Global Information Assurance Certification Paper

Shedding Light on Cross Domain Solutions

GIAC (G SEC) Gold Certification

Author: Scott D Smith, Smith24197@gmail.com

Advisor: Richard Carbone

Accepted: November 6, 2015

Abstract

Never before has the need for collaborative data become more important, nor more vulnerable. APTs are becoming increasingly more advanced at jumping between systems. Meanwhile, today's users are adept enough with computer systems to employ workarounds to traditional network separation requirements. These growing external and internal threats have an accumulative impact in creating and exploiting vulnerabilities. Correcting the aftereffects of this problem is a zero-sum game. To address the root of the problem, system architects must weigh information sharing capabilities against security controls between different network enclaves. This paper will explore the concept of cross-domain solution types by discussing their respective capabilities, common architectures, and critical considerations.

1. Introduction

As a general practice for information security assurance, a need-to-know model has dominated the industry. It has led to the development of information silos which, while ideal from a defense-in-depth perspective, insulates context and makes cross-referencing cumbersome if not impossible. Would it not be great if information, once categorized with a security designation, was not restricted to a single processing environment? Should a solution not be available that enables data to flow between isolated networks, after undergoing a series of checks and balances? The concept sounds simple enough.

A permissive network inherits certain risks that much be weighed against the value of transferable data capabilities. Perimeter security controls need to be placed between enclaves according to classification hierarchy. These boundaries authenticate transaction queries, regulate information flows, and mitigate the impact of compromises.

It then becomes necessary to break down the legacy information silos in terms of security classifications and releasibility. The silos become domains, which are information processing environments defined by the level of information sensitivity. The Government of Canada (GoC) categorizes information sensitivities as Classified, where unauthorized disclosure can be described in terms of national injury, and Protected, where the injury is described according to a person or organization [PWGSC, 2015].

Table 1 – GoC Information Sensitivities [Source: Department of National Defence (DND)]

Information Sensitivity	Classified (i.e. National interest)	Protected (i.e. Individual or Organization)
Unauthorized disclosure could cause exceptionally grave injury	Top Secret	Protected C
Unauthorized disclosure could cause exceptionally serious injury	Secret	Protected B
Unauthorized disclosure reasonably expected to cause injury	Confidential	Protected A

These sensitivity labels can be mapped to domains based on security controls. Going back to GoC as an example, domains are classified as Top Secret, Secret, Designated and Unclassified. Domains can further be broken down by caveats, which are releasable sensitivities restricted to access by certain groups or categories. Practical examples include Secret Releasable, which may be provided for disclosure to trusted third parties, or Unclassified Public Access, which may host all internet-facing services.

Table 2 – GoC Domains and Caveat Examples [Source: DND]

Domain	Information Sensitivity	Caveat Examples
Top Secret (TS)	Top Secret	Joint Worldwide Intelligence Communication System (JWICS)
Secret (S)	Secret, Confidential, Protected C, Protected B	Consolidated Secret Network Infrastructure (CSNI) – Canadian Eyes Only (CEO)
		Secret Internet Protocol Router Network - Releasable (SIPR Rel)
Designated (D)	Protected A	Defence Wide Area Network (DWAN)
		GoC Network (GCNet)
		Non-Secure Internet Protocol Router Network – Releasable (NIPR Rel)
Unclassified (U)	N/A	General Purpose Network (GPNet)

The relationships between domains and caveats must further be described. Should a universal policy be adopted throughout? Perhaps some domains must favor integrity (i.e. Biba model [Biba, 1975, p.19]) over confidentiality (i.e. Bell-LaPudula model [Bell, 2005, p.11])? Once these questions have been answered, one can begin reviewing the most suitable solutions to enable access to and transfer of data between domains.

2. Definitions

2.1. Types of CDS

The Committee on National Security Systems (CNSS) defines A Cross Domain Solution (CDS) as "a form of controlled interface that provides the ability to manually and/or automatically access and/or transfer information between different security domains [CNSSI 4009, p.23]." CDSs can therefore be broken down into three types: Access, Transfer, and Multi-level solutions (MLS).

2.1.1. Access Solution

An access solution describes a user's ability to view and manipulate information from domains of differing security levels and caveats. In theory, the ideal solution respects separation requirements between domains by preventing overlaps of data between domains, which ensures data of differing classifications cannot 'leak' (i.e. data spill) between networks at any host layer of the OSI/TCP model. In practice, however, data spills are an ever-present concern that system designers attempt to mitigate within acceptable risk levels. For this reason, data transfer is addressed as a separate CDS.

Figure 1 provides a comparison between access and transfer solutions as complimentary CDSs. Note the access solution is located between the user and differing domains.

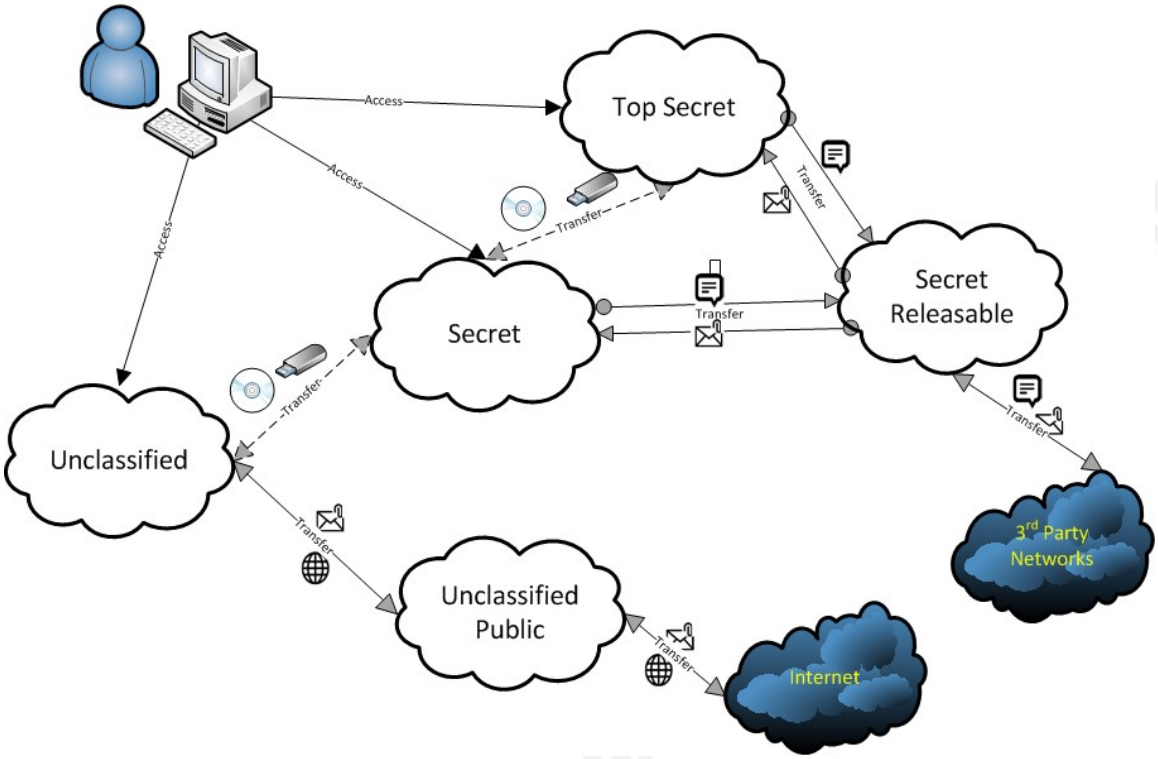

Figure 1 – Access to Domain-Pair Relationships Map [Source: DND]

2.1.2. Transfer Solutions

A transfer solution provides the ability to move information between domains of differing security levels and caveats. Transfer solutions must respect data sensitivities and governable policies of each domain or caveat to prohibit operational security incidents.

Figure 2 provides a comparison between access and transfer solutions, with an emphasis on the relations and services between domains. Note the transfer solution is located between differing domains. Each domain-pair's relationship and security controls will determine their appropriate transfer services. The overall solution may be comprised of different policies, such as the direction of data flow and transfer protocol, for each domain-pair's relationship. These may be as simple as unfiltered email and web browsing between the internet and unclassified public domains, bidirectional FTP between all caveats within the designated domain, or unidirectional fixed-length XML messaging from Top Secret to Secret.

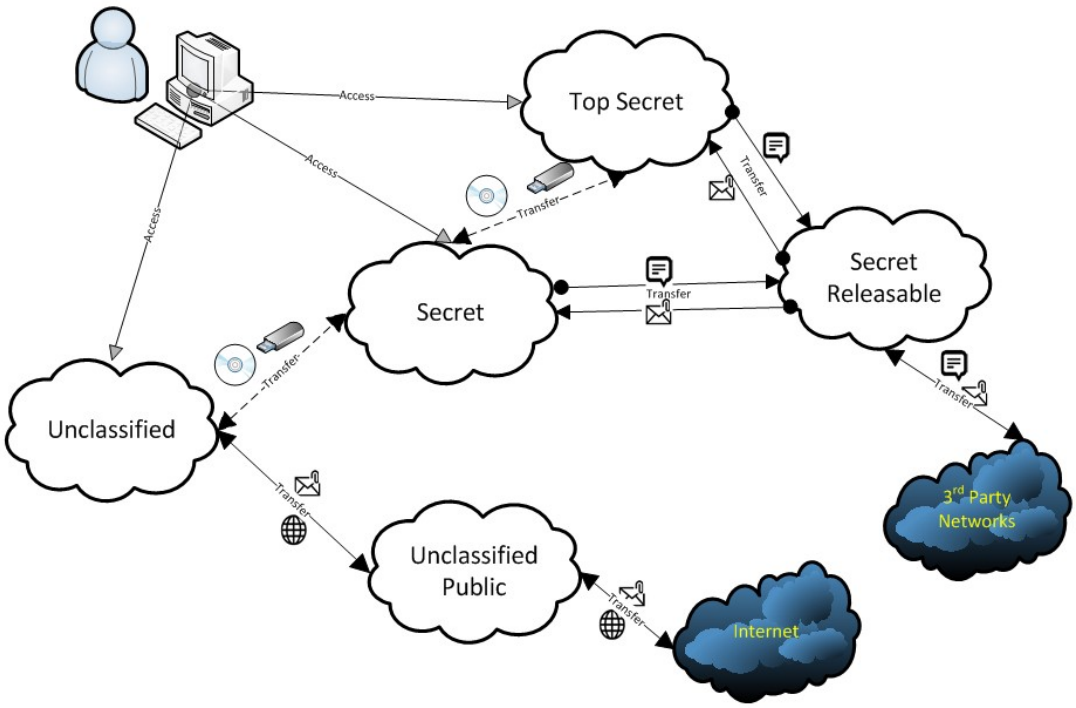

Figure 2 – Domain-Pair Transfer Relationships Map [Source: DND]

While there are many different types of transfer solutions, they can all be described in general terms relative to common network devices. Primarily, a transfer CDS must control addressed connections between domains similar to a firewall. A content inspector then needs to examine the data according to approved policies and determine the appropriate handling instructions. A diode will ensure directionality of data flow at the physical layer, preventing data spills. Finally, just because a CDS has ensured the right people are sending proper messages in the correct directions does not mean the data is usable. Like a gateway, protocol translation, signal conversion, and even impedance matching may be necessary to ensure Cross Domain systems are interoperable. Figure 3 illustrates the four functions below:

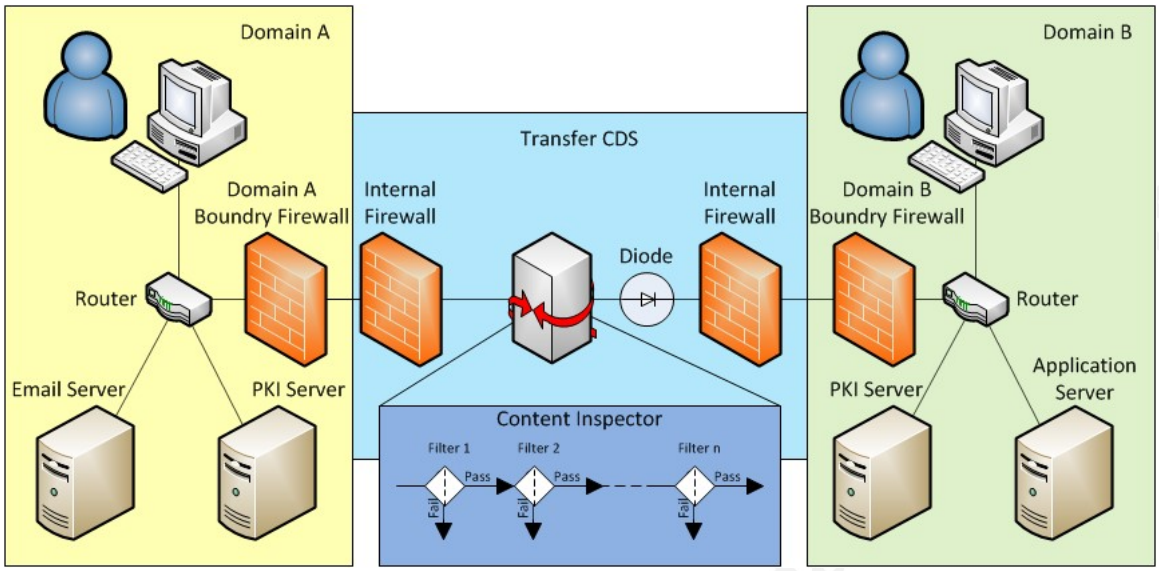

Figure 3 – Illustration of Transfer CDS Functions [Source: DND]

From a defense-in-depth perspective, all four common network device functions are essential to network security. Table 3 below breaks down the vulnerabilities each appliance addresses alone as compared to a full transfer CDS. Note the High Assurance Guard (HAG) is effectively a proxy server capable of deep content inspection.

Table 3 – Appliance Protection Comparison Chart [Source: DND]

Threat	Fire wall	Anti-Virus	Data Diode	HAG	Full CDS
Data Leakage			✓	✓	✓
Network Access Control	✓		✓	✓	✓
Data Label Inspection				✓	✓
Protocol Break				✓	✓
Hidden/Inappropriate Content				✓	✓
Malicious Code/Malware		✓		✓	✓
Delete Steganography				✓	✓
Embedded File Inspection				✓	✓
Metadata Inspection				✓	✓
Zero Day				✓	✓

✓ Somewhat Protects Against ✓ Consistently Protects Against

2.1.3. Multi-level Solutions

Access and transfer solutions rely on multiple single level (MSL) systems that maintain the separation of domains; this architecture is considered multiple individual levels of security (MILS). A multi-level solution (MLS) differs from MILS architecture by storing all data in a single domain. The solution uses trusted labeling and integrated Mandatory Access Control (MAC) schema to parse data according to user credentials and clearance in order to authenticate read and right privileges. In this manner, an MLS is considered an all-in-one CDS, encompassing both access and data transfer capabilities.

The concept can lead to significant performance advantages over conventional CDS models by the sheer reduction in processes necessary to access and manipulate data. The trusted data labeling and consolidation of domains removes the need for content inspection, filtering, and sanitization operations. Likewise, synchronization and replication errors are eliminated as all clients have access to the same server. In practice, however, an MLS is exceptionally difficult and expensive to develop [Chen, 2010]. For this reason, the remainder of this paper will focus on conventional transfer and access solution architectures.

3. Architectures

3.1. Access Solutions

3.1.1. Isolated Domains

The isolated domains approach is based on the segregation of information processing environments. Operators employ a unique workstation for each domain to which they require access, as shown in Figure 4 below. Members of the Canadian Armed Forces refer to this set up as ‘swivel-chair.’

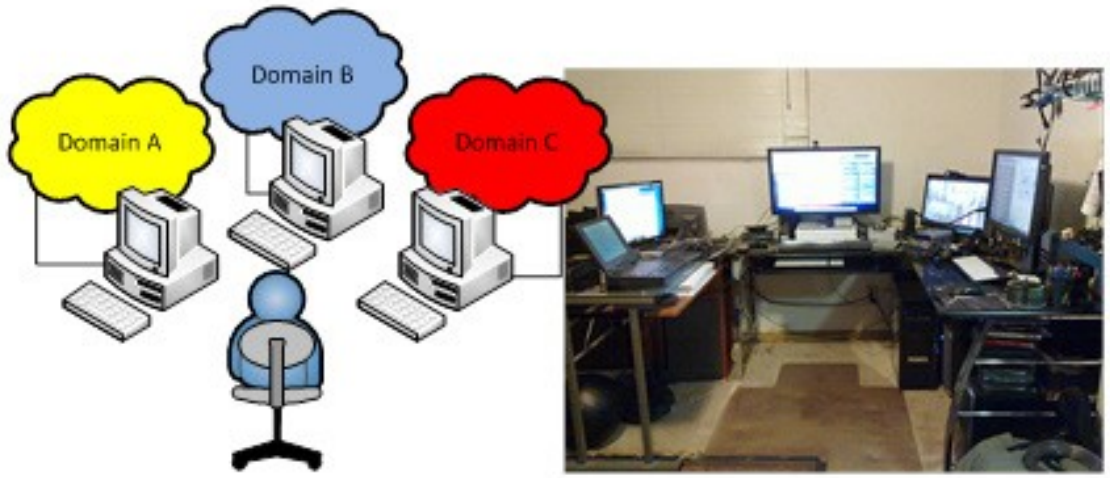

Figure 4 – Isolated Domains [Source: DND]

Domains can be isolated by virtualization through physical infrastructure, although common VLAN attack methods make physical infrastructure the preferred solution [Roullier, 2004, p.23]. Independent cabling and workstations are required in either distribution method. Since each domain is physically separated from one another, the isolated domain approach boasts high redundancy and excellent defence-in-depth. This also makes the isolated domain approach the most expensive architectures in terms of capital, administrative resources, and power requirements. For this reason, it is typically only seen in mission-critical government and/or industries, and is considered a legacy security model.

3.1.2. Periods Processing

In a periods processing design the network as a whole changes its security classification at specific times. All connected devices must be sanitized between periods to ensure no residual data remains on the hosts or network storage devices. Users therefore require only a single workstation on a single network to access multiple domains, as illustrated in Figure 5 below. There is an argument, however, about whether periods processing can be considered a true access solution since the workstation does not traverse domains so much as the entire network is reclassified.

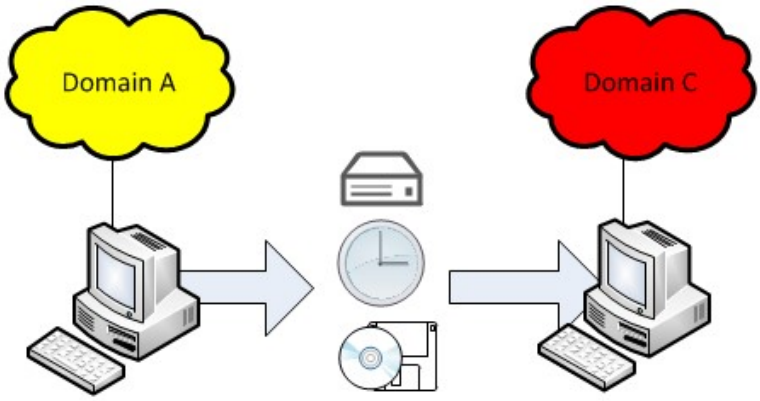

Figure 5 – Periods Processing [Source: DND]

The sanitization cycle is of key importance to this access solution. Typical systems rely on built-in read and write processes to clear data for future use by erasing the address where the data resides. The data, despite lacking pointers to it, remains unaltered until eventually overwritten by new data. Clearing is considered insufficient for sanitizing sensitive information and cannot be relied on to ensure data cannot spill between domains, therefore purging processes must be considered. These may include block erase, in which the target data is overwritten prior to erasing pointers, and cryptographic erase, whereby all data is encrypted at rest state and the crypto keys are erased with address pointers [Kissel, 2014, p. 24].

While periods processing is more economical than isolated domains, it is the least convenient and productive access solution for users. By nature of the design, users are constrained to work around the period processing schedule. Productivity is lost prior to domain switching, when hard drives may need to be swapped and data backed up on removable media. Furthermore, the sanitization cycle may take some time to purge the network as a whole and return it to an operable state.

3.1.3. KVM Switching

Keyboard, Video monitor, and Mouse (KVM) switches are an economic hardware solution to access independent domains from a single workstation. The switch connects at the physical layer to ensure separation between domains, as illustrated in Figure 6

below. While the user is limited to accessing one domain at a time, they can be switched at the user's convenience.

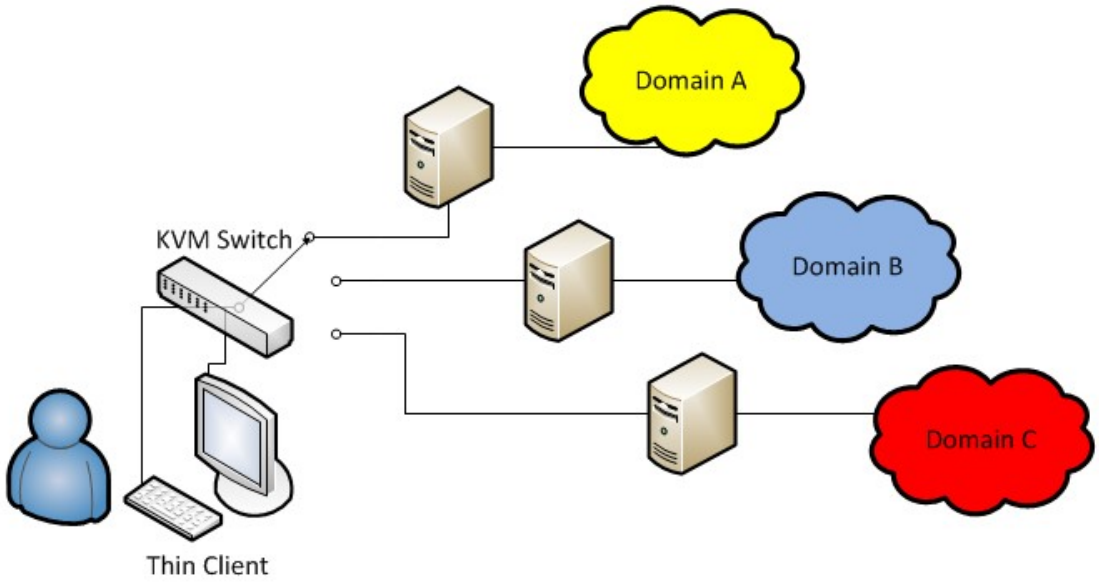

Figure 6 – KVM Switch [Source: DND]

These switches can accommodate a mix of thick and thin clients, and modern switches support automated crossover of other user devices, such as speakers and card readers. This feature can be considered a major risk of data spills when switching between domains without sanitization processes, especially for memory sticks and thick client hardware that remain connected. For this reason, most of today's KVM-based access solutions favor thin clients with Endpoint USB authorization solutions.

KVMs are not immune to vulnerabilities. As the pivot point between domains, software-based versions are the ideal ingress vector for attackers. KVM's in access solutions should therefore be hardware-based and developed from discreet electrical components; this does not remove the risks from the workstation itself. Many of today's mice, monitors, and especially keyboards feature re-programmable components that could be compromised to transfer data from switched domains [Adder Technology]. Discreet diodes placed between the user interface devices and the switch will successfully protect against authorized data spills, although the threat of malware-injection remains present.

3.1.4. Partitioned Workstation

Partitioned workstations rely on virtualization to access two or more domains simultaneously. The host operating system (OS) ensures separation of the virtual OSs (virtual machines, or VMs), which interact with their respective domains, as shown in Figure 7 below. Note the network interface card (NIC 1) is shared between the VMs, and logically described as NIC 2 for Domain A and NIC 3 for Domain B.

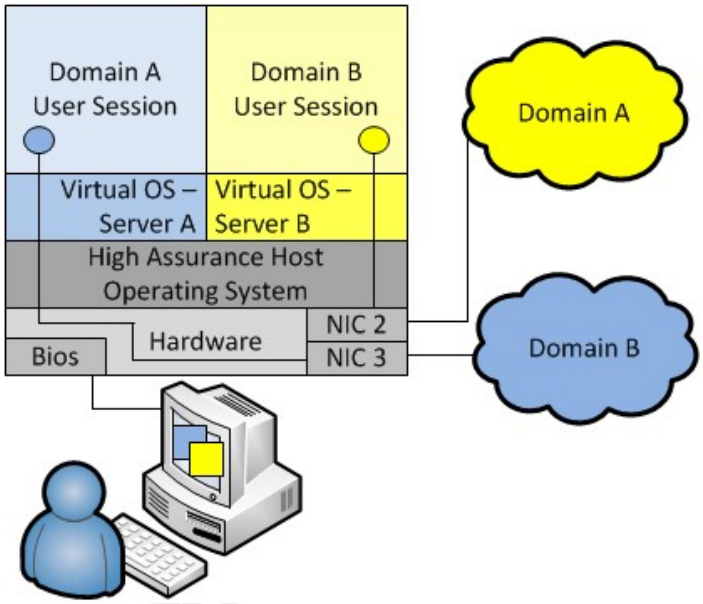

Figure 7 – Partitioned Workstation [Source: DND]

VM-based access solutions require a specialized host OS to ensure domains remain separated. The kernel in particular must be hardened to ensure hardware is segmented and sanitized between each VM's processes to prevent data spills, as well as the hypervisor for thin client distributions. This can slow down performance, or lead to incompatibilities between third-party applications.

Virtualized platforms must also be hardened and patched against VM-based vulnerabilities that exploit common virtual OS relations with host processes to escape the virtual environment. VEMON, a vulnerability in virtual floppy drive code found in many VM platforms, was the first exploit to compromise default configuration of multiple VM platforms discovered in 2015 [CrowdStrike, 2015]. Other vulnerabilities specific to distinct VM platforms have been reported as early as 2007.

3.2. Transfer Solutions

3.2.1. Air Gap

An air gap indicates two domains that are physically isolated from one another. In order to pass information between domains the user is required to first copy the electronic data onto removable media, and then insert the media into a workstation connected to the second domain. The file transfer method between air-gapped domains is colloquially referred to as 'sneakernet.'

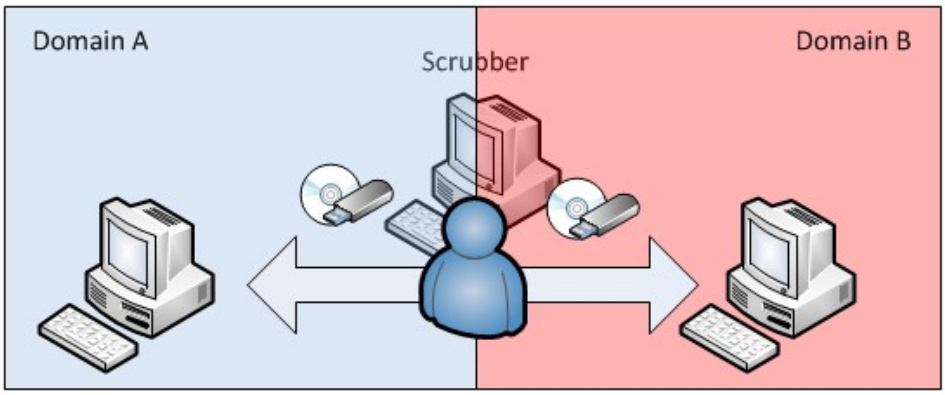

Figure 8 – Air Gapped Domains [Source: DND]

Air gapped systems are characterized as favoring throughput over latency. As Andrew Tanenbaum once wrote, "Never underestimate the bandwidth of a station wagon full of tapes hurtling down the highway" [Tanenbaum, 1989, p.57].

From an information security standpoint, they are a cost effective means to ensure confidentiality and integrity of data. While unauthorized access is nearly impossible to remote adversaries, cases of exploited vulnerabilities have been reported. Electronic eavesdropping has been achieved by intercepting low order electromagnetic frequency (EMF) signals generated from computer system monitors [Kuhn, 2004, p.17]. Advanced malware, such as Stuxnet, have been discovered to jump between network devices by hitching a ride on infected USB flash drives.

These external threats can largely be mitigated via common security measures such as physical separation requirements, EMF shielding, 'scrubbers' (i.e. standalone workstations with specialized malware and content inspection located between processing zones) and hardened network perimeter access. The internal vulnerabilities,

however, are much more difficult to address. A lack of centralized logging and auditing of file transfers leads to practically no oversight of the user community. Even the best human-enforced security policies and practices can be defeated internally, as proven by ex-US Army Intelligence Analyst Chelsea Manning's contribution to WikiLeaks in 2010 [Shanker, 2010].

3.2.2. Data Diodes

Also referred to as a unidirectional transfer solution, diodes restrict data flows in one direction at the physical layer. Typically, the low classified domain (i.e. low side) is permitted to send data to the high side only, as illustrated in Figure 9 below. This enables essential services, such as patching and data base replication, to flow to the high side without the risk of data spills.

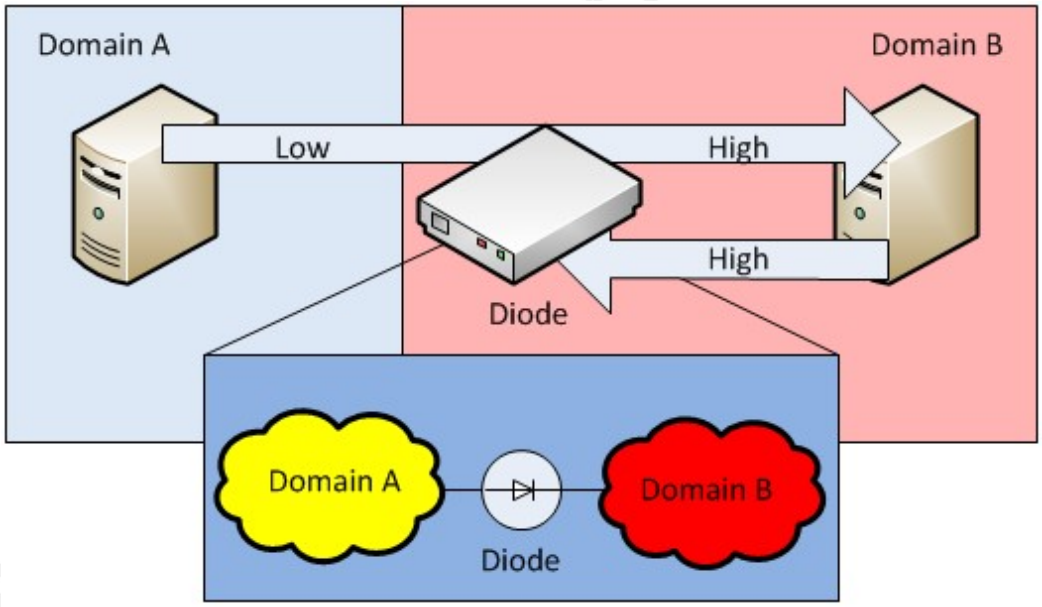

Figure 9 – Data Diode [Source: DND]

While many would argue a diode is, at best, only half a transfer solution, this model is useful in ensuring that data cannot leave the high domain even in the event of compromise. The inability to bi-directionally transfer data limits its application with certain protocols, such as TCP/IP that requires the three-way handshake [Scott, 2015, p.12]. Despite eliminating conventional attack patterns which require feedback, such as probing and reconnaissance activities, successful attacks bypassing diodes have been

documented, such as Stuxnet in 2010 [Sitnica, 2014, p.2]. Within a CDS environment, it is theoretically possible that an attacker with some knowledge of the high side's file structure and data labeling could craft malware capable of modifying or removing data.

3.2.3. Bi-directional Guard

A guard is a single appliance that provides all the functionality of the four common network devices that describe a transfer CDS: firewall, diode, content inspector and diode. They are usually purpose-built for the processing environment with emphasis on content inspection. Since the guard is effectively an all-in-one transfer CDS, the bi-directional data flow is implied, as shown in Figure 10 below.

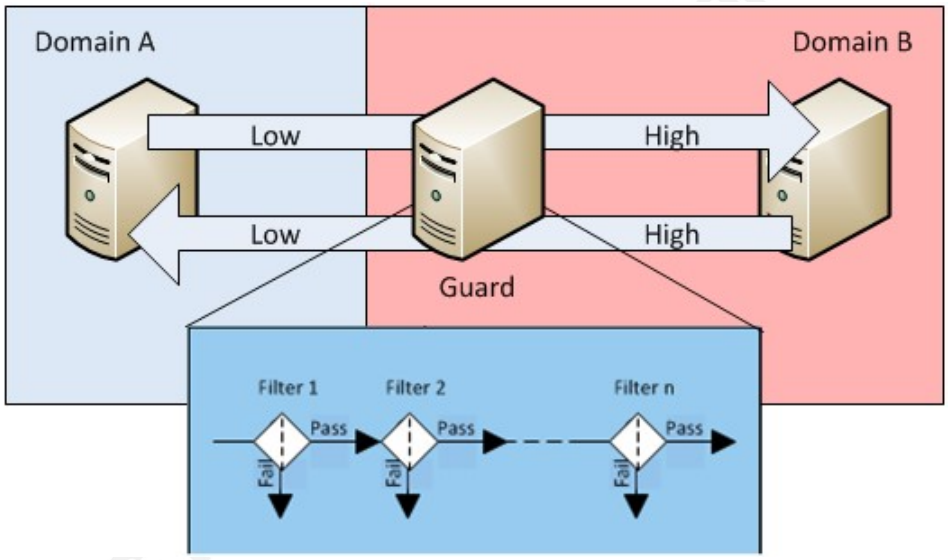

Figure 10 – Bi-directional Guard [Source: DND]

While firewalls examine individual packets at the TCP/IP layer, a guard's deep content inspection is capable of completely assembling multiple packets into a message inside a proxy environment. The message is then filtered for content (dirty/clean word checks) and embedded files, such as attachments and macros. Message handling instructions are attached to each filter, enabling the guard to pass the message on, redirect it, drop it, or even quarantine the packets for forensic review [Maney, 2004, p.6].

Although guards can usually handle hundreds of file types, the risk of steganography is ever-present. Images, video and audio files can all be modified to covertly pass data that would otherwise be flagged. In practice, guards mitigate this risk

by translating media files into a common uncompressed file type before inspection. While this solution cannot always prevent data spills, it is successful at removing malicious code and macros.

4. Critical Considerations

4.1. Credentials for Access Solution

Identity, Credential and Access Management (ICAM) is essential to an effective access CDS. In the isolation approach, each domain is responsible for maintaining a list of user credentials and authentication process. An access solution such as KVM switching or partitioned workstations may not be capable of simultaneously storing multiple unique credentials (i.e. two or more PKI certificates). While there are solutions to storing multiple credentials in a single container, such as the Oracle Wallet Manager, the criteria for an access solution precludes transfer of data between domains. Common credentials across multiple domains is also the wrong approach, since if any one domain is breached, all domains to which the user has access are compromised. For this reason, most government agencies have moved towards two factor authentication models with a shared token, such as a common access card with unique passwords per domain. This ensures one credential is shared across domains proving identity, while the domains maintain unique authentication measures.

4.2. Data Transfer Solution Review Process

4.2.1. Human Review

The simplest and most versatile review process involves a human operator. This is most useful in environments characterized by dynamic and unstructured data flows that have a low volume and/or speed of traffic. Most practical implementations rely on two operators to authenticate data in order to mitigate judgement inconsistencies over time and enforce accountability. John Woodward described the human review roles in terms of guard operators, which perform sanitization duties, and security watch officers, which approve or deny high-to-low transfers [Woodward, 1979, p.323].

Of course, the addition of a second pair of eyes further reduces throughput. To compensate, the low-to-high review is often omitted entirely in practice, as detailed in

Figure 11 below. Note the inherent risks of low-to-high malware and high-to-low data spills in this particular design.

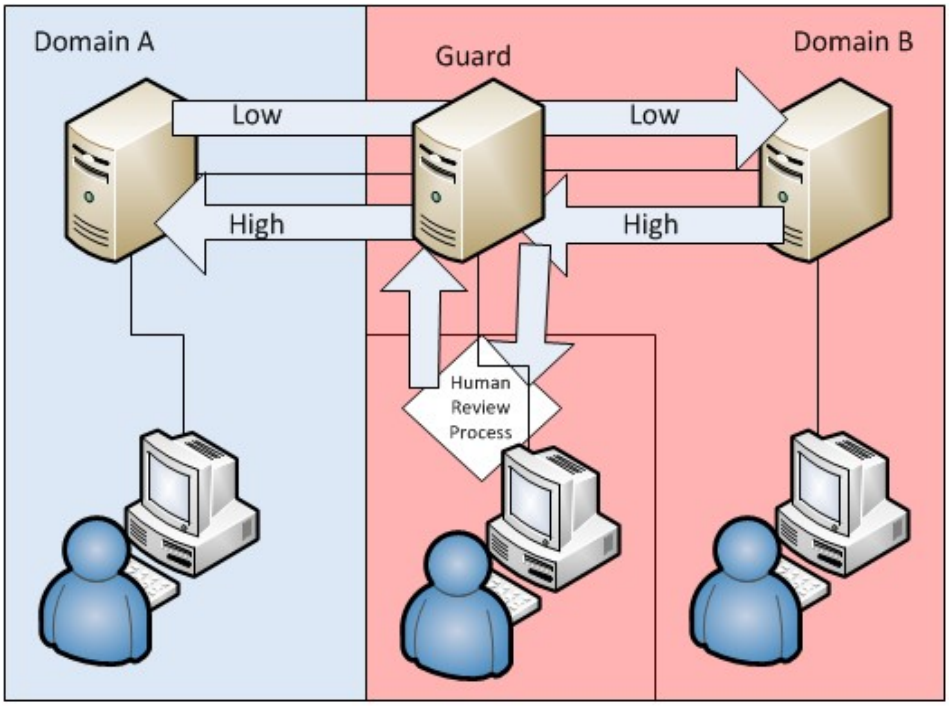

Figure 11 – Human Review (Source: DND [9])

4.2.2. Automated Review

To address static environments with structured, high volume and traffic data flows we turn to automated guards. This review model is fast, scalable and consistent. On the other-hand, it is also easy to circumvent most automated content inspection processes by a human client. Dirty word checks can be defeated with little imagination, and steganography in file types such as images, audio, and video can be virtually impossible to catch [Maney, 2004, p.8]. In practice, those file types may be dropped or re-formatted to prevent passage of malicious macros and/or disclosure of unauthorized information.

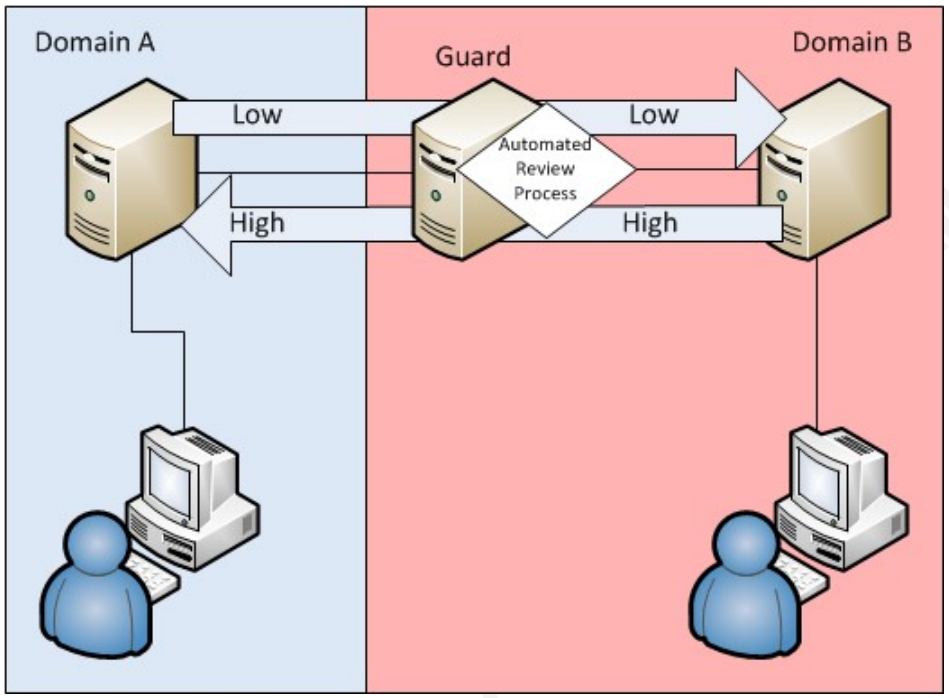

Figure 12 – Automated Review [Source: DND]

4.2.3. Hybrid Review

The ideal model for a dynamic, high volume and high traffic, unstructured data environment, the hybrid blends the best features of both human and automated review. Rather than outright reject content that does not pass the filters, this model can send select data transfer requests, such as images and flagged dirty words, to a human operator. The remaining common requests are reviewed and handled in near-real time. The balance of human to automated review can be tweaked overtime to develop a streamlined solution specific to the domains' data flows.

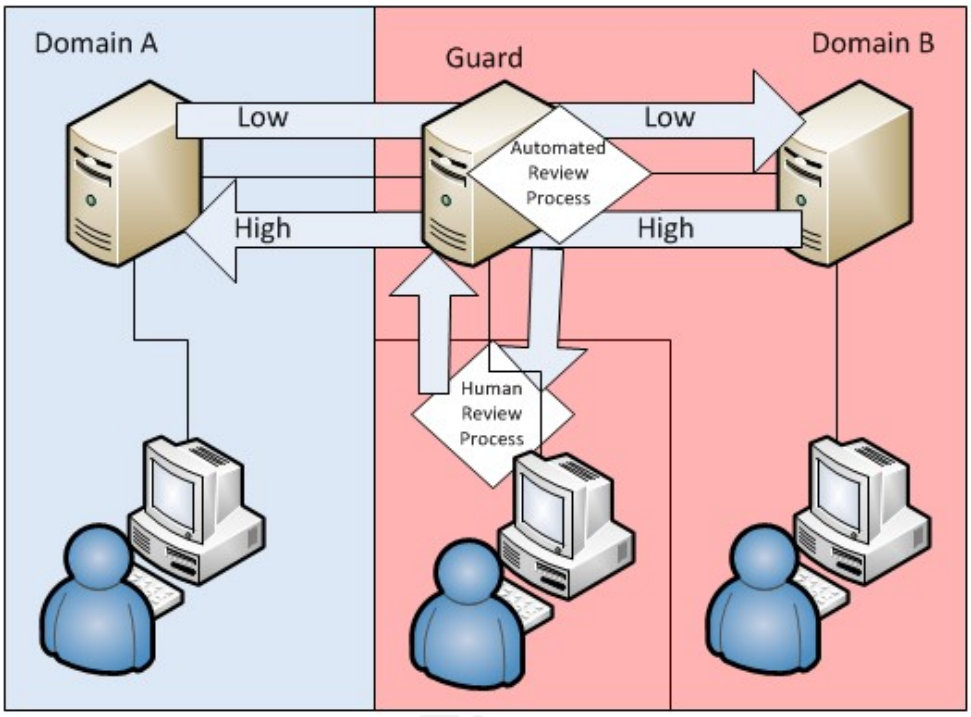

Figure 13 – Hybrid Review [Source: DND]

4.3. Accreditation

The United Cross Domain Services Management Office (UCDSMO, formerly the UDCMO) is currently the only organization that manages accreditation of CDSs. First established in 2006, the office manages all cross-domain initiatives across the Department of Defence (DoD) and intelligence communities [Takai, 2012]. The UCDMSO also provides a baseline list of validated solutions rated for deployment processing environments as classified by the Defense Intelligence Agency (DIA). The environments are defined as either Top Secret SCI (Sensitive Compartmented Information) and Below Interoperability (TSABI) or Secret and Below Interoperability (SABI). A selection of CDSs from the 2013 UDCMO baseline list is included in the appendix.

For CDSs to be employed in a TSABI environment, The Director of Central Intelligence Directive 6/3 (DCID 6/3) breaks down the approved domain interfaces by Protection Levels 1 through 5, each of which has hundreds of security controls that must be met in the storage, processing and communication of data [DCID 6/3].

SABI standards are approved via the Defence Information Security Agency's (DISA) Risk Decision Authority Criteria (RDAC) [DISA].

While the baseline list and accreditation levels are a good starting point, more consideration is required to ensure selection of the right product or solution. In most cases, the UDCMSO simply conducts oversight while the vender self-checks their solutions against compliance standards. As such, the test conditions may vary significantly from one accreditation to the next. Independent testing in dedicated lab environment is necessary to ensure a common standard is met.

5. Conclusion

While many of the UCDSMO baselined solutions are nationally controlled for government and defense use, the general concepts can be applied to commercial networks. In fact, as the IT industry continually moves towards additional features and improved functionality in common network devices, more commercial options are becoming available. An organisation may accept a standard guard, or decide on application aware firewalls paired with a deep-content inspection enabled gateway as an acceptable solution. The appropriate cross-domain solution will take into account the economic impact and accepted risk in addition to the information environment and associated security policies.

As such, the architectures as described in this paper are hardly exclusive. Nonetheless, they should serve as guide to the general types, flavors and considerations behind the CDSs in use today.

6. References

[1] Adder Technology (N.D.). *Secure KVM switching.* Retrieved from Amplicon Benelux website: http://www.ampliconbenelux.com/docs/secure-kvm-switching-white-paper.pdf

[2] Bell, D. E. (2005). *Looking back at the Bell La Pudula model*. Retrieved from: http://www.acsac.org/2005/papers/Bell.pdf

[3] Biba, K. J. (1975). *Integrity considerations for secure computer systems* (522B). Retrieved from Mitre Corporation website: http://seclab.cs.ucdavis.edu/projects/history/papers/biba75.pdf

[4] Calloni, B. (2011). *Connecting the RDAC to its environment*. Cherry Hill, NJ: Lockheed Martin

[5] Chen, P. (2010). *Multilevel secure systems and cross domain solutions: challenges and solutions*. Salt Lake City, UT: IEEE Systems and Software Technology Conference.

[6] Committe on national Security Systems. (2010). *National information assurance glossary* (CNSSI 4009). Retrieved from http://CNSSI 4009, National Information Assurance Glossary

[7] CrowdStrike (2015). *Virtualized environment neglected operations manipulation*. Retrieved from: http://venom.crowdstrike.com/

[8] Defence Information Security Agency (N.D.) *Connection process guide – cross domain solutions*. Retrieved: http://www.disa.mil/Network-

Services/Enterprise-Connections/Connection-Process-Guide/DISN-Service-Appendices/Cross-Domain-Solutions

[9] Department of National Defence (2015). *Information brief: cross domain and caveat solutions* (internal draft).

[10] Director of Central Intelligence Directive 6/3 (N.D.). *Protecting sensitive compartmented information within information systems.* Retrieved from: https://www.fismacenter.com/DCID%206_3%20Appendices.pdf

[11] Durante, R. J., & Woodruff, J. C. (2012). *SecureView: Government/industry collaboration delivers improved levels of security, performance, and cost savings for mission-critical applications* (88ABW-2012-6533). Air Force Research Laboratory.

[12] Kissel. R et al (2014) *Guidelines for media sanitization (800-88)*. Retrieved from National Institute of Standards and Technology website: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-88r1.pdf

[13] Kuhn, M. G. (2004). *Electromagnetic eavesdropping risks of flat-panel displays*. Retrieved from the University of Cambridge website: http://www.cl.cam.ac.uk/~mgk25/pet2004-fpd.pdf

[14] Shanker, T. (2010, July 8,). *Loophole may have aided theft of classified data*. The New York Times. Retrieved September 25, 2015, from http://www.nytimes.com

[15] Maney, C. (2004). *Security issues when data traverses information domains: do guards effectively address the problem?* Retrieved from SANS reading room website: https://www.sans.org/reading-room/whitepapers/assurance/security-

issues-data-traverses-information-domains-guards-effectively-address-problem-1418

[16] Public Works and Government Services Canada. (2015). *Security levels - industrial security program.* Retrieved from http://iss-ssi.pwgsc-tpsgc.gc.ca/outils-tools/ns-sl-eng.html

[17] Roullier, S (2004). *Virtual LAN security:weaknesses and countermeasures.* SANS Institute reading room: https://www.sans.org/reading-room/whitepapers/networkdevs/virtual-lan-security-weaknesses-countermeasures-1090

[18] Scott, A. (2015). *Tactical data diodes in industrial automation and control systems*. Retrieved from SANS Institute website: http://www.sans.org/reading-room/whitepapers/firewalls/tactical-data-diodes-industrial-automation-control-systems-36057

[19] Takai, T. M. and Tarasiuk, A. (2011). *Use of unified cross domain management office (UCDMO) baseline cross domain solutions (CDSs).* Retreived from semantic community website: http://semanticommunity.info/DoD_Chief_Information_Officer/Unified_Cross_Domain_Management_Office_(UCDMO)#Mission

[20] Tanenbaum, A. S. (1989). Computer Networks*. New Jersey: Prentice-Hall. ISBN 0-13-166836-6*.

[21] Sitnica, A (2014). *Case study: energy and utilities defense response based on 2014 attack pattern statistics.* Retrieved from SANS Institute website: https://www.sans.org/reading-room/whitepapers/scada/energy-utilities-defense-response-based-2014-attack-pattern-35657

[22] Woodward, J (1979). *Applications for mulitlevel secure operating systems*. Bedford, MA: The Mitre Corporation