DCSS Cyber Assessment Report

Mike Sues
Cord3 Innovation

Glen Henderson, Alan Magar and Alan Clason
Cord3 Innovation

Prepared By:
Cord3 Innovation
464 Besserer St.
Ottawa, Ontario K1S 5N4

DRDC Project Number: 3775-2005-15bs (CST C4ISR 3)
PWGSC Contract Number: W87714-08FE01/Bell
Contract Technical Authority: Daniel Charlebois, Defence Scientist.

Disclaimer: The scientific or technical validity of this Contract Report is entirely the responsibility of the Contractor and the contents do not necessarily have the approval or endorsement of the Department of National Defence of Canada.

Contract Report
DRDC-RDDC-2016-C102
March 2015

IMPORTANT INFORMATIVE STATEMENTS:

Aviation Management Interoperability for Emergency Response and Recovery Project, # CSSP-2014-CP-2005, was supported by the Canadian Safety and Security Program which is led by Defence Research and Development Canada's Centre for Security Science (DRDC-CSS), in partnership with Emergency Management British Columbia (EMBC). The project was led by DRDC-CSS.

Canadian Safety and Security Program is a federally-funded program to strengthen Canada's ability to anticipate, prevent/mitigate, prepare for, respond to, and recover from natural disasters, serious accidents, crime and terrorism through the convergence of science and technology with policy, operations and intelligence.

DCSS Cyber Assessment Report

Prepared by: Mike Sues,
Cord3 Innovation

Date: March 31, 2015

Revision: Final v1.0

Confidentiality

This document is UNCLASSIFIED.

Disclaimer

A number of sections of the report are comprised of material contributed by multiple authors, most notably members of the Bell Canada / Cord3 Innovation team

Authors

Cord3 Innovation Team	Role
Glen Henderson	Lead System Architect
Alan Magar	Information Security Analyst
Alan Clason	Quality Assurance Specialist

Review

Cord3 Innovation Team	Role
Will Coxon	Technical Project Manager

Executive Summary

DCSS is a *security overlay*: a set of interconnected services that communicate through the exchange of messages on top of an existing network deployment. Any network security or application security based environment can be enabled for data-centric protection without the need to remove or de-emphasize existing security protections. In this way, an environment protected by DCSS retains the existing safeguards that were in place prior to the deployment of the security overlay. Additionally, DCSS itself is able to leverage existing physical, administrative, network and application safeguards as part of its deployment profile. Most significantly, the deployment of DCSS does not change the underlying accreditation of the target network by modifying the security architecture that was initially certified.

This study examined the architecture of the security overlay represented by the DCSS components and how they may be compromised or used in the context of the CWIX 2015 environment to compromise other partners or derive information about other partner's activities. For each attack scenario found in the study procedures to test for the existence of the vulnerability and its impact were also developed along with the means to mitigate each issue.

In this way we have been able to enumerate the potential attack surface of the DCSS when deployed in the context of the CWIX 2015 environment and provide a means of both testing and resolving these issues to decrease the attack surface and operational risk to partners employing the DCSS security overlay.

The next step in this process should be to execute the developed procedures, assess the risk of each scenario, apply the resolution and then validate the issue has been fixed by regression testing the mitigation steps.

Table of Contents

Table of Figures

NO TABLE OF FIGURES ENTRIES FOUND.

1. Introduction

This document provides information in support of the Canadian Data Centric Security Solution (DCSS) program participation at the Coalition Warfighter Interoperability Experiment to be held in June 2015 (CWIX2015). At this experiment, the DCSS program team will be delivering new capabilities that build upon and extend the Data Centric Security (DCS) information protection model. Specifically, these extensions will enable the exchange of information assets between national partners in a coalition environment in a manner that adheres to the principles of a DCS-based information protection architecture. These principles include:

1. The ability to apply national and cross coalition information handling policy decisions to individual assets;

2. The ability to ensure that information assets are not released to requesting users or nations unless explicitly permitted by the national security policy; and

3. The ability to record the application of security policy decisions in a trusted and tamper-resistant audit store in support of incident analysis and forensic investigation.

This document describes the needed enhancements and extensions to the current DCS information protection model that will support cross coalition information sharing and that will be demonstrated at the CWIX2015 experiment. That is, this document will identify how DCS-based cross coalition information exchange can be achieved by leveraging and combining:

- The trust that is established as part of the standard operating procedure for standing up a coalition network environment; and

- The trust that is established within a national network when information protection practices based on data-centric security principles are used to secure information assets.

The DCS protection model provides a degree of trust by ensuring that information assets are responsibly protected at a national level. The cross coalition trust model builds upon this trust to enable information sharing across partner nations in a coalition environment. This layered approach to leveraging established trust is the principle upon which DCS cross coalition information exchange is founded.

1.1. Document Purpose

This document presents vulnerability test cases to confirm designed security controls or possible gaps in security controls.

The CWIX2015 technology target is a subset of the overarching DCS cross coalition information protection model that is intended to:

1. Validate the trust model design decisions;
2. Provide evidence in support of the viability of the DCS-based approach for information exchange in a coalition environment;
3. Identify new, or refine existing, requirements for the components in the proposed solution; and
4. Gather evidence in support of the acceptability of the proposed cross coalition information exchange capability.

This document will serve as the baseline definition for the technology target that will be delivered for the CWIX2015 experiment. This document provides:

- High level architectural description of the new DCS cross coalition capability that will be demonstrated;
- Identification of the enhancements to the existing national DCS model needed to support the cross coalition capability; and
- A description of how the new cross coalition DCS model will be deployed as an overlay to existing coalition network and application infrastructures.

This document identifies the information standards that have a bearing on the deployed solution including:

- Data format standards for information assets;
- Security labelling standard for applying security metadata to information assets;
- Security standards for the application of cryptographic protection on information assets; and
- Networking protocols used to exchange information assets at a national level and between coalition networks.

Additionally, this document provides a linkage between the technology target and the CWIX2015 exercise quality assurance, testing and scenario development documentation. Specifically, the two main support documents that will present the findings from the CWIX2015 exercise are:

1. The *CWIX2015 Test Plan* **[Reference 1]** that will define the practices that will be used to test, demonstrate and evaluate this capability; and

2. The *CWIX2015 Trial Report* **[Reference 2]** that will present the results of the exercise both in terms of the results gained from CWIX2015 participation and an assessment of the technology in the context of the overarching exercise scenarios.

1.2. Document Audience

The intended audience for this document are security practitioners that have an interest in DCS information protection and have a specific need to build or evaluate the proposed cross coalition information exchange model. This includes security professionals with a need for:

1. A description of the DCS cross coalition capability and the CWIX2015 technology target that will serve to prove the viability of DCS-based cross coalition information exchange; and

2. An understanding of the design, development, and testing requirements that will lead to the technology target deployment at the CWIX2015 exercise.

1.3. Data Centric Security Program History

The DCS program has been assigned the mandate to develop and transition DCS concepts towards the eventual target of operational deployment. It derives much of its information protection architecture from research conducted by Defence Research and Development Canada (DRDC) and the Secure Access Management for a Single Operational Network (SAMSON) technology demonstration. The goal of this research was to leverage DCS capabilities to reduce the need for multiple SECRET networks through community of interest (COI) separation. Ultimately, the intent was to rationalize the DCS architectural specification in order to define the next generation of secure networks through its ability to address the following four basic challenges in information assurance.

1. The specification, application and enforcement of a unified and holistic security policy across all information assets;
2. The restriction of transactions against information assets to only those communities with the policy right to perform those actions;
3. The ability to provide assurance that information is only released to those users that have a policy right to access it; and

4. The use of a trusted audit facility that records the details related to the release of information in a tamper-resistant form.

These four design priorities were enacted through the development of the DCS core solution components, which include the following:

- A centralized Policy Decision Point (PDP) that brokers access to information assets based on the security attributes of the requesting user, the security attributes of the targeted information asset, the operation requested on the asset and any environmental conditions that impact the decision-making process (e.g. time of day, location, user role);

- Application aware Policy Enforcement Points (PEPs) that intercept operations on information assets and ensure that the requested action is in compliance with the security policy. Vetting the action against the security policy of the PDP derives this assurance;

- Asset-level Cryptographic Transformation Service (CTS) protection where each asset is uniquely encrypted with its own symmetric key; a key that is only accessible to the PEP and only applied when the requested operation is compliant with the security policy; and

- A Trusted Audit Service (TAS) that creates a record of each requested action on the information asset, the decision that was made to permit or deny the action and a rationale for the decision that was made.

These DCS capabilities were delivered as a security overlay to existing application, network and information architectures. This national DCS information protection model is presented in the context of the cross coalition information exchange model in section 2.1:Data Centric Security within National Networks.

1.3.1. The CWIX2014 Exercise

In June 2014, the DCS program participated in the Coalition Warfighter Interoperability Experiment (CWIX2014). The intent of this experiment was to gather evidence to prove the viability of a new mechanism to exchange information between organizational domains. The challenge presented in existing information architectures is that organizations typically select their own preferred security-labelling format. The lack of a common approach to defining inter-organization security labels and associating information assets with those labels significantly hinders the ability to exchange information in a controlled manner. In contrast, when a common labelling mechanism is in place for the exchange of data assets, information can be shared between organizations (and handled and processed) according to mutually agreed to memoranda of understanding (MOU). The goal of this experiment was to create, deploy and test an implementation of a cross organization information exchange solution that would support the use of MOUs between organizations through a common labelling standard for information in transit.

The programmatic and technical objectives for the CWIX2014 experiment were:

- To prove the viability of a cross organization information exchange mechanism where:

    - Security metadata from assets sent from a source domain are used to create a transport security label;
    - Information assets and their associated transport security labels are placed in a container for transmission between domains;
    - At the receiving domain, the security metadata on the transport container is validated against the security metadata on the original asset to ensure that it was not modified in transit;

- To support a set of security labelling formats on the original information assets including:

    - Property-based security attribute files (e.g. Titus Document Classification);
    - XML Security Policy Information File (SPIF);
    - NCIA Security Label Attributes (SLAB) format;

- To exchange information using unmodified client and server software while supporting the capability to encapsulate information assets into a transport container in transit;

- To demonstrate cross coalition information sharing using email as the information exchange method; and

- To gather supporting information in the form of log files to demonstrate the effectiveness of the proposed solution.

This experiment included participation by the following nations and organizations:

1. United States;
2. United Kingdom;
3. Germany;
4. Finland; and
5. Netherlands.

While the detailed test results are documented in the *SD003 CWIX2014 Test Report* **[Reference 3]**, the following observations can be made about the results from the experiment.

- All solution elements performed as expected.

- The approach to hardware and software deployment, configuration and operation allowed the experiment to take place without any technical difficulties;

- All testing was successful: no failures were encountered during the test cycle procedures; and

- The solution was successfully demonstrated in operation to partners and VIP representatives.

1.4. CWIX2015

Based on the success of the CWIX2014 exercise, CWIX2015 will continue to evolve the concept of cross coalition information exchange by introducing some new capabilities and enacting scenarios that include a full end-to-end DCS-based information exchange between partner nations.

This cross coalition capability is intended to achieve the following objectives:

1. To prove the viability of the DCS cross coalition trust model and information exchange framework as proposed in this document;

2. To validate the over-arching design decisions that form the core of the DCS cross coalition information exchange functional model;

3. To observe the behaviour of the solution in an operational context with respect to:
   a. Ease of implementation and deployment;
   b. Robustness of the security protection mechanisms;
   c. Operational performance;

4. To observe the proposed solution from a functional perspective and assess the degree to which application support, information exchange and integration with established processes can be done seamlessly and transparently;

5. To gather evidence to justify further exploration, examination and evaluation of a DCS cross coalition model; and

6. To mitigate risks associated with transitioning this capability to an operational environment.

1.5. Assumptions

This report makes the following assumptions:

- Security Classification & Clearance – It is assumed that for the CWIX2015 experiment all networks will be at the same classification level and all users will have the proper clearance to access networks at that classification level. Specifically, for CWIX2015 it is assumed that all information exchange occurs at the SECRET level and that all users are cleared to this level; and
- Original Security Label – It is assumed that for CWIX2015 all participants will wish to retain the original security metadata on information assets that are transferred between domains.

1.6. Glossary

In order to avoid confusion and improve clarity, a number of key terms will be defined within this section of the report. Specifically, this section will define how the following terms will be used within the context of this report:

- Community Of Interest (COI) – The term COI is used throughout this report. In the context of this report the term is applied to both users and data. In terms of users, the term denotes membership within the COI. In terms of assets, the term is used to restrict access to a specific COI. The term COI is synonymous with caveats and warning terms;

- Domain – Within this report the terms domain (short for security domain) and national network are used interchangeably. In reality, a security domain may be comprised of a number of networks all governed by the same security policy. However, in the context of CWIX2015, and this report, the term domain and national network are synonymous as each of the national networks will have a unique security policy;

- Security Label – A security label is used primarily to denote the sensitivity of an information asset, but may contain additional security metadata. Security labels are sometimes referred to as confidentiality labels. The term security label will be used throughout this report;

- Security Labelling Policy - A security *labelling policy* in the context of this report refers to the syntax and semantics for the construction of security labels on information assets. For example, a new classification value denotes a change to the national security labelling policy; and

- Trust - Trust has different meanings in different contexts. In the context of this document trust refers to inherited trust as opposed to provable trust. The trust that is achieved within a domain through the use of DCS is being extended cross-coalition. Specifically, the information protections provided in the originating domain are replaced by equivalent information protections during transit between domains. Consequently, when the recipient domain receives the asset they can be confident that it has not been altered or disclosed. However, in this information exchange the originating domain is ultimately trusting the recipient domain to handle the information that was exchanged in a responsible manner commensurate with the assigned security label and in compliance with the eMoU in place between the two nations. It should be noted that this inherited trust improves upon existing paper-based mechanisms and will drastically improve information sharing between coalition nations by automating the process.

1.7. Document Outline

This document is structured into the following sections:

- Section 2:Scope
- Section 3:Threat Scenarios
- Section 4:Test Procedures;
- Section 5: Conclusions and Recommendations;
- Section 6: References;

2. Scope

This scope of this Cyber assessment focused on the insider threats specific to the environment modelled in the CWIX2015 exercise. This constitutes threat agents that are partners in the coalition who, though work together as part of the coalition, have varying levels of trust in one another and possess reasonably advanced Computer and Network Exploitation capabilities. Their goal may be to derive and/or collection information about other partner's activities, change their perception of communications by misleading them or even deny them access to communications. Their methods of attack are expected to employ their connectivity with the partners and in some cases, use their DCSS servers and policies to their advantage.

3. Threat Scenarios

This section provides a description of the Threat Scenarios and Agents that were considered in the development of the test cases as well as a justification for their inclusion.

3.1. Scenario Goals

The goal of the scenarios are to employ the access through the trust level in the coalition and use of Data-Centric Security Services (DCSS) to attempt to gain intelligence about other partners or gain access to partner's domain networks.

3.2. Threat Agents

In a coalition environment it is expected that, even though partners will have signed memorandums of agreement as to how they will work together and perhaps interact, not all will or should be trusted equally by their partners. Some will use this partnership and level of trust to conduct intelligence operations against their partners to improve their position in the coalition or for political purposes. It is a fair assumption that each partner in the coalition will have a Computer Network Exploitation capability of varying levels of sophistication and will use this capability against their partners.

Organizations with such Computer Network Exploitation capabilities are typically well-funded and motivated though clearly within a coalition environment there will be partners with better funding and capabilities. Moreover, some, but not all partners will possess an industry base which the other partners utilize for their hardware and software requirements and it is a fair assumption that partners will use their industry base in a cooperative manner to insert engineered vulnerabilities or software implants to provide remote access and data exfiltration capabilities. We do not consider this form of access through their industry base in the development of the threat scenarios but instead scenarios in which non-engineered vulnerabilities are exploited.

3.3. Scenarios

Follows are each of the threat scenarios attacking key DCSS assets or data flows along with recommended remediation techniques.

3.3.1. Clients Configured to Bypass Policy Enforcement

Client data in a domain environment must traverse Policy Enforcement Points so that the domain policy can be confirmed for information releasability to the COI at the message's classification level. An insider with the ability to change their client software configuration could instead bypass the Policy Enforcement Points and communicate directly with the end servers (e.g. email or XMPP servers) to attempt to bypass policy checks.

However, since emails on the server are encrypted the client accessing the server directly would only be able to access encrypted emails, which could reveal some information about traffic patterns and volumes. Note that since the destination policy and COI verification is performed in the destination domain, an attacker would not be able to bypass these checks by sending email directly to their own email server, except if they were sending it within their own domain.

Note that configuration and hardening of the clients to restrict these sort of changes is outside of the scope of the DCSS though this can be mitigated by encrypting the messages on the servers so that only access through the Policy Enforcement Points would decrypt the messages.

3.3.2. Information Leakage through Email Message Attachments

Though the user applies a security label to the message it is possible that additional information has been hidden in fields of the email message attachment format that are of a higher classification than the security label. This could be used by a software implant as a method of exfiltrating data to another partner, which upon receipt at the destination or collected enroute, the hidden data is extracted.

Note that neither Data Loss Prevention nor protection from covert channels is a design requirement for the DCSS nor is the DCSS considered nor designed to be a Cross-Domain Solution.

3.3.3. Information Leakage through Email Message Content

Though the user applies a security label to the message it is possible that additional information has been hidden within the content or format of the email message that is of a higher classification than the security label. This could be used by a software implant as a method of exfiltrating data to another partner, which upon receipt at the destination or collected enroute, the hidden data is extracted.

Note that neither Data Loss Prevention nor protection from covert channels is a design requirement for the DCSS nor is the DCSS considered nor designed to be a Cross-Domain Solution.

3.3.4. Information Leakage through SMTP Headers/Protocol

Though the user applies a security label to an email message it is possible that additional information has been hidden within the SMTP headers that are under the control of the email client or the domain's email server that is of a higher classification than the security label. This could be used by a software implant as a method of exfiltrating data to another partner, which upon receipt at the destination or collected enroute, the hidden data is extracted.

Note that neither Data Loss Prevention nor protection from covert channels is a design requirement for the DCSS nor is the DCSS considered nor designed to be a Cross-Domain Solution.

3.3.5. Information Leakage through Chat Room Names

Chat room names hosted on XMPP servers could be used to exfiltrate data between domains, different COI's or even from a higher to lower security label. This could be used by a software implant to exfiltrate data to another partner, within different COI's or even within attacks in the same domain but with a lesser access to the chat rooms.

Note that neither Data Loss Prevention nor protection from covert channels is a design requirement for the DCSS nor is the DCSS considered nor designed to be a Cross-Domain Solution.

3.3.6. Information Leakage through Chat Messages

Though the user applies a security label to the message it is possible that additional information has been hidden within the content or format of the chat message that is of a higher classification than the security label. This could be used by a software implant as a method of exfiltrating data to another partner, which upon receipt at the destination or collected enroute, the hidden data is extracted.

Note that neither Data Loss Prevention nor protection from covert channels is a design requirement for the DCSS nor is the DCSS considered nor designed to be a Cross-Domain Solution.

3.3.7. Information Leakage through the XMPP Protocol

Though the user applies a security label to a message it is possible that additional information has been hidden within the XMPP protocol that are under the control of the email client or the XMPP server that is of a higher classification than the security label. This could be used by a software implant as a method of exfiltrating data to another partner, which upon receipt at the destination or collected enroute, the hidden data is extracted.

Note that neither Data Loss Prevention nor protection from covert channels is a design requirement for the DCSS nor is the DCSS considered nor designed to be a Cross-Domain Solution.

3.3.8. Email Probing From a Domain to Derive Security Policy Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at the destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the inbound message based on the policy (e.g. Classification), a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different security labels could be sent to users in another partner to enumerate a domain's security level policies for each of their users.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.9. Email Probing From Outside of a Domain to Derive Security Policy Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at the destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the inbound message based on the policy (e.g. Classification), a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different security labels could be sent from an external connection to the Domains to users in a partner to enumerate a domain's security level policies for each of their users.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.10. Email Probing From a Domain to Derive COI Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the inbound message based on the COI (e.g. not a member), a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different COI's could be sent to users in another partner, collecting the rejection responses to enumerate a domain's COI policies for each of their users.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.11. Email Probing From Outside of a Domain to Derive COI Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the inbound message based on the COI (e.g. not a member), a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different COI's could be sent from an external connection to users in a partner, collecting the rejection responses to enumerate a domain’s COI policies for each of their users.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.12. Email Probing From a Domain Using Email Address Spoofing to Derive Security Policy Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the policy (e.g. Community of Interest) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different security labels could be sent to users in a domain with a source spoofed from another domain, collecting the rejection responses from another part of the network. The goal of this probing would be to enumerate a domain’s security level policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.13. Email Probing From Outside of a Domain Using Email Address Spoofing to Derive Security Policy Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the policy (e.g. Community of Interest) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different security labels could be sent from an external connection to users in a domain with a source spoofed from another domain, collecting the rejection responses from another part of the network. The goal of this probing would be to enumerate a domain’s security level policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.14. Email Probing From a Domain Using Email Address Spoofing to Derive COI Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the

COI (e.g. not a member) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails different COI's could be sent to users in a domain with a source spoofed from another domain, collecting the rejection responses from another part of the network. The goal of this probing would be to enumerate a domain's COI policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.15. Email Probing From Outside of a Domain Using Email Address Spoofing to Derive COI Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the COI (e.g. not a member) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails different COI's could be sent from an external connection to users in a domain with a source spoofed from another domain, collecting the rejection responses from another part of the network. The goal of this probing would be to enumerate a domain's COI policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent,

3.3.16. Email Probing From a Domain with SMTP Header Spoofing to Derive Policy Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the policy (e.g. level of security label) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different security labels could be sent to users in a domain with a source and SMTP routes spoofed from another domain, collecting the rejection responses as they are sent back to the SMTP routes. The goal of this probing would be to enumerate a domain's security level policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.17. Email Probing From Outside of a Domain with SMTP Header Spoofing to Derive Policy Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the policy (e.g. level of security label) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different security labels could be sent from an external connection to users in a domain with a source and SMTP routes spoofed from another domain, collecting the rejection responses as they are sent back to the SMTP routes. The goal of this probing would be to enumerate a domain's security level policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.18. Email Probing From a Domain with SMTP Header Spoofing to Derive COI Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the COI (e.g. not a member) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different COI's could be sent to users in a domain with a source and SMTP routes spoofed from another domain, collecting the rejection responses as they are sent back to the SMTP routes. The goal of this probing would be to enumerate a domain's COI policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.19. Email Probing From Outside of a Domain with SMTP Header Spoofing to Derive COI Information of a Partner Domain

When emails are sent to a COI with a security label a check is performed at a destination domain to confirm the recipient is able to receive it based on the COI and security label. If the Border Policy Enforcement Point rejects the message based on the COI (e.g. not a member) a message is sent back to the sender with a high level description of the reason for rejection. Multiple emails with different COI's could be sent from an external connection to users in a domain with a source and SMTP routes spoofed from another domain, collecting the rejection responses as they are sent back to the SMTP routes. The goal of this probing would be to enumerate a domain's COI policies for each of their users for the spoofed domain.

This can be mitigated through monitoring of inbound messages and matching up policy decisions (i.e. accept or reject) as well as the destination to which rejection messages are sent.

3.3.20. Email Forwarding With a High Label Sent With a Lower Level

An email with an initial higher security label could be relabeled and forwarded with a lower level security to another user either on purpose or accidentally to disclose higher classified data to users with access to only lower levels.

This can be mitigated by cryptographically binding the original security label to the messages and basing policy decisions on the original label. Note that if this is not supported by the security labeling software then mitigation must take the form of user awareness training and procedural controls.

3.3.21. Email Forwarding To a Different Community of Interest

An email with an initial COI could be forwarded with a different COI at the same security level either on purpose or accidentally to disclose information to users beyond the initial COI.

This can be mitigated by cryptographically binding the original COI to the messages and basing policy decisions on the original label. Note that if this cannot be supported then mitigation must take the form of user awareness training and procedural controls.

3.3.22. Buffer Overflows in Security Label Processing on Policy Enforcement Points

The processing of messages and verification against policies on Policy Enforcement Points could suffer from buffer overflows that would allow an attacker to execute arbitrary code on the PEP and compromise the system or crash the PEP processing and cause a denial of service condition. As well, once compromised, the attacker could also modify the policies being applied the traffic flows and, even though message may be encrypted, collect traffic to derive information about the partner’s activities and user roles from Traffic Analysis.

This could be mitigated through code reviews and vulnerability analysis (i.e. fuzzing) of the Policy Enforcement Points DCSS interfaces.

3.3.23. Compromise of the Policy Server

A compromise of a domain’s Policy Server or supporting database would allow an attacker to compromise the confidentiality, integrity and availability of the policies served by this DCSS component. This level of access could allow an attacker to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data.

This can be mitigated through code reviews and vulnerability analysis (i.e. fuzzing) of the Policy Domain Points DCSS interfaces as well as standard hardening practices applied to the platform.

3.3.24. Spoofing As National Policy Enforcement Points

Spoofing as a National Policy Enforcement Point would allow an attacker to enforce a different set of policies to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data. Note that messages will be encrypted so its use to collect data will not yield message information but volumes and traffic patterns (i.e. Traffic Analysis) which could yield information about partner activities and user roles.

Note it is not the role of the DCSS to stop spoofing attacks but controls should be in place at the data layer to account for such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.25. Spoofing As Border Policy Enforcement Points

Spoofing as a Border Policy Enforcement Point would allow an attacker to enforce a different set of policies to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data. Note that message will be encrypted so its use to collect data will not yield message information but volumes and traffic patterns (i.e. Traffic Analysis) could yield information about partner activities and user roles.

Note it is not the role of the DCSS to stop spoofing attacks but controls should be in place at the data layer to account that such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.26. Spoofing As the Policy Server

Spoofing as a Policy Server would allow an attacker to deliver a different set of policies to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data.

Note it is not the role of the DCSS to stop spoofing attacks but controls should be in place at the data layer to account that such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.27. Man in the Middle Attacks against National Policy Enforcement Points

Conducting man in the middle attacks against the National Policy Enforcement Point would allow an attacker to collect message as well as enforce a different set of policies to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data. Note that messages will be encrypted so its use to collect data will not yield message information but volumes and traffic patterns (i.e. Traffic Analysis) could yield information about partner activities and user roles.

Note it is not the role of the DCSS to stop man in the middle attacks but controls should be in place at the data layer to account that such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.28. Man in the Middle Attacks between Domains

Conducting man in the middle attacks against Border Policy Enforcement Points would allow an attacker to collect message as well as enforce a different set of policies to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data. Note that messages will be encrypted so its use to collect data will not yield message information but volumes and traffic patterns (i.e. Traffic Analysis) could yield information about partner activities and user roles.

Note it is not the role of the DCSS to stop man in the middle attacks but controls should be in place at the data layer to account that such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.29. Man in the Middle Attacks against the Policy Server

Conducting man in the middle attacks a domain’s Policy Server would allow an attacker to view the policies as well as deliver a different set of policies to block more messages and conduct a denial of service attack or permit more messages to pass and attempt to leak additional data.

Note it is not the role of the DCSS to stop man in the middle attacks but controls should be in place at the data layer to account that such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.30. Man in the Middle Attacks against the Trusted Audit Service

Conducting man in the middle attacks against the Trusted Audit Service would allow an attacker to view syslog messages and derive DCSS activity information. Moreover, this could be used to block or insert false messages to affect the integrity of the audit data in the Trusted Audit Service.

Note it is not the role of the DCSS to stop man in the middle attacks but controls should be in place at the data layer to account that such attacks and provide assurance of endpoint origin. This can be mitigated through encryption and origin authentication controls at the data level to ensure that only the client and Policy Enforcement Point can decrypt data and provide assurances in the endpoints.

3.3.31. Downgrading Cryptographic Connections to National Policy Enforcement Points

Downgrading the cryptographic connections to National Policy Enforcement Points would allow an attacker to cryptanalytically attack and then view encrypted messages and their envelopes but not the contents of messages. This could be used to perform Traffic Analysis to derive information about partner activities and user roles based on traffic volumes and patterns.

This can be mitigated through reconfiguration of the services to prohibit support for weaker ciphers and security protocols.

3.3.32. Downgrading Cryptographic Connections to Border Policy Enforcement Points

Downgrading the cryptographic connections to Border Policy Enforcement Points would allow an attacker to cryptanalytically attack and then view encrypted messages and their envelopes but not the contents of messages. This could be used to perform Traffic Analysis to derive information about partner activities and user roles based on traffic volumes and patterns.

This can be mitigated through reconfiguration of the services to prohibit support for weaker ciphers and security protocols.

3.3.33. Downgrading Cryptographic Connections to the Policy Server

Downgrading the cryptography in connections to a Policy Server would allow an attacker to cryptanalytically attack and then view the requested policies.

This can be mitigated through reconfiguration of the services to prohibit support for weaker ciphers and security protocols.

3.3.34. Downgrading Cryptographic Connections to the Trusted Audit Service

Downgrading the cryptography in connections to the Trusted Audit Service would allow an attacker to cryptanalytically attack and then view syslog messages and derive some information about DCSS activities.

This can be mitigated through reconfiguration of the services to prohibit support for weaker ciphers and security protocols.

3.3.35. Spoofing Audit Records to Insert False Messages in the Trusted Audit Service

Spoofing false messages to the Trusted Audit Service would allow an attacker to affect the integrity of audit records used for diagnosis of issues and Incident Response.

This can be mitigated through encrypting audit records sent to the Trusted Audit Services and building in origin authentication controls at the data layer.

3.3.36. Information Disclosure through Trusted Audit Service Information Sharing Between Partners for Incident Response

The sharing of information from Partner’s Trusted Audit Services to conduct Incident Response activities would reveal sensitive information about the operation of the DCSS as well as the partner’s activities and patterns employing the DCSS.

This could be mitigated through procedural controls to either work through a Trusted Third Party or perform stringent sanitization on audit records shared with others partners.

4. Test Procedures

This section provides a high-level description of the test procedures to check for the presence of the vulnerabilities and assess the level of risk for each Threat Scenario documented in the previous section.

4.1. Test Procedures

The following are high level test procedures to check for the conditions required for the each attack scenario documented in the previous section.

4.1.1. Clients Configured to Bypass Policy Enforcement

1. Reconfigure the client software to point to the email/XMPP server instead of the National Policy Enforcement Point.
2. Start client and retrieve data from the server.
3. Review the data that could be retrieved by the client.

4.1.2. Information Leakage through Email Message Attachments

1. Create a Microsoft Word document and change the title in the document properties to contain a “secret” string.
2. Attach the Word document to an email and send it to an approved COI with an appropriate security label that represents a lower security label than the string in the document’s title.
3. Send email
4. Confirm the email and attached file were properly received at the destination.

4.1.3. Information Leakage through Email Message Content

1. Create an email for an approved COI with an appropriate security label.
2. Insert into the email message a string that represents data from a higher security.
3. Send the email.
4. Confirm the email was properly received at the destination.

4.1.4. Information Leakage through SMTP Headers/Protocol

1. Employ or develop an email client that permits manipulation of the SMTP headers sent by the client.
2. Create an email for an approved COI with an appropriate security label.
3. Modify the SMTP X-Mailer header to insert a string that represents data from a higher security.
4. Send the email.
5. Confirm the email was properly received at the destination.

4.1.5. Information Leakage through Chat Room Names

1. Start chat client and connect to the XMPP server at an approved security label/COI.
2. Create a new chat room with a name whose string is a higher classification that the approved label.

4.1.6. Information Leakage through Chat Messages

1. Start chat client and connect to the XMPP server at an approved security label/COI.
2. Create a new chat message containing data which represents a higher classification that the approved label.

4.1.7. Information Leakage through the XMPP Protocol

1. Employ or develop a chat client that permits direct manipulation of the XMPP XML message sent by the client.
2. Start chat client and connect to the XMPP server at an approved security label/COI.
3. Create a new chat message
4. Create and send a chat message.
5. Confirm the chat message was properly posted to the XMPP server.

4.1.8. Email Probing From a Domain to Derive Security Policy Information of a Partner Domain

1. Create an email for a person in another domain within an approved COI with a security label that is not approved for the person.
2. Send the email.
3. Review the resulting rejection message.
4. Repeat for each user in the other domain.

4.1.9. Email Probing From Outside of a Domain to Derive Security Policy Information of a Partner Domain

1. From an external connection outside of a domain create an email for a person in another domain within an approved COI with a security label that is not approved for the person.
2. Send the email.
3. Review the resulting rejection message.
4. Repeat for each user in the other domain.

4.1.10. Email Probing From a Domain to Derive COI Information of a Partner Domain

1. Create an email for a person in another domain who is not within an approved COI with a security label that is approved for the person.
2. Send the email.

3. Review the resulting rejection message.
4. Repeat for each user in the other domain.

4.1.11. Email Probing From Outside of a Domain to Derive COI Information of a Partner Domain

1. From an external connection outside of a domain create an email for a person in another domain who is not within an approved COI with a security label that is approved for the person.
2. Send the email.
3. Review the resulting rejection message.
4. Repeat for each user in the other domain.

4.1.12. Email Probing From a Domain Using Email Address Spoofing to Derive Security Policy Information of a Partner Domain

1. Employ an email client that permits the spoofing of source email addresses and ensure the email server will permit relaying spoofed emails.
2. Collect network traffic in between domains.
3. Create an email, spoofed from an address in another domain, for a person in another domain within an approved COI with a security label that is not approved for the person.
4. Send the email.
5. Collect the resulting rejection message.
6. Repeat for each user in the other domain.

4.1.13. Email Probing From Outside of a Domain Using Email Address Spoofing to Derive Security Policy Information of a Partner Domain

1. Employ an email client that permits the spoofing of source email addresses and an email server will permit relaying spoofed emails.
2. Collect network traffic in between domains.
3. From an external connection outside of a domain create an email, spoofed from an address in another domain, for a person in another domain within an approved COI with a security label that is not approved for the person.
4. Send the email.
5. Collect the resulting rejection message.
6. Repeat for each user in the other domain.

4.1.14. Email Probing From a Domain Using Email Address Spoofing to Derive COI Information of a Partner Domain

1. Employ an email client that permits the spoofing of source email addresses and ensure the email server will permit relaying spoofed emails.
2. Collect network traffic in between domains.

3. Create an email, spoofed from an address in another domain, for a person in another domain who is not within an approved COI with a security label that is approved for the person.
4. Send the email.
5. Collect the resulting rejection message.
6. Repeat for each user in the other domain.

4.1.15. Email Probing From Outside of a Domain Using Email Address Spoofing to Derive COI Information of a Partner Domain

1. Employ an email client that permits the spoofing of source email addresses and an email server will permit relaying spoofed emails.
2. Collect network traffic in between domains.
3. From an external connection outside of a domain create an email, spoofed from an address in another domain, for a person in another domain who is not within an approved COI with a security label that is approved for the person.
4. Send the email.
5. Collect the resulting rejection message.
6. Repeat for each user in the other domain.

4.1.16. Email Probing From a Domain with SMTP Header Spoofing to Derive Policy Information of a Partner Domain

1. Employ an email client that permits the spoofing of source SMTP routing headers and spoofed source email address and ensure the email server will permit relaying spoofed emails.
2. Create an email, spoofed from an address in another domain and spoofing the final SMTP routing header as the true source domain's SMTP server, for a person in another domain within an approved COI with a security label that is not approved for the person.
3. Send the email.
4. Review the resulting rejection message as it arrives at the domain's SMTP server.
5. Repeat for each user in the other domain.

4.1.17. Email Probing From Outside of a Domain with SMTP Header Spoofing to Derive Policy Information of a Partner Domain

1. Employ an email client that permits the spoofing of source SMTP routing headers and spoofed source email address and use an email server will permit relaying spoofed emails.
2. From an external connection outside of a domain create an email, spoofed from an address in another domain and spoofing the final SMTP routing header as the true source domain's SMTP server, for a person in another domain within an approved COI with a security label that is not approved for the person.
3. Send the email.

4. Review the resulting rejection message as it arrives at the domain's SMTP server.
5. Repeat for each user in the other domain.

4.1.18. Email Probing From a Domain with SMTP Header Spoofing to Derive COI Information of a Partner Domain

1. Employ an email client that permits the spoofing of source SMTP routing headers and spoofed source email address and ensure the email server will permit relaying spoofed emails.
2. From an external connection outside of a domain create an email, spoofed from an address in another domain and spoofing the final SMTP routing header as the true source domain's SMTP server, for a person in another domain within an approved COI with a security label that is not approved for the person.
3. Send the email.
4. Review the resulting rejection message as it arrives at the domain's SMTP server.
5. Repeat for each user in the other domain.

4.1.19. Email Probing From Outside of a Domain with SMTP Header Spoofing to Derive COI Information of a Partner Domain

1. Employ an email client that permits the spoofing of source SMTP routing headers and spoofed source email address and an email server will permit relaying spoofed emails.
2. Create an email, spoofed from an address in another domain and spoofing the final SMTP routing header as the true source domain's SMTP server, for a person in another domain within an approved COI with a security label that is not approved for the person.
3. Send the email.
4. Review the resulting rejection message as it arrives at the domain's SMTP server.
5. Repeat for each user in the other domain.

4.1.20. Email Forwarding With a High Label Sent With a Lower Level

1. Create an email to an approved COI and security label.
2. Send email.
3. Retrieve email at destination and forward to another COI and assign a lower security label.
4. Send email.
5. Confirm email was received at destination.

4.1.21. Email Forwarding To a Different Community of Interest

1. Create an email to an approved COI and security label.
2. Send email.

3. Retrieve email at destination and forward to another COI that is approved to receive the original security label.
4. Send email.
5. Confirm email was received at destination.

4.1.22. Buffer Overflows in Security Label Processing on Policy Enforcement Points

1. Identify input formats and network interfaces for the DCSS Policy Enforcement Points.
2. Develop a fuzzing model based on this input format.
3. Execute fuzzing cases and review server and service for unusual behaviour such as crashes or unusual log file entries.

4.1.23. Compromise of the Policy Server

1. Identify input formats and network interfaces for the DCSS Policy Server service.
2. Develop a fuzzing model based on this input format.
3. Execute fuzzing cases and review server and service for unusual behaviour such as crashes or unusual log file entries.
4. Run a Vulnerability Scanner against the Policy Server to identify any platform issues.

4.1.24. Spoofing As National Policy Enforcement Points

1. Setup a rogue National Policy Enforcement Point with a different security policy.
2. Employ a spoofing technique such as ARP poisoning or DNS redirection to masquerade as the true National Policy Enforcement Point.
3. Create and send an email to a COI with a security label that would not be approved by the true National PEP but is permitted by the rogue point.
4. Confirm the email was received at the destination.

4.1.25. Spoofing As Border Policy Enforcement Points

1. Setup a rogue Border Policy Enforcement Point with a different security policy.
2. Employ a spoofing technique such as ARP poisoning or DNS redirection to masquerade as the true Border Policy Enforcement Point.
3. Create and send an email to a COI with a security label that would not be approved by the true Border PEP but is permitted by the rogue point.
4. Confirm the email was received at the destination.

4.1.26. Spoofing As the Policy Server

1. Setup a rogue Policy Server with a different security policy.
2. Employ a spoofing technique such as ARP poisoning or DNS redirection to masquerade as the true Policy Server.
3. Create and send an email to a COI with a security label that would not be approved by the true Policy Server but is permitted by the rogue point.
4. Confirm the email was received at the destination.

4.1.27. Man in the Middle Attacks against National Policy Enforcement Points

1. Setup a rogue National Policy Enforcement Point with a different security policy but will also pass connections on to the true National Policy Enforcement Point, modifying only key elements of the policy.
2. Employ a man in the middle technique such as ARP poisoning or DNS redirection to masquerade as the true National Policy Enforcement Point.
3. Create and send an email to a COI with a security label that would not be approved by the true National PEP but is permitted by the rogue point.
4. Confirm the email was received at the destination.

4.1.28. Man in the Middle Attacks between Domains

1. Setup a rogue Border Policy Enforcement Point with a different security policy but will also pass connections on to the true Border Policy Enforcement Point, modifying only key elements of the policy.
2. Employ a spoofing technique such as ARP poisoning or DNS redirection to masquerade as the true Border Policy Enforcement Point.
3. Create and send an email to a COI with a security label that would not be approved by the true Border PEP but is permitted by the rogue point.
4. Confirm the email was received at the destination.

4.1.29. Man in the Middle Attacks against the Policy Server

1. Setup a rogue Policy Server with a different security policy but will also pass connections on to the true Policy Server, modifying only key elements of the policy.
2. Employ a spoofing technique such as ARP poisoning or DNS redirection to masquerade as the true Policy Server.
3. Create and send an email to a COI with a security label that would not be approved by the true Policy Server but is permitted by the rogue point.
4. Confirm the email was received at the destination.

4.1.30. Man in the Middle Attacks against the Trusted Audit Service

1. Setup a rogue Trusted Audit Service that will also pass connections on to the true Trusted Audit Service, to delete and insert selective syslog messages.
2. Employ a spoofing technique such as ARP poisoning or DNS redirection to masquerade as the true Trusted Audit Service.
3. Perform a DCSS action that will generate an audit event that the rogue Trusted Audit Service will delete.
4. Confirm the event was not recorded in the true Trusted Audit Service.

4.1.31. Downgrading Cryptographic Connections to National Policy Enforcement Points

1. Run a Vulnerability Scanning tool against the National Policy Enforcement Point server to assess the supported cryptographic protocols and ciphers.

4.1.32. Downgrading Cryptographic Connections to Border Policy Enforcement Points

1. Run a Vulnerability Scanning tool against the Border Policy Enforcement Point server to assess the supported cryptographic protocols and ciphers.

4.1.33. Downgrading Cryptographic Connections to the Policy Server

1. Run a Vulnerability Scanning tool against the Policy Server to assess the supported cryptographic protocols and ciphers.

4.1.34. Downgrading Cryptographic Connections to the Trusted Audit Service

1. Run a Vulnerability Scanning tool against the Trusted Audit Service server to assess the supported cryptographic protocols and ciphers.

4.1.35. Spoofing Audit Records to Insert False Messages in the Trusted Audit Service

1. Employ a tool that spoofs the sending of syslog messages.
2. Send spoofed syslog messages to the Trusted Audit Service.
3. Confirm the spoofed syslog messages were received and inserted into the audit trail.

4.1.36. Information Disclosure through Trusted Audit Service Information Sharing Between Partners for Incident Response

1. Not applicable.

5. Conclusions and Recommendations

A number of potential attack scenarios were found in the DCSS architecture and configuration planned for the CWIX 2015 exercise for which procedures to test for the existence and impact where also derived. Moreover, means of mitigating each potential attack scenario were also developed to quickly manage their risk in the event that a scenario is confirmed.

The next step in this ongoing analysis should be to execute the procedures to test for the existence and impact of each attack scenario. To do so will require the development of a lab or test environment in which the DCSS vulnerabilities can be investigated independent of any DCSS development, functional testing or trials.

Moreover, during the course of testing, additional means of mitigating each attack scenario may also be discovered that may represent lower costs steps to implement than the recommended approaches. The mitigation steps could then be implemented by the DCSS development team, propagated back to the test environment and the test procedures executed a final time to validate the fix properly corrected the vulnerability. In this way, vulnerability testing of the DCSS becomes a step integrated into the development cycle with the goal of reducing the attack surface and vulnerabilities earlier in the development process rather than after release.

6. References

[Reference 1] *CWIX2015 Test Plan, December 23, 2014*;

[Reference 2] *CWIX2015 Trial Report, in progress, March 2015*;

[Reference 3] *SD003 CWIX2014 Test Report, March 2013*;

[Reference 4] *SD006 CWIX2014 Trial Report*, Bell Canada, July 2014; and

[Reference 5] *Confidentiality Labelling and Binding for Joint Coalition Information Sharing*, Edition A Version 1, NATO Standardization Agency, September 2014.