provide this level of assurance to the networks of the U.S. Government, intelligence community (IC), and defense branches.

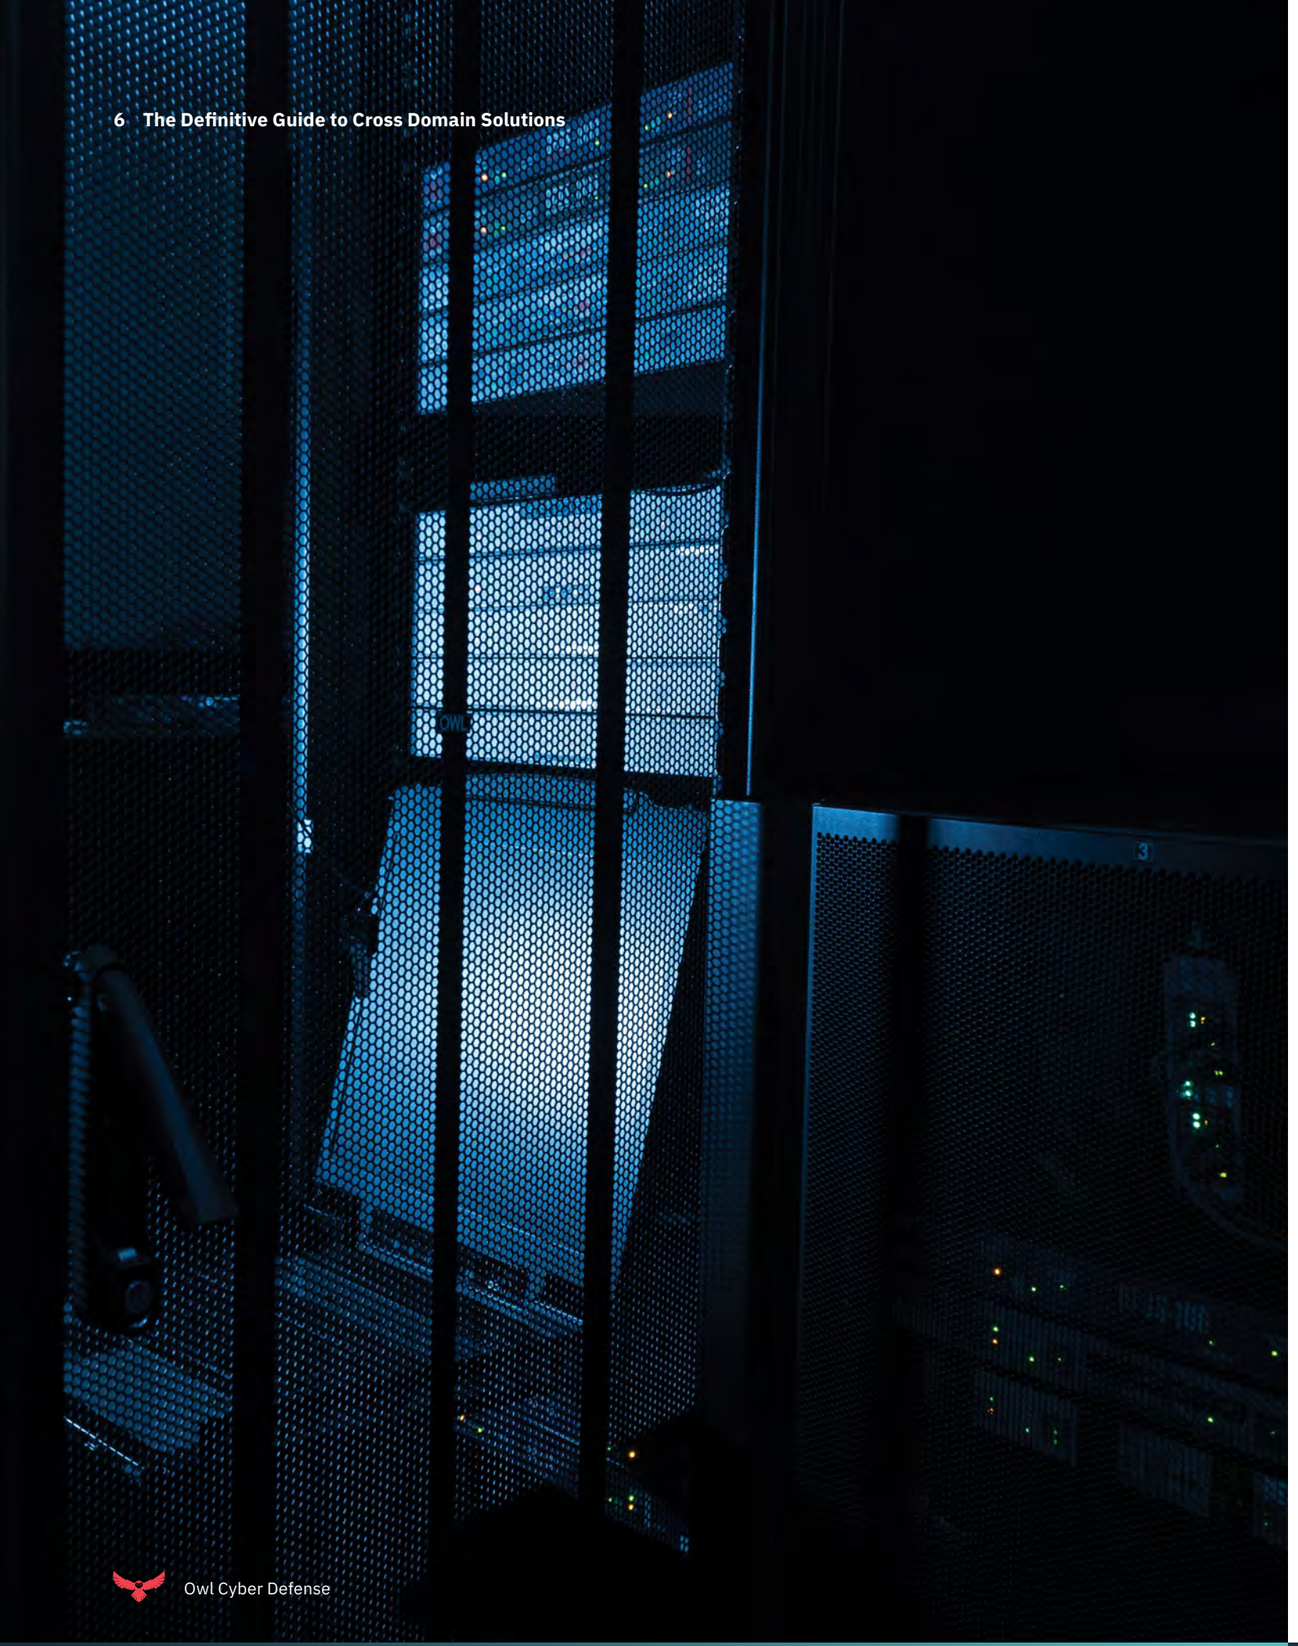

What is a CDS?

While definitions may vary between U.S. and international usage, the National Institute of Standards and Technology (NIST) defines cross domain solutions as:

> **"A form of controlled interface (a boundary with a set of mechanisms that enforces the security policies and controls the flow of information between interconnected information systems) that provides the ability to manually and/or automatically access and/or transfer information between different security domains."** [1]

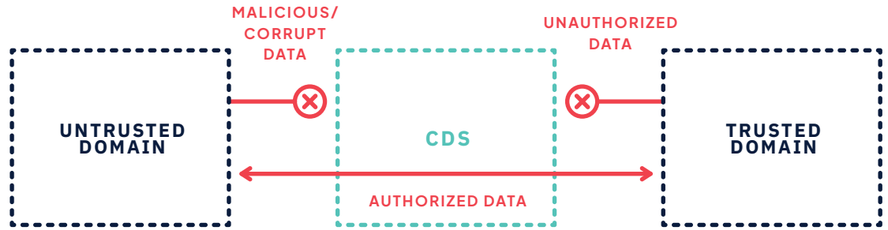

Indeed, while they are strictly defined within the context of U.S. Government regulations, cross domain solutions can include a variety of functionality and feature sets which both set them apart from other secure data transfer and boundary protection systems and each other. There are two functional classes of CDSs: transfer CDSs and access CDSs. This document will focus on transfer cross domain solutions.

Essentially, a transfer CDS is an information security system which allows or denies data transfer by enforcing an organization's security policies. CDSs are designed to control and restrict the flow of information both to and from trusted and untrusted domains ("high to low" or "low to high"). While often used in a unidirectional capacity, CDSs may have one-way-only, two-way, bidirectional, or multidomain connectivity. Some CDSs implement only software enforcement, while others provide additional enforcement through hardware.

The following is intended to help guide you through the various types, technologies, benefits, uses, and misconceptions of transfer cross domain solutions. We hope you find it useful and educational.

1 https://csrc.nist.gov/glossary/term/cross_domain_solution

Background & Principles of a CDS

The subject of cross domain security is paramount, as the need for more secure information sharing between highly secure domains within highly sensitive entities (defense, intelligence, etc.) continues to escalate in parallel with an increasingly rich and sophisticated threat environment. This involves both the increase in efficiency for securely sharing information, as well as the corresponding increase in security assurance.

History of CDs

The concept of a trusted network security solution was devised as a part of a U.S. Department of Defense (DOD) and National Security Agency (NSA) initiative in the 1980's to define and build "Trusted Operating Systems" and "Trusted Networking" components. The outcome of that initiative was the **Trusted Computer System Evaluation Criteria**[1] (TCSEC), published in 1983, also known as the "Orange Book" in the DOD "Rainbow Series" of computer standards and best practices.

Among many other facets of network security, the TCSEC defined the Policy, Accountability, Assurance, and Documentation standards, as well as a classification system for network security assurance levels. This of course did not include all of the modern functionality and sophisticated software present in modern network security, but the refinement of the TCSEC criteria became the foundation for what we know today as cross domain solutions.

More recently, the TCSEC was superseded by the Common Criteria (CC) international standard, originally published in 1994, and established as the ISO/IEC 15408 standard in 1999.[2] However, the CC standard was never really carried forth (with respect to CDSs) in maintaining what is known as a Protection Profile or Security Target, outside of the trusted operating systems used within CDSs. As such, since the mid-2000's, the standard for cross domain solutions has been maintained substantially within the DOD and intelligence community (IC). With the publication of Committee on National Security Systems Instruction (CNSSI) 1253 - "Security Categorization and Control Selection for National Security Systems"[3] in 2012, specifically Appendix F, Attachment 3 (commonly referred to as the "CDS overlay"), the standards of CDS development were established more clearly for the first time since the early 2000's.

Today, the accreditation of cross domain solutions in the U.S. is primarily the responsibility of the National Cross Domain Strategy and Management Office (NCDSMO) within the NSA. There are still pockets of the IC that don't rely on NCDSMO guidance, though the prevailing trend is toward compliance with NCDSMO directives. Over the last decade, the base of solid network security

1 https://csrc.nist.gov/csrc/media/publications/conference-paper/1998/10/08/proceedings-of-the-21st-nissc-1998/documents/early-cs-papers/dod85.pdf
2 https://www.commoncriteriaportal.org/iccc/ICCC_arc/history.htm
3 http://www.cnss.gov/CNSS/issuances/Instructions.cfm

concepts has grown to include new data filtering standards, new trusted operating systems, additional layers of security requirements and redundancy, and the marriage of specialized network security hardware and software. The NCDSMO, in December 2018, published the first set of mandated controls and processes specifically for CDSs - a mandate is commonly referred to as "Raise the Bar (RTB)."

Security Domains

As cross domain solutions involve the security and transfer of data between domains (thus "cross domain"), it's important to understand what is meant by the word "domain" in the context of security. NIST defines a security domain as, "A domain that implements a security policy and is administered by a single authority."[4]

There are a few important elements involved in this definition, the first of which is that the domain itself has its own security policy. This policy may or may not be unique to that domain and may or may not represent a general security classification or a unique subset of that classification (e.g. TS/SCI). The final point to emphasize is that the security policy within the domain is administered independently of external domains by a single authority – isolated control within the trusted environment is crucial to ensure integrity.

Security domains may also be layered within each other, sorted hierarchically, or organized in other fashions, depending on the security requirements, in order to form a larger domain with a subdomain structure. In this way, individual segments of the larger domain may also have individual policies specific to their use and contents.

Boundary

The boundary of any domain or system is the logical and/or physical perimeter which is set around it. It encompasses all those components of the system that are to be accredited and excludes

4 https://csrc.nist.gov/glossary/term/security_domain

separately accredited systems to which the system is connected (e.g. via cross domain solution). The accreditation serves as an assurance of separation or segmentation from external domains, even if the domain resides within another domain.

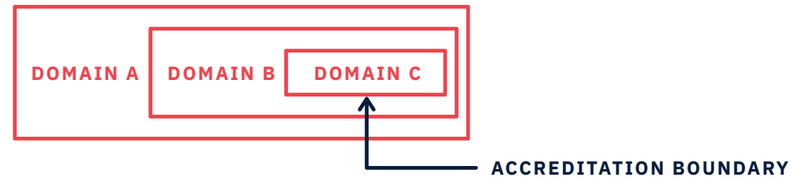

Enclave

Defined by NIST as "A set of system resources that operate in the same security domain and that share the protection of a single, common, continuous security perimeter,"[5] enclaves may be better understood as the lowest level of a security domain structure, wherein everything inside that domain shares the same security profile, policy, and administration. In this way, a security domain may be a collection of enclaves and/or subdomains, which in turn comprise other subdomains and/or enclaves.

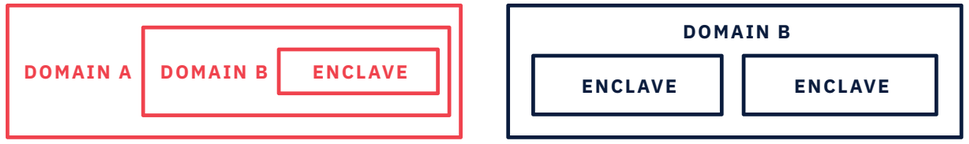

High/Low

Because the very definition of cross domain solutions comes out of the **Trusted Computer System Evaluation Criteria**[6] (US DOD "Orange Book"), "high" and "low" security are relative terms given to the respective trusted and untrusted environments in the cross-domain transfer. Data can be transferred in both directions, to or from the trusted network (or bidirectionally), and security policy within the CDS should be modified accordingly.

5 https://csrc.nist.gov/glossary/term/enclave
6 https://en.wikipedia.org/wiki/Trusted_Computer_System_Evaluation_Criteria

In general, the "high" security domain is the one with a more stringent security policy where the CDS is installed, while the "low" domain has a less strict security policy, although it is not always that simple. The receiving domain may require specific security controls that the sending domain does not, although there may be other controls on the sending domain that do not apply to the receiving domain. Classified networks can simplify this a bit, with the "high" network of a higher security classification than the "low" network, although when transferring data between networks of equal clearance level, this again becomes somewhat less clear cut.

The simplest, most definitive terms would be to consider the trusted domain (yours) the "high" side, while the untrusted (external) domain is the "low" side.

Air Gap

While many would consider an "air gap" to be a complete physical and logical separation, the definition within the CDS context is somewhat different. An air gap is an interface between two domains or systems at which:

(a) they are not connected physically and

(b) any logical connection is not automated.

Zones (DMZ/SOC/NOC)

Network zones are typically demarcations of groups of connected, intercommunicating systems and/or domains, separated by security perimeters. In the case of network DMZs or security operations centers (SOC) or network operations centers (NOC), these "zones" act as buffers between trusted and untrusted networks, providing an additional layer of security for data communications. DMZs are a mechanism to provide additional functionality to the connected domain while limiting exposure to the internal systems. NOCs & SOCs are specialized network zones designed to manage and monitor network assets.

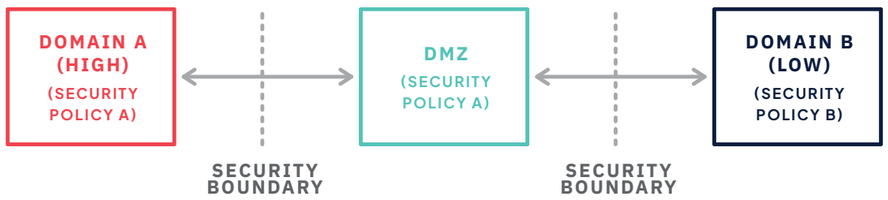

Confidentiality, Integrity, and Availability (CIA)

As the core of all information security, the "CIA triad" of confidentiality, integrity, and availability have been in place since the 1970's. These core concepts encompass the basic building blocks of any secure system and are key to enabling trust within network domains.

- **Confidentiality:** the assurance that information is disclosed only to authorized entities
- **Integrity:** the assurance that information and systems are not modified without authorization
- **Availability:** the assurance that systems are accessible and useable by authorized entities when required[7]

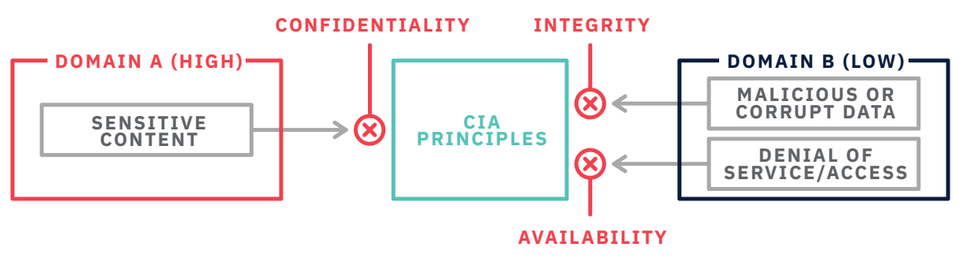

Certified cross domain solutions are high assurance tools designed to implement the CIA policy for data transfers between domains of differing trust levels. Data transfers may be trusted to untrusted (High-to-Low), untrusted to trusted (Low-to-High), or bidirectional, and may not necessarily include a difference in classification (e.g. Secret to Unclassified).

Like a firewall, a CDS is used to protect the integrity of the systems and prevent unauthorized access, while ensuring data availability to authorized parties. The integrity and confidentiality of data must also be extremely reliable, as data sharing may be mission critical.

However, firewalls are software-only solutions designed primarily as "all-or-nothing" security tools to protect against unauthorized access and known threats/malware. While some "next-generation" firewalls offer some forms of data loss prevention and content inspection, they lack the hardware-enforcement, defense-in-depth architecture, trusted software, and sophisticated content filtering of a CDS.

Cross domain solutions implement content filtering (of various types) according to the strictly defined security policy of the trusted domain. These filters are tasked with both preventing confidential data

7 https://www.cyber.gov.au/publications/fundamentals-of-cross-domain-solutions

from leaving the trusted domain and preventing malicious or unauthorized data from entering the trusted domain.

Authenticity & Non–Repudiation

In addition to the classic CIA triad, cross domain solutions must also include enforcement elements for authenticity and non-repudiation (AKA accountability).

Authenticity: the assurance that the identity of a user, process or device is known and verified, as a prerequisite to allowing authorized access to resources in a system

Non-repudiation: 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[8]

Authenticity is now a fairly standard feature for any network security system, often requiring multifactor (e.g., password, token, time sensitive code, etc.) or biometrics (e.g., fingerprint, face scan, etc.) to validate identity.

Non-repudiation encompasses elements of both accountability and provenance to ensure both that the sender and receiver of a data transfer can be verified, and once verified that a completed data transfer between them cannot be subsequently denied or altered by either party. This means that the verified receiving party of a completed data transfer cannot subsequently deny having received it. Non-repudiation also includes the prevention of replay by the sending party (duplicating authorized data and submitting it twice), such as resubmitting an authorized purchase or account transfer, or re-firing a missile.

"Insurance vs. Assurance"

Most network security technologies are insurance measures that are rooted in the idea that they will defend to a certain extent, but that there is an acceptable level of risk that threats can (and will) eventually find their way through. Federal government, defense, and intelligence organizations cannot afford the same levels of risk. CDSs are based on the fundamental assurance that the solution has been evaluated and hardened such that available avenues of exploitation are reduced as near to non-existence as possible, and that there is a potential threat on both sides of the device, in infiltration and exfiltration.

Do I Actually Need a CDS?

Before going into more detail about the various technologies and features of cross domain solutions, you may be wondering if you actually need a CDS in the first place. While the following sections may provide more disambiguating detail on the many other types of network security technologies

8 https://www.cyber.gov.au/publications/fundamentals-of-cross-domain-solutions

available today, such as next-generation firewalls (NGFW), data diodes, or one-way transfer (OWT) systems, there are two fundamental questions you should ask yourself first:

1. Are you transferring data between domains that do not trust each other (two - or more - parties)?

As previously discussed, "trust" in this context has little to do with the known properties of the parties from or to whom you may be transferring data, but rather the ability to verify and validate those domains and data. If you cannot, that domain is typically rendered untrusted for the purposes of high-assurance network security, and everything that is transferred to or from it is subject to security controls.

How do the applications communicate (email, etc.)?

2. Do you require content filtering?

One of the primary methods of controlling information flow is access control. However, access only goes so far. There may be hidden threats within connections to authorized parties, and there may be restricted data that could leak to unauthorized parties. In order to detect and eliminate these threats, cross domain solutions utilize advanced content filtering to ensure just not secure connectivity, but more importantly, secure information sharing without risk of data exposure, systems infiltration, or data leakage.

What kind of data do you need to transfer?

What Makes a "Cross Domain Solution"?

Technologically, cross domain solutions are a unique combination of hardened and trusted software and/or hardware components, architected to enforce a meticulously designed security policy on data transferred to or from sensitive trusted domains. As the epitome of network security assurance, every facet of a CDS is designed to protect the trusted domain and the system itself with the strongest available tools and technologies, layered to create a sophisticated matrix of security controls. These key elements are what elevates cross domain solutions above other network security solutions and places them into a class of their own.

Although the tools and technologies used within a CDS may be selected from and/or designed by trusted sources, the complete solution still must be rigorously tested to verify its security and operational assurance. Because they are intended for use within classified domains, true "cross domain solutions" within the U.S. government context are export restricted and not available for sale outside of verified or approved government and defense organizations within the "Five Eyes" (the United States, United Kingdom, Australia, Canada, and New Zealand).

How is a CDS Different from Other Network Security Solutions?

As a fairly unique and esoteric technology, cross domain solutions are not well known or understood by many as distinct from firewalls or other network security technologies. While not truly an "apples to apples" comparison, as CDSs are truly in a class of their own as regards their sophistication and level of security, one can take a side by side look at them amongst other solutions to more clearly define the differences.

Firewalls

Firewalls are software-enforced transfer technologies that provide a logical separation of network and application infrastructures, through network and protocol filtering. It is critical technology for reducing attack vectors between network and application infrastructures at the protocol level between those infrastructures and is thus often employed as a part of a transfer CDS. Firewalls are usually based on general purpose operating systems (e.g., Linux, Solaris) and are always located in trusted environments as they provide no domain separation capabilities or "protocol break" in the data transfer. One of the weaknesses in a firewall is that when it allows a service, e.g., FTP, it defaults to allowing all of the functions of this service to work, while a CDS can strongly restrict those functions.

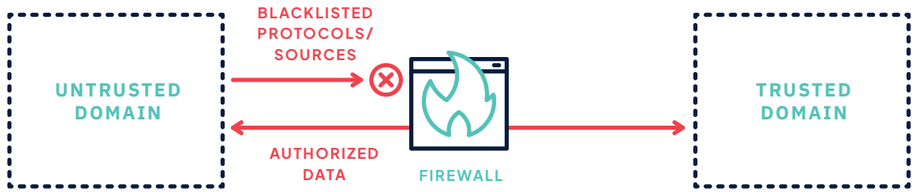

Data Diodes

Data diodes are hardware-enforced one-way data transfer technology (e.g., optocoupler, FPGA) that provides assured separation of network infrastructures. They are designed to transfer data one-way across an air gap, with a protocol termination to hide all network routing information from untrusted external users and systems. It is a critical technology for eliminating infiltration (in high-to-low transfer) or exfiltration (in low-to-high transfer) attack vectors at the network level between network infrastructures, and often used within transfer CDSs to provide a hardware-enforced security element and network separation.

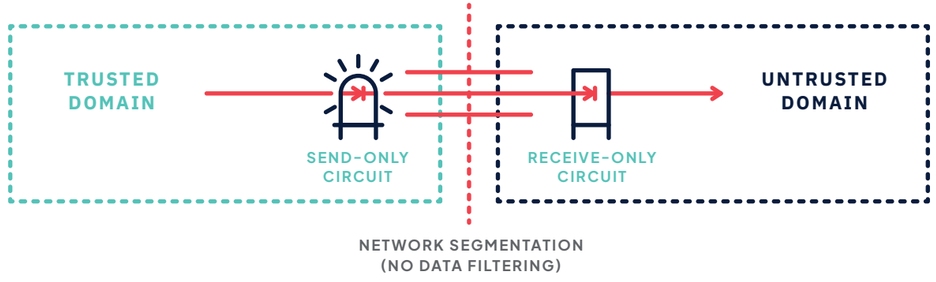

OWT Systems

One-Way Transfer (OWT) Systems are hardware-enforced one-way transfer technology bound with software-enforced transfer technology. The software-enforced technology includes a separate processing environment dedicated for each network infrastructure. Source infrastructure is the "Pitcher" processing environment and Destination infrastructure is the "Catcher" processing environment. These processing environments may exist in a single chassis with data diode technology connecting the two processing environments or in two separate servers with data diode technology as the connection between the servers. Those processing environments support multiple protocols and may provide some network, protocol, and content filtering capabilities and are usually based on general purpose operating systems (e.g., Linux, Windows).

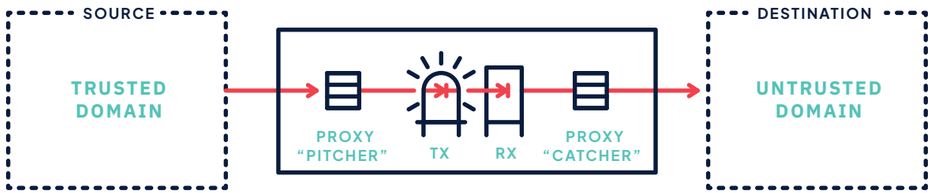

High-Assurance Software Guards

Software guards provide a content inspection layer and filtering similar to an access CDS without many of the other security mechanisms in a typical transfer CDS (e.g., a hardware-enforced transfer mechanism). Although they provide additional assurance well beyond a traditional firewall, they are essentially proxy servers capable of deep content inspection (DCI).

Access / Multilevel CDSs

Access CDS (desktop technology) and Multi-level CDS (usually Database technology) are able to concurrently host multiple domains but provide no transfer capability of content between those domains. These technologies are used primarily to converge assets of multiple domains on a single platform, making them available to authorized users without ever transferring sensitive data between domains.

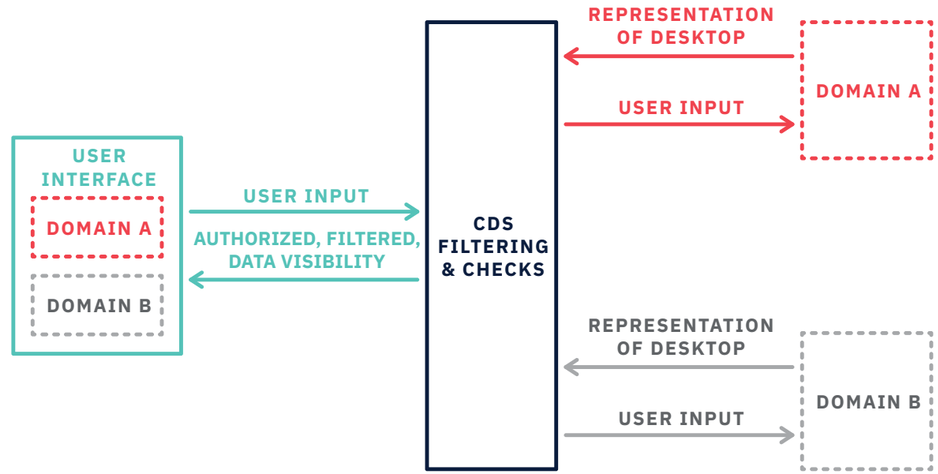

Transfer CDSs

Typically, the term “CDS” will refer to a transfer CDS – the subject of this document. As previously mentioned, a transfer CDS is software-enforced transfer technology that provides protocol, network, and content separation and filtering for network and application infrastructures. It is a critical technology for significantly reducing attack vectors of network, application, and application content infrastructures. CDS are always based on Trusted Operating Systems (e.g., Security-Enhanced Linux [SELinux], Solaris with Trusted Extensions). CDSs concurrently support the separation of multiple domains.

Transfer CDSs are used to provide the state-of-the-industry security technology controls for transferring content between two or more domains (e.g., Company/Internet, TS/S, S/U, Company 1/Company 2, Coalitions, Consortiums) where there must be the highest degree of assurance that only authorized content is collected, distributed, or disseminated between authorized sources and destinations within those domains. There are a few different variations of transfer CDS, but they all follow the same basic principles.

One-Way

One-way CDSs transfer data in only one direction, either into or out of a trusted domain (Low-to-High or High-to-Low). This type of CDS typically operates with the hardware-enforcement of a data diode to assure no possibility of transfer in the other direction. CDSs that do not feature a hardware-enforced element will require an add-on hardware-enforced OWT element in order to connect to high-threat networks (HTN), e.g. the Internet, and to prevent data leakage, infiltration, or other exploitation of the return path. This type of CDS is common in data collection and monitoring applications as they rarely require return acknowledgments or corrective measures.

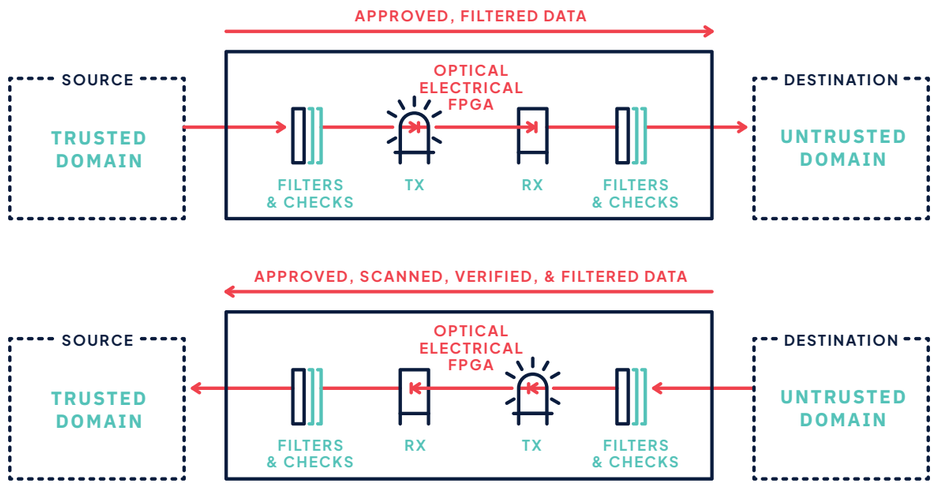

Two-Way

Two-way CDSs serve a very similar functional purpose to one-way CDSs except that they include two separate pathways, each to transfer data in only one direction - one in and one out of a trusted domain, rather than one or the other. This is commonly achieved by employing two data diodes within the CDS – one for each data path – with no overlap in connectivity or logical functions.

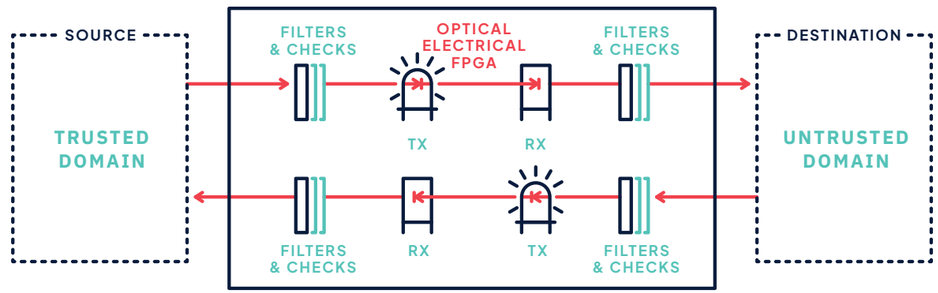

Bidirectional

A bidirectional CDS enables a round-trip pathway for application level query/response and acknowledgment communications on a common platform. As such, bidirectional CDSs must feature both High-to-Low and Low-to-High security policies and enforce each on the appropriate data path. This also necessitates physically separate High-to-Low and Low-to-High paths to enforce separation of data flows while enabling a single round trip path (Low-to-High-to-Low or High-to-Low-to-High) via paired proxies which coordinate the two one-way paths into a bidirectional data flow.

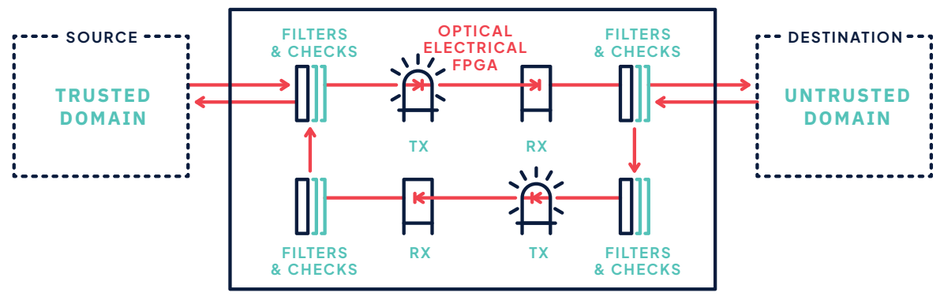

Multi-Domain

Recent advances in technology have enabled the ability to connect to multiple domains through a single transfer CDS. In fact, some vendors have promised the ability to connect over 20 domains. In practice, however, the ability to accommodate more than a handful of domains quickly becomes extremely complicated. Requirements such as hardware-enforced separation in any HTN application and for each connection to utilize a physically separate network interface controller (NIC) make the prospect of connecting double digit domains within a single set of enterprise servers impractical. All security is provided only by software enforcement and not augmented by hardware enforcement.

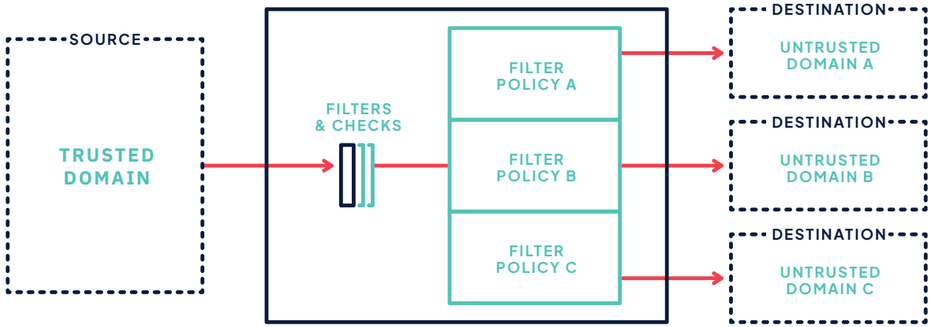

Hybrid CDSs

A hybrid CDS is a CDS composed of multiple elements to create a complete solution. For example, transfer CDS software-enforced technology combined with hardware-enforced one-way transfer technology (OWT), or a transfer CDS combined with an access CDS.

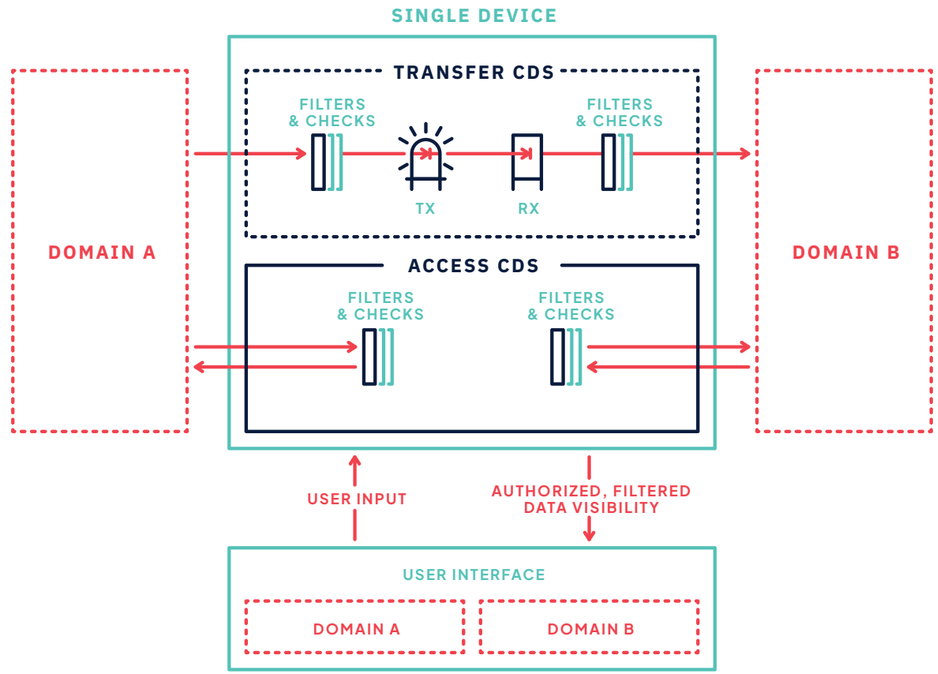

Exportable CDSs

There are now network security solutions that employ similar or analogous technologies, principles, and methodologies but are not built on the same code base as cross domain solutions designed for use in U.S. government/defense applications. The purpose of these devices is to provide the assurance of a CDS to the non-export-restricted marketplace, such as international governments outside of the Five Eyes (Australia, Canada, New Zealand, United Kingdom, and United States) and critical infrastructure.

‘Raise the Bar’ (RTB)

The 2018 publication of the National Cross Domain Strategy and Management Office (NCDSMO) **“Raise the Bar”** (RTB) mandate is causing a positive transformation of the cross domain solution (CDS) market. However, like any increase in government regulation on an industry, there will be some impacts to customers that will be seen in cost, performance, and operational procedures, as well as the loss of some vendor participation in the CDS market.

The RTB mandate’s focus on improving the security of CDS data flow and content filtering controls has produced new, more complex CDS architectures. The overall result in these architectures is much more rigorous processes and greater assurance for system and data owners concerning the protection of their mission and enterprise dataflows that are transferring content.

One aspect of the increased assurance is the recognition that hardware-enforced one-way transfer mechanisms are an essential component for system and data owners that need cross domain transfer of information with either the source or destination being a less trusted (e.g., Unclassified, High Threat, Internet) network.

Hardware-enforced one-way transfer mechanisms prevent any and all aspects of back channel communications originating from the receiving network, including infiltration or exfiltration. System and data owners must employ hardware-enforced one-way transfer mechanisms for any touchpoint that they deem has any substantive degree of risk to ensure their infrastructures are protected from unauthorized connectivity.

It’s also important to note that the RTB standard is a carefully managed, evolving set of requirements and its baseline is continually being revised with additional guidance and requirements. These revisions will include “phases” of compliance and ongoing refinements to current technologies.

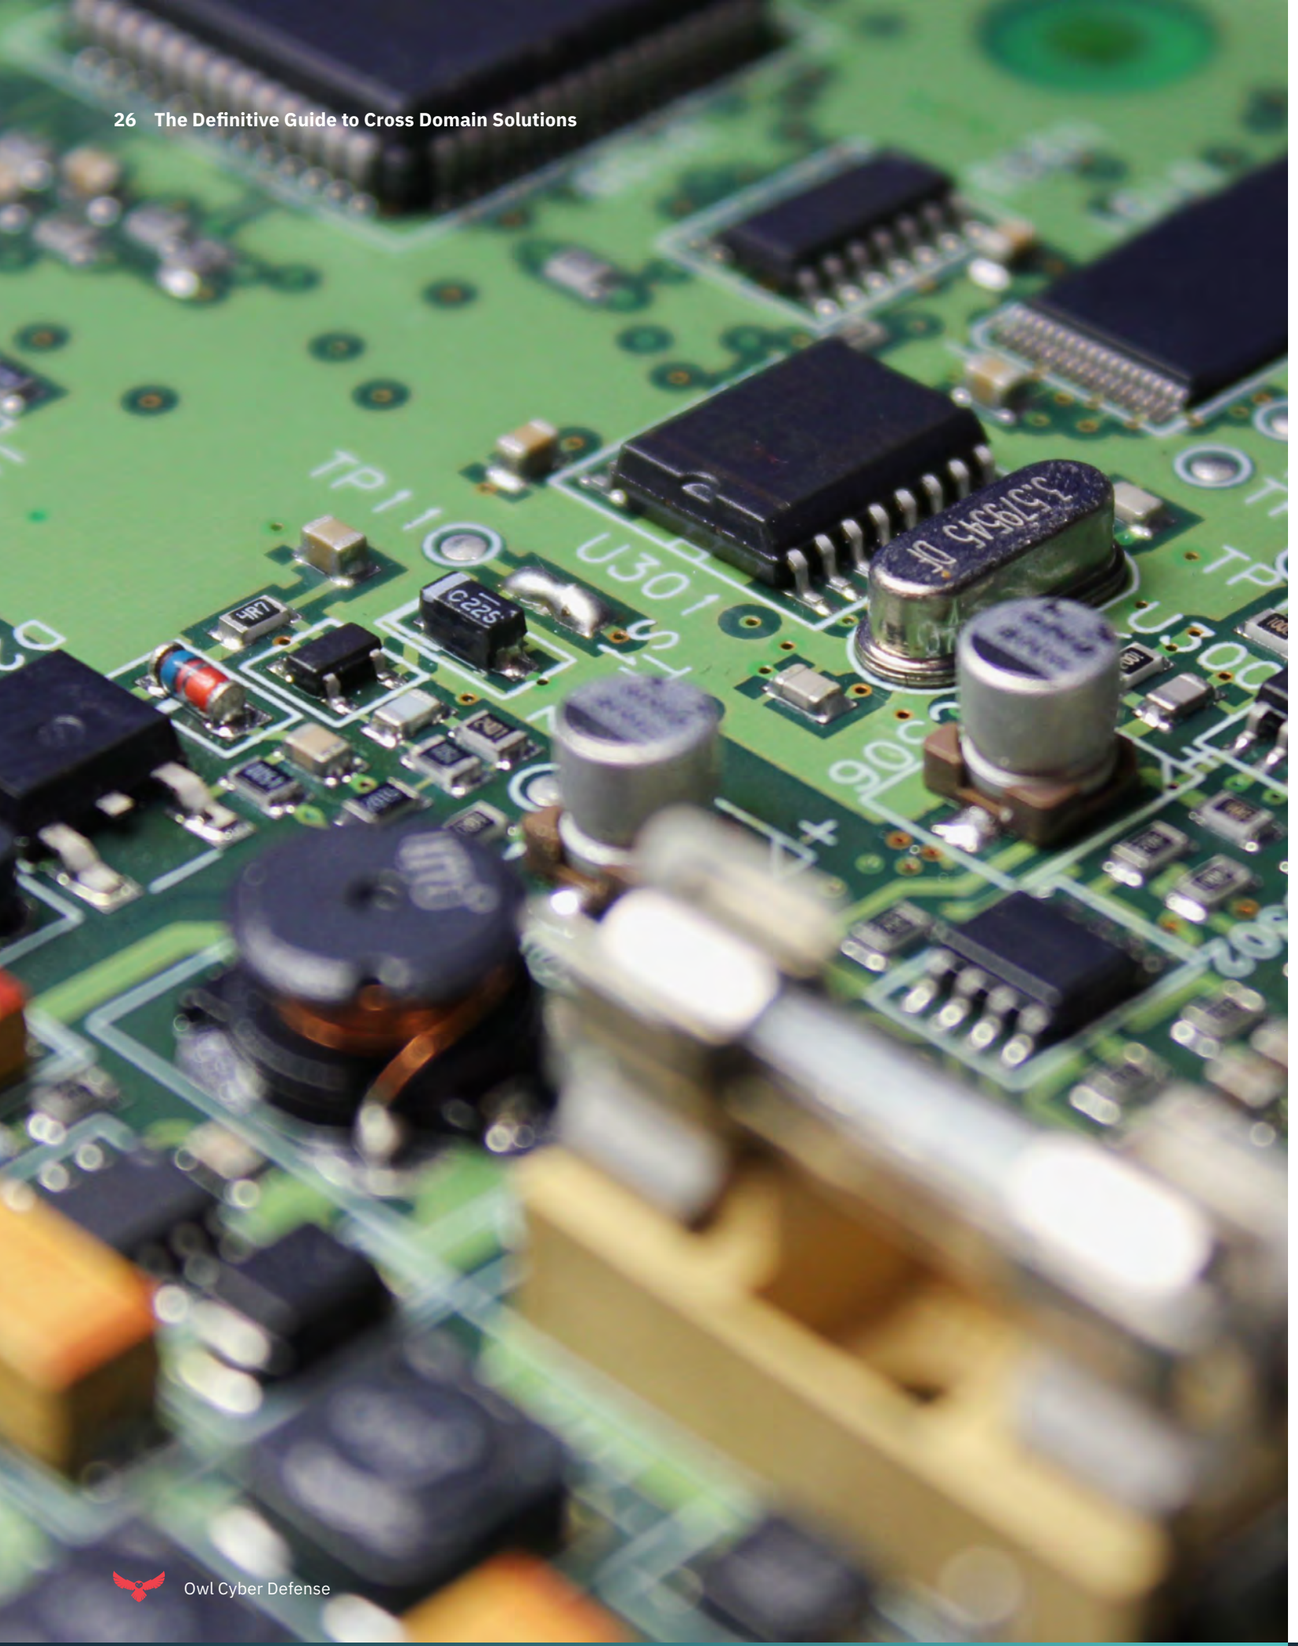

Common CDS Technologies

Cross domain solutions by their nature incorporate a number of security and functional technologies. These include firewalls, anti-virus, data diodes, data loss prevention (DLP), and role-based access controls (RBAC), among others. Although many of these technologies are commonly in use independently or within other systems in a network, in those cases their use within a CDS provides a redundancy for added security. In addition to these common technologies, CDSs employ a number of other unique software and hardware features.

Not all CDSs are created the same, however. Some types of CDSs are natively integrated hardware and software enforcement, while others are manual combinations of separate software-only and hardware-enforced solutions. This is not to confuse components for their role in the CDS – for example, a commodity server, while a hardware element, provides no hardware enforcement of security or policy. For the purposes of this paper, there is no distinction between the various types of CDSs in terms of the technologies they employ, however the assumption is regardless of their composition, they must include a baseline of capabilities. One should also keep in mind that integrated solutions are almost always more secure in terms of seamless operation.

CDS Software

From the hardened software platform to data filtering and numerous integrated applications, software is the brain of a CDS. Though the numerous, redundant software systems and applications used in a layered, overlapping architecture are what gives them collective strength, each individual technology should still be best in class in any accredited CDS.

Content Filtering

CDSs combine content/data filtering at the application level with protocol level filtering for individual data ports for a multifaceted toolbox of capabilities. These capabilities may be a part of a software-only solution with an add-on hardware enforcement or one-way transfer (OWT) solution.

Content filtering software generally falls into three categories: Fixed Format, Streaming, and Complex.

Both structured and unstructured content filtering require the use of inspection, sanitization and quarantine functions to prevent the transfer of unauthorized or malicious content.

- **Inspection**

Verify and validate that transferred content complies with defined policy and format standards. The CDS will require decryption of any encrypted content for inspection. If in accordance, transfer content to destination or quarantine content (based on additional policies).

- **Sanitization**

Redact portions of transferred content that does not comply with defined policy and format standards. This is a point of distinction for CDSs among other types of network security, as it does not require a complete quarantine of a file or document. If in accordance, can transfer redacted content to destination or quarantine content (based on additional policies).

- **Quarantine**

Content that does not comply with defined policy and format standards and is quarantined and not transferred to destination. Can either be kept on secured memory within CDS or transferred to an organization's forensics environment (e.g., Defensive Cyber Operations [DCO]) for further analysis and processing.

Structured Content Filtering

Structured content or "fixed format" messages, including text or schema-based XML, is filtered using a process known as **linear pipeline.** This straightforward approach applies a series of filters and checks in order, each separated into isolated, independent tasks, with handoffs from one to the next.

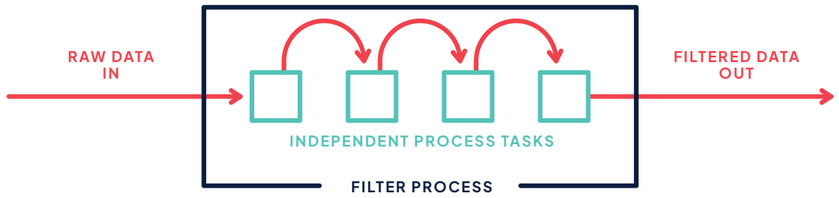

Complex Content Filtering

Complex content or unstructured data, e.g., MS Office, PDF, or imagery, is broken down into more basic elements and filtered using a process known as **recursive decomposition.** This divide and conquer approach involves normalizing or decomposing data so that it can be inspected using standard content filters. In some cases, custom filters may still be required for specific data formats.

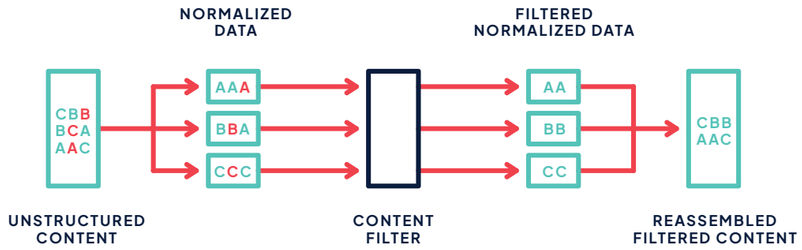

Streaming Content Filtering

Filtration of streaming data types, such as UDP, video, or sensor data, typically involves **parsing streams** to restrict identifying metadata or remove corrupted or unnecessary data from the stream. As the filtration is typically needed at full "line rate," the parameters are typically static with options to turn specific filters on or off depending on mission requirements.

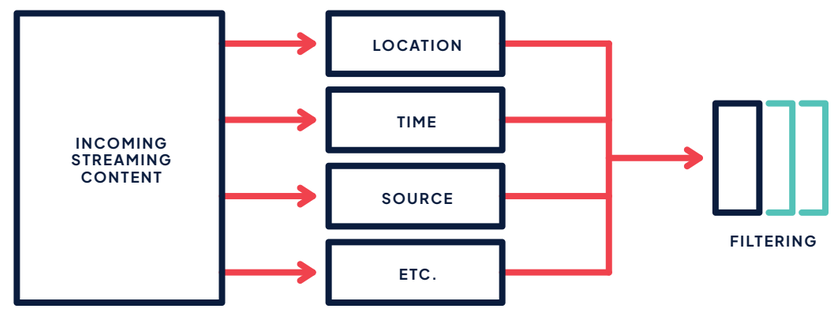

Trusted Operating Systems

A trusted operating system (TOS) is an operating system that provides layered, multilevel security capabilities, and the assurance to be "trusted" to perform according to specifications and policy.

In accredited CDSs, TOS capabilities must be sufficient to meet the requirements of the NIAP Operating System Protection Profile (OSPP) within the NSA and the Common Criteria standards. This includes supporting certain OS capabilities, including mandatory access control (MAC). "Mandatory" in this case means that access control policy is set outside the control of the user or object owner. A prime example of MAC is in the application and enforcement of data classification within the IC and military, wherein all objects are classified by a central authority. The data owner does not decide whether a user has any particular clearance; neither can they change the classification of an object (e.g., from Secret to Top Secret).

Unlike regular operating systems, the design of a trusted operating system requires far greater scrutiny and careful consideration, involving implementation of appropriate and consistent policy features assembled together with a high degree of assurance.

Key features of a TOS include:

- User identification and authentication
- Mandatory access control (MAC)

Discretionary access control (DAC)

- Object reuse protection
- Complete mediation
- Trusted path
- Audit
- Audit log reduction
- Intrusion detection

There are only a handful of TOSs that can be used as the foundation for accredited CDSs. TOSs must be EAL4+ Common Criteria and/or NIAP OSPP compliant. Example TOSs include:

- Linux/SELinux
  - Red Hat Enterprise Linux (RHEL)
- UNIX
  - Oracle Solaris w/Trusted Extensions
  - HP/UX
- Secure Trusted Operating Program (STOP)
  - BAE STOP

Anti-Malware Software

Anti-virus (AV) is a key component of any network security solution. In addition to detecting known threats and variants, modern AV can also apply heuristics to detect anomalies or suspicious code. Any suspected malware or unexpected content is automatically quarantined.

Data Loss Prevention & Insider Threat Detection Software

While the focus within many other network security tools lies heavily on the prevention of infiltration by bad actors, the exfiltration of sensitive content is of particular concern in high-to-low transfers between confidential or classified environments. Data loss prevention (DLP) software is designed to detect unauthorized transfers of data out of a particular environment (typically by an inside actor), and prevent that data from being transferred.

Identity & Access Management Infrastructure

As with any secure system, CDSs require a comprehensive role-based access control (RBAC) framework and infrastructure to maintain a separation of duties between users and administrators. RBACs are also often supported by separating data and admin ports and/or prohibiting administration while the device is in use.

Logging & Audit Systems

Audit log capabilities are vital to identifying any malicious or accidental attempts to manipulate or bypass the CDS. These logs must be heavily secured against any potential disabling or alteration in order to maintain absolute integrity. In addition, there must be a comprehensive review and analysis process in place to take appropriate corrective actions.

Application Layer Firewalls

Unlike their predecessors (standard firewalls), next-generation firewalls (NGFW) and web application firewalls (WAF) and other network "guards" provide a means to filter at the application layer, meaning that they can control access to, from, or by an application or service. This capability is important to maintain a highly restricted access environment and such firewalls are integral front-line pieces of the layered CDS architecture.

CDS-related Hardware

If software is the brain of a CDS, then the hardware is the musculature and skeleton, enabling additional capabilities and holding it all together. This includes both the chassis and onboard hardware-enforced elements which provide extremely secure mechanisms beyond what software can provide on its own.

Data Diodes

As hardware-enforced network separation and one-way data transfer devices, data diodes provide an ideal means to enable physical and logical domain separation. Indeed, such hardware-enforced separation is now required according to current NSA RTB guidelines, and no CDS can be accredited without it.

Data diodes come in a few different types, depending on the method used to enable the data transfer, including optical (optocoupler), electrical (dielectric isolator), and field-programmable gate array (FPGA). There are some performance, efficiency, and cost advantages to each, although the market in general has moved toward smaller, faster technology such as micro optocouplers and FPGAs.

There are a handful of established data diode vendors (full disclosure, Owl Cyber Defense is a prominent data diode vendor) which supply CDS vendors as well as government and defense customers to fulfill this requirement.

COTS Secure Mechanisms

- **TPM 2.0 and Trusted Boot**

  Trusted platform modules (TPM) are a passive hardware-based security technology, typically chips soldered directly onto a motherboard or other cards, that enable cryptographic functions. They are tamper-resistant, and immune to software-based attacks. TPMs are useful for storing trusted boot states and securing other operating system functions.

- **Intel® Boot Guard**

  Boot Guard is a hardware-based platform security technology designed to prevent repurposing of the platform to run unauthorized software, including the execution of boot blocks. They are interoperable with TPMs and are rooted in protected hardware infrastructure.[1]

- **UEFI Secure Boot**

  Unified extensible firmware interface (UEFI) secure boot is a multi-vendor (including open source) verification tool designed to protect a system against malicious code being loaded and executed early in the boot process, before the operating system has been loaded. This is to prevent malicious software from installing a “bootkit” and taking control over a computer. If an invalid binary is loaded, the user is alerted, and the system will refuse to boot the binary.[2]

Hardware Platforms

Appliances & Cots Servers

Vendor-built rack-based appliances often include a number of self-protection and tamper-proof or tamper-evident mechanisms to prevent physical manipulation of the device. These include tamper switches, tamper-evident tape, and operational fail-safe mechanisms that check to ensure all systems are functioning normally.

Ruggedized

In addition to hardened and tamper-proof equipment, a Mil-Std-810G or later certification is required by the U.S. government for many mobile and tactical applications. These ruggedized solutions are typically utilized in defense-oriented environments, such as vehicles, aircraft, and ships.

1 https://www.intel.com/content/www/us/en/intelligent-systems/shark-bay/security-technologies-4th-gen-core-retail-paper.html
2 https://www.linuxjournal.com/content/take-control-your-pc-uefi-secure-boot

Security Principles of Cross Domain Solutions

In order to better understand the context of cross domain solutions, one must look first at the principles that guide their design and implementation. While the software and hardware components of CDSs define technical capabilities, the principles and functional requirements define their utility and create a baseline feature set.

As multifaceted, high-assurance network security devices, CDSs are uniquely equipped with a variety of policy-based, contextual, and operational security features that set them apart from other network technologies. While the following is not an exhaustive catalogue, it represents the standard list of transfer CDS functions and mechanisms. This includes context-appropriate security-enforcing mechanisms, layered secure architectures, and secure operation features.

Context-Appropriate Security-Enforcing Mechanisms

In order to enforce security measures and differentiate authorized transfers of data and information from unauthorized ones, security-context-appropriate mechanisms must be implemented. A security context in a domain is defined by a domain security policy. Access clearance is first given to users and processes accessing, manipulating, or transferring data based on rules defined in an attached security context. These rules may include a variety of authentications, filtering, transformation, protocol manipulation, access controls, and other enforcement mechanisms. In the Linux security module (LSM) in SELinux, the security context is an extended attribute.

Policy Enforcement

Security policy encapsulates an array of controls all determined by the security rules set forth from the organization itself. This provides a security context for every user and process based on the actions available in the system. Enforcing this policy is key to maintaining security within the system, domain, and across the organization itself.

Known-Good

“Known-good” is a security principle similar to whitelisting, in which the system blocks any unexpected data, protocols, etc. and only allows users, actions, and data which are known to be authorized and comply with the defined structure and content. This also applies to the channels and ports which only allow data which is expected/requested on the appropriate pathway. It does not, however, assume trust – all data is subject to content inspection and filtering.

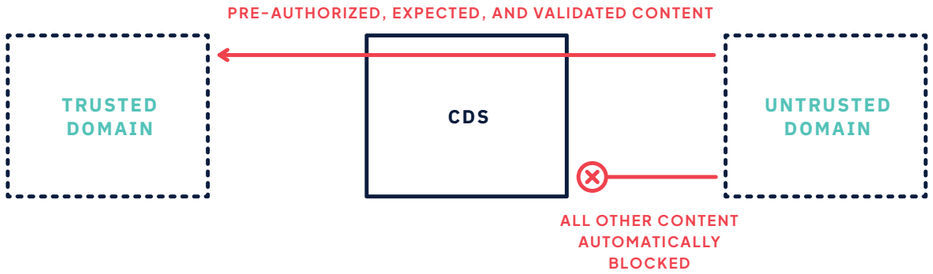

Data Transformation & Normalization

Prior to inspection and filtering, unstructured or complex data types, such as MS Office documents or other packaged files, must be transformed in order to create a deconstructed data set, wherein the complex data is broken down into simpler file types, or even binary or text. The CDS filters can then detect underlying or hidden elements that may otherwise be missed upon higher level inspection. As a subset of this type of transformation, video can also be transcoded to help reduce and mitigate the threat of embedded steganography.

Diverse data sets are also normalized in order to increase the speed and efficiency of analysis and filtering. This involves adapting the schema to accommodate the data set and removing redundancies and other anomalies within the data.

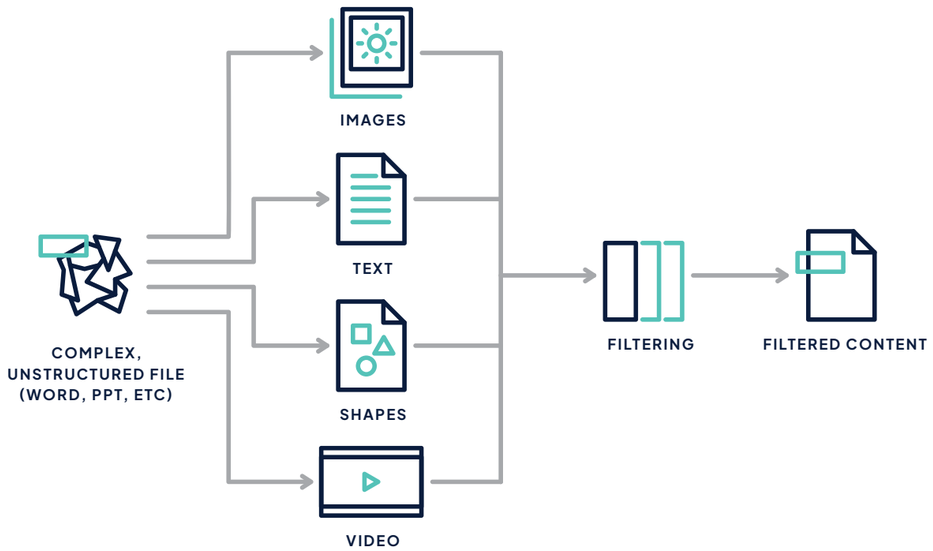

Content Filtering & Quarantine

Data filtering is one of the key differentiators of cross domain solutions from other network security solutions. There are both standard data filters that have been developed by government agencies and standards bodies, and custom filters which can be designed for fit-for-purpose applications. Data filters vary in type for structured and unstructured data, depending on the need to normalize or transform the data to detect embedded threats or unauthorized/unexpected content.

In addition to data filters, cross domain solutions also apply validation checks, to ensure the input is well-formed, and data sanitization, which can involve removing or changing content in data that passes validation checks. If the data fails any of the checks or filters, it is either dropped or quarantined.

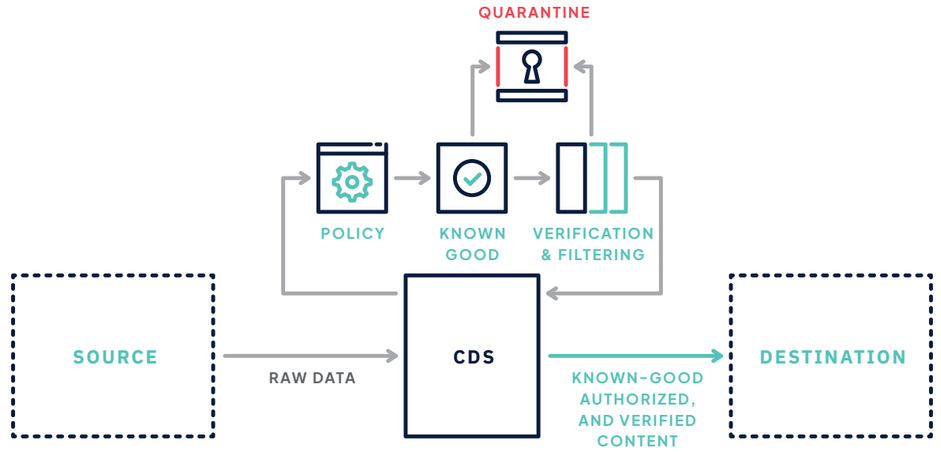

Data Loss Prevention

In addition to preventing threats from entering the organization, cross domain solutions are designed to prevent data loss and unauthorized data exfiltration. Any unauthorized data is filtered and prevented from transfer, before it leaves the source (trusted) domain.

Data Provenance

Provenance refers to the chronology of the origin, development, ownership, location, and changes to transferred data. It may also include personnel and processes used to interact with or make modifications to the data. Verifying data provenance is key to evaluating risk and identifying potential threat vectors or modifications introduced both at the point of data creation and at any point after.

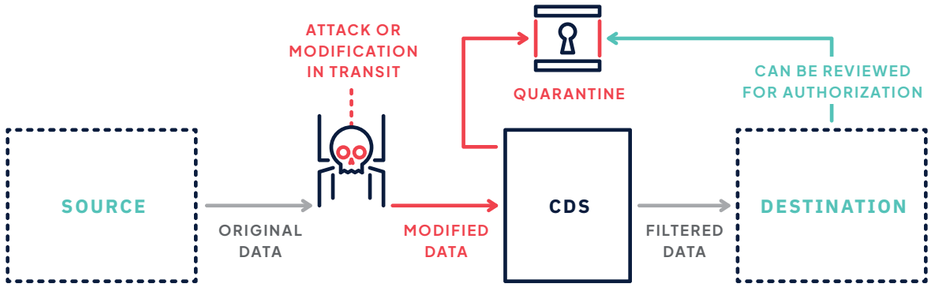

Domain Separation & Protocol Break

In order to assure true domain separation, cross domain solutions incorporate a hardware-enforced network segmentation and protocol break via data diodes. Data flows are transmitted between domains on one-way transfers, with a protocol termination on the send side, and a protocol resume on the receive side. Acknowledgments and protocol translation are handled on each side via proxy.

For bidirectional use cases, the CDS can be configured to route acknowledgments and other data through a separate return path. Both paths still feature a protocol break, which eliminates the exposure of network routing information of each domain from the other.

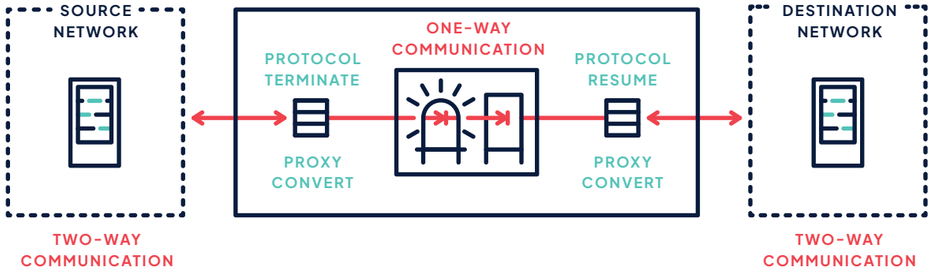

Flow Control

Flow control (or “information flow control”) is a mandatory access control (MAC) mechanism to ensure that data passes through processes in a particular order, without bypass. A key component to flow control in cross domain solutions is a type enforcement (TE) system which enforces the security policy by applying a security type to data, depending on its sensitivity. Data flows are then restricted regardless of encryption, role-based access controls, or other complementary mechanisms.

Secure Architecture and Design

Beyond the functions and capabilities, also vital to security are the design principles and architecture of the system itself. Keep in mind that cross domain solutions are designed for use against the most advanced, persistent, and intelligent threats in existence today. Unless the system is planned, designed, built, and programmed with the strongest possible security at every point, it runs the risk of compromise.

Defense in Depth

As a cybersecurity best practice for all systems, it’s vital to implement multiple layers of diverse defenses to prevent compromise from a single point of failure. This includes everything from the OS and hardware to the applications and data filters. So long as the attacker cannot make it through every security layer, the system and the domain will remain secure. Defense in depth also acts as a check on user error and insider threats. The more layers of security in place, the less likely any potential data leakage or unauthorized exposure of the secure domain and its contents.

Secure by Design

While it may seem fairly obvious, or perhaps contradictory, as cross domain solutions are by their definition “add-on” security devices, the nature of their high assurance is based on the adherence to the highest possible levels of security in the design of their software and hardware, from the ground up. By utilizing the best available security in every possible aspect, cross domain solutions continuously enforce the necessary authentication, authorization, confidentiality, integrity, accountability, and availability requirements according to the security policy, even when the system is under attack.

RAIN

The “RAIN” acronym in cross domain solutions stands for Redundant, Always Invoked, Independent, and Non-Bypassable. Now the foundation of CDSs itself, the RAIN concept is based on the concept of a reference monitor within operating systems, and was brought back to the fore with the new “Raise the Bar” (RTB) mandate from the NSA.

Redundancy, especially within a secure data transfer system, menas that security-relevant components (including filters and domain separation) are invoked twice for each unique function.

By **always invoking** security, the intended filtering will always occur, reducing the chance that a threat could sneak through under the guise of a trusted file or data stream.

Each function within the transfer and filtering must be created **independently,** reducing or eliminating a single point of failure by compromising a single component or programming code.

CDS security measures must be **non-bypassable,** including within the software and data stream, such that a threat cannot find or exploit "backdoors" or other circumvention methods. This concept can also be extended to hardware, and even the physical environment.

System Assurance & Secure Operation

Lastly, in order to guarantee trust in the system, there must be ways to verify system assurance and maintain secure operation, in the face of attacks of unknown origin, persistence, and sophistication. In addition to securing the domain and maintaining strict data control, each element of the CDS operation itself must be selected carefully, designed, delivered, and manufactured securely, locked down tightly, tested rigorously, and maintained regularly. The RTB mandate also now imposes strict security requirements on both the supply chain and development environment.

Trusted Platforms & Components

While the specific components are discussed in more detail in What Makes a Cross Domain Solution, trusted platforms are vital to the secure operation of cross domain solutions. A trusted system must have the necessary security functions and assurance that the security policy will be enforced and be able to process a range of information sensitivities. Again, the security policy of the device is paramount to secure operation. Trusted platforms are often based on an added "security module" to the operating system, while trusted components can be anything from trusted platform modules (TPM) to data diodes.

Secure Administration

Proper administration is the key to managing any system. However, administrative access is also one of the first targets of any bad actor, so steps must be taken to ensure that such access is secured against compromise. Methods of secure administration include, but are not limited to: multifactor authentication, least privilege access to all users and processes, highly restricting the number of administrators, dividing administrative duties to two or more personnel to diminish the power of each, creating secured privileged access points, and automated audit logs.

Self-Protection

In order to prevent physical or digital tampering or manipulation, it's important that high assurance devices, such as cross domain solutions, include self-protection mechanisms. These include, but are not limited to: trusted boot processes, separate administration and operation modes, tamper switches, internal process verification checks, and other intrusion detection and prevention functionality.

Secure Failure

Unless the system has a secure fail-state, any attempt to disable or turn off the device or other potential hardware failure could compromise the security of the trusted domain. As such, cross domain solutions typically have a "no-transfer" fail state, which does not allow any data to pass into or out of the trusted domain.

Opaque Operation

In the interest of proper use and functionality, the operation of the cross domain solution will likely be known to the trusted domain users. However, to maintain security, details of its operation must not be leaked to untrusted domain, nor should any failed commands or attempts to probe the system from the untrusted domain provide any domain- or security-related information.

Independent Assessments and LBSA

In the United States, cross domain solutions are subject to certification by the U.S. Government, administered by an organization within the National Security Agency (NSA) called the National Cross Domain Strategy and Management Office (NCDSMO). This certification is called the "Lab Based Security Assessment (LBSA)."

The LBSA certification requires meticulous technology and processing assessment that involves thoroughly testing every aspect of the device. Once passed, the device can be eligible for the "Baseline List" of solutions certified for U.S. intelligence and defense use.

As the testing under the NCDSMO LBSA is far more rigorous than other standards, such as Common Criteria EAL, it typically supersedes any other certifications.

Regular Maintenance, Training, and Support

System upkeep and patching is a vital part of keeping any system secure, and cross domain solutions are no different. As such, in the United States, they are mandated to be kept up to date with the latest versions of software and any available vulnerability patches. In addition, product training and support must be made available to the user in order to guarantee proper use and continued security.

How are Cross Domain Solutions Used?

CDSs are used across government, intelligence community, and defense networks to control information flow and maintain the assurance of trusted environments. They are employed for a variety of data collection, distribution, and dissemination uses, from offloading maintenance information from sensitive vehicles to disseminating information to multiple coalition forces.

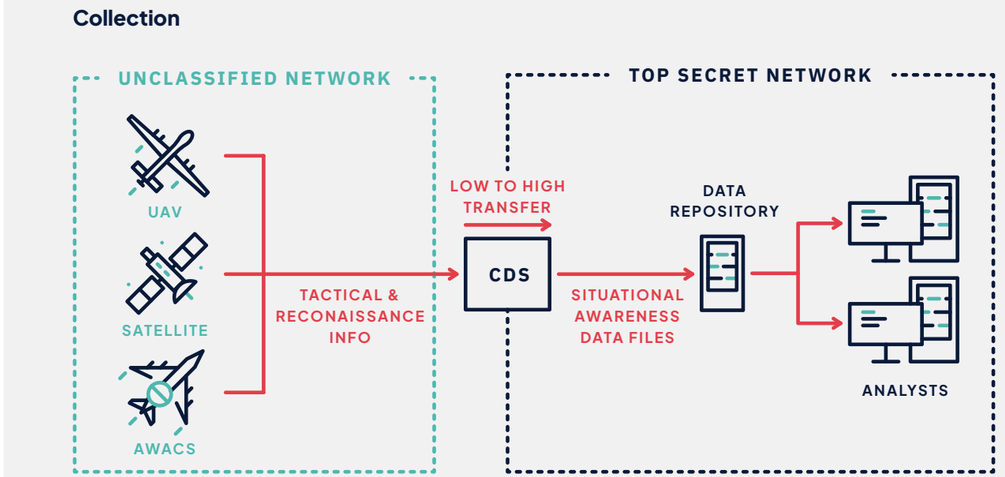

- UAV sensor data and video
- Ground force protection video
- Open source data, imagery, video
- Exploitation data, imagery, video
- Satellite imagery
- Telemetry data
- Audit events and record data

Distribution

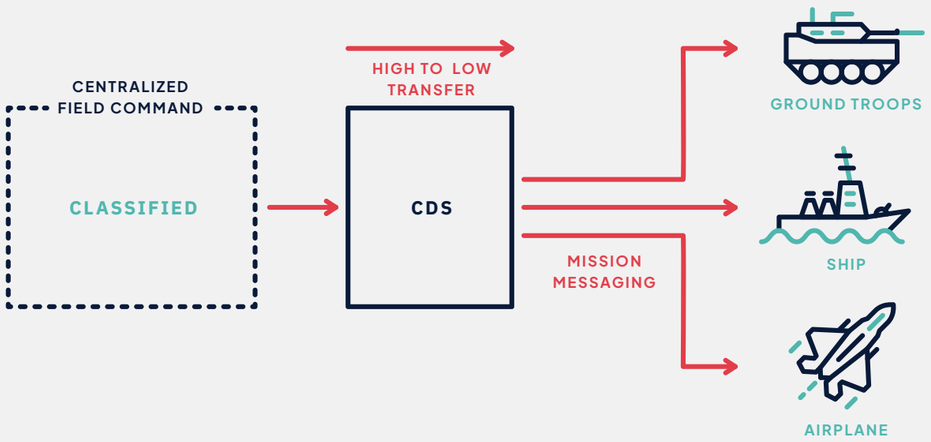

- Fixed Format messaging
- Mission messaging
- Repository query/response
- Alerts and tips
- Command and control
- Information/Intelligence processing content

Dissemination

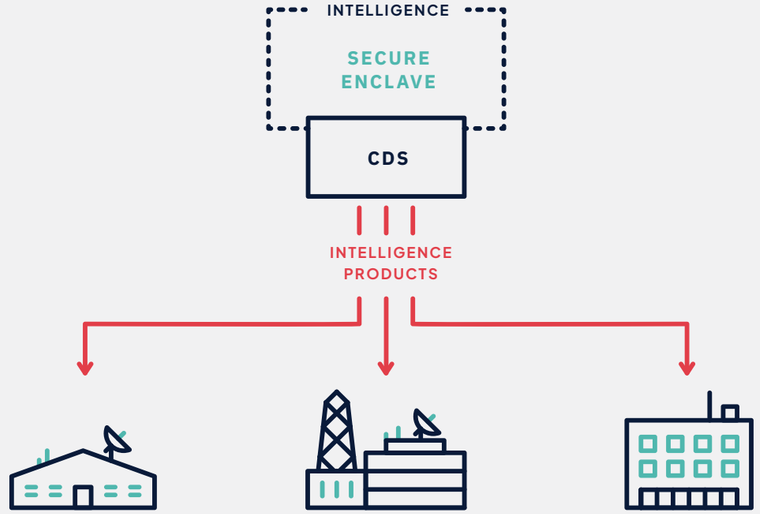

- Information products
- Intelligence products
- Situation reports

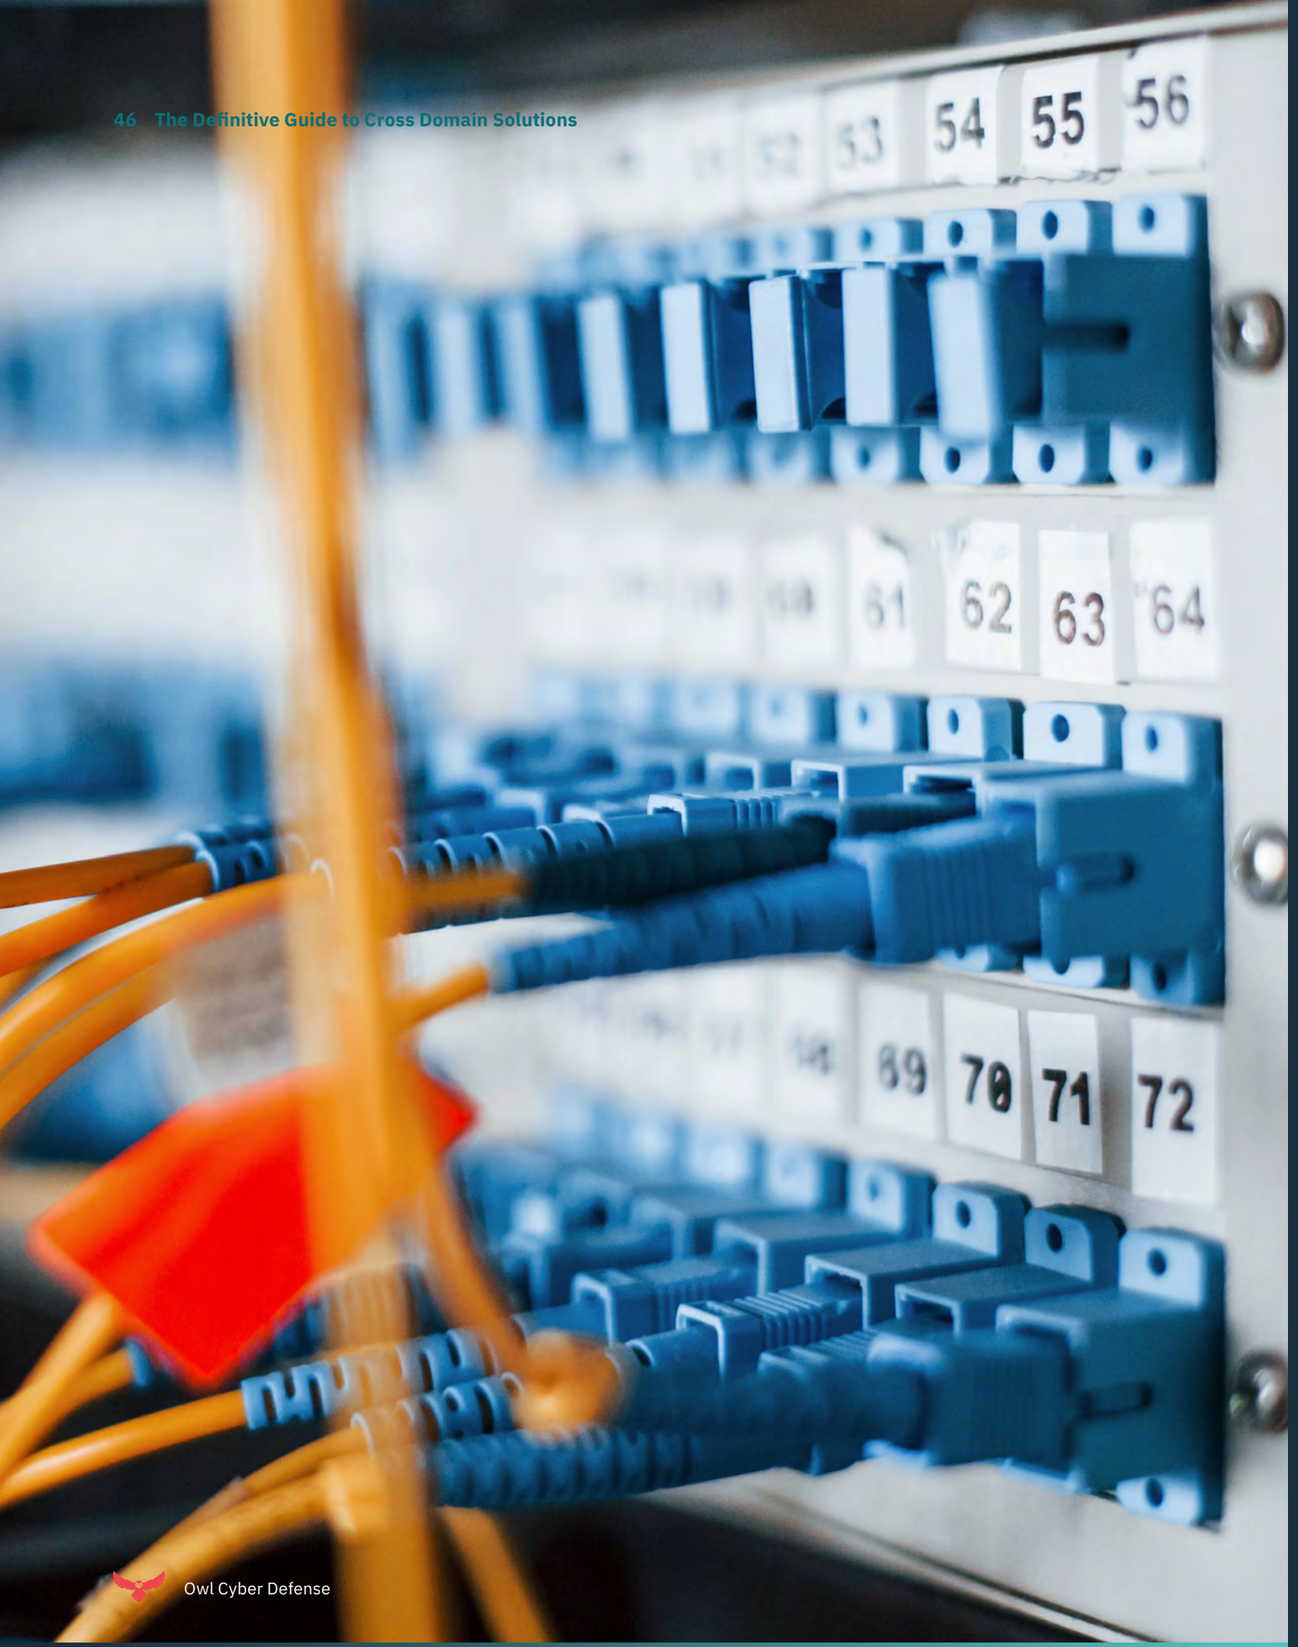

Summary

While certainly not necessary in every network, transfer cross domain solutions provide the highest assurance means to ensure trust in critical environments. There simply is no greater data transfer security technology available today, and the latest solutions are only becoming more adept and effective at mitigating even the most advanced, persistent threats. As these threats evolve and move to new targets, CDSs will be there, with similar technologies being built for critical infrastructure and commercial markets even today. As the saying goes – “If you want insurance, get a firewall. If you want assurance, get a CDS.”

For more information on Owl Cyber Defense cross domain solutions and available technologies, visit **www.owlcyberdefense.com.**

OWL Cyber Defense™

Owl Cyber Defense Solutions, LLC, headquartered in Columbia, MD, leads the industry in data diode and cross-domain network cybersecurity solutions for faster, safer and smarter decision making. We create solutions tailored for high-risk sectors including the military, government and critical infrastructure. Our advanced technologies enable secure, near-instantaneous collaboration, bridging network barriers to protect critical missions. With a focus on scalability and interoperability, Owl ensures that organizations can maintain secure, reliable, and compliant communication channels against evolving cyber threats.

Visit **www.owlcyberdefense.com** or contact us at **info@owlcyberdefense.com** for more details.

203-894-9342 | owlcyberdefense.com