NIST SPECIAL PUBLICATION 1800-6B

Domain Name System-Based Electronic Mail Security

Volume B:
Approach, Architecture, and Security Characteristics

Scott Rose
Information Technology Laboratory
National Institute of Standards and Technology

William Barker
Dakota Consulting
Silver Spring, MD

Santos Jha
Chinedum Irrechukwu
The MITRE Corporation
McLean, VA

Karen Waltermire
National Cybersecurity Center of Excellence
National Institute of Standards and Technology

January 2018

This publication and its additional content is available free of charge from:
https://doi.org/10.6028/NIST.SP.1800-6

DISCLAIMER

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST or NCCoE, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

National Institute of Standards and Technology Special Publication 1800-6B, Natl. Inst. Stand. Technol. Spec. Publ. 1800-6B, 72 pages, (January 2018), CODEN: NSPUE2

All comments are subject to release under the Freedom of Information Act (FOIA).

National Cybersecurity Center of Excellence
National Institute of Standards and Technology
100 Bureau Drive
Mailstop 2002
Gaithersburg, MD 20899
Email: nccoe@nist.gov

NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

The National Cybersecurity Center of Excellence (NCCoE), a part of the National Institute of Standards and Technology (NIST), is a collaborative hub where industry organizations, government agencies, and academic institutions work together to address businesses' most pressing cybersecurity issues. This public-private partnership enables the creation of practical cybersecurity solutions for specific industries, as well as for broad, cross-sector technology challenges. Through consortia under Cooperative Research and Development Agreements (CRADAs), including technology partners—from Fortune 50 market leaders to smaller companies specializing in IT security—the NCCoE applies standards and best practices to develop modular, easily adaptable example cybersecurity solutions using commercially available technology. The NCCoE documents these example solutions in the NIST Special Publication 1800 series, which maps capabilities to the NIST Cyber Security Framework and details the steps needed for another entity to recreate the example solution. The NCCoE was established in 2012 by NIST in partnership with the State of Maryland and Montgomery County, Md.

To learn more about the NCCoE, visit https://nccoe.nist.gov. To learn more about NIST, visit https://www.nist.gov.

NIST CYBERSECURITY PRACTICE GUIDES

NIST Cybersecurity Practice Guides (Special Publication Series 1800) target specific cybersecurity challenges in the public and private sectors. They are practical, user-friendly guides that facilitate the adoption of standards-based approaches to cybersecurity. They show members of the information security community how to implement example solutions that help them align more easily with relevant standards and best practices and provide users with the materials lists, configuration files, and other information they need to implement a similar approach.

The documents in this series describe example implementations of cybersecurity practices that businesses and other organizations may voluntarily adopt. These documents do not describe regulations or mandatory practices, nor do they carry statutory authority.

ABSTRACT

This document proposes a reference guide on how to architect, install, and configure a security platform for trustworthy email exchanges across organizational boundaries. The project includes reliable authentication of mail servers, digitally signing and encrypting email, and binding cryptographic key certificates to sources and servers. The example solutions and architectures presented here are based upon standards-based and commercially available products. The example solutions presented here can be used by any organization implementing Domain Name System-based electronic mail security.

KEYWORDS

authentication; data integrity; digital signature; domain name system; electronic mail; encryption; internet addresses; internet protocols; named entities; privacy

ACKNOWLEDGMENTS

We are grateful to the following individuals for their generous contributions of expertise and time.

Name	Organization
Bud Bruegger	Fraunhofer IAO
Victoria Risk	Internet Systems Consortium
Eddy Winstead	Internet Systems Consortium
Paul Fox	Microsoft Corporation
Janet Jones	Microsoft Corporation
Nate Lesser	National Cybersecurity Center of Excellence
Karen Waltermire	National Cybersecurity Center of Excellence
Doug Montgomery	NIST ITL Advanced Networks Technologies Division
Ralph Dolmans	NLnet Labs
Benno Overeinder	NLnet Labs
Joe Gersch	Secure64
Saksham Manchanda	Secure64

The Technology Partners/Collaborators who participated in this build submitted their capabilities in response to a notice in the Federal Register. Respondents with relevant capabilities or product components were invited to sign a Cooperative Research and Development Agreement (CRADA) with NIST, allowing them to participate in a consortium to build this example solution. We worked with:

Technology Partner/Collaborator	Build Involvement
Fraunhofer IAO	Configuration of DNS Services products and Mail Transfer Agent
Internet Systems Consortium	DNS Services software
Microsoft Corporation	Mail User Agent, Mail Transfer Agent, and DNS Services products
NLNet Laboratories	DNS Services products and configuration of Mail Transfer Agent
Secure64	DNS Services and Mail User Agent products and configuration of Mail User Agent and Mail Transfer Agent

Contents

List of Figures

List of Tables

1 Summary

This National Institute of Standards and Technology (NIST) Cybersecurity Practice Guide addresses the challenge of providing digital signature technologies to provide authentication and integrity protection for electronic mail (email) on an end-to-end basis, and confidentiality protection for email in transit between organizations. It implements and follows recommendations of NIST Special Publication 800-177 (SP 800-177), *Trustworthy Email*. Detailed protocol information and implementation details are provided in SP 800-177. Domain Name System[1] protection features are consistent with SP 800-81-2, *Secure Domain Name System (DNS) Deployment Guide*.

The NIST Special Publication 1800-6 series of documents contain:

- rationale for and descriptions of a Domain Name System-based (DNS-based) email security platform that permits trustworthy email exchanges across organizational boundaries and
- a series of How-To Guides, including instructions for installation and configuration of the necessary services, that show system administrators and security engineers how to achieve similar outcomes

The solutions and architectures presented are built upon standards-based, commercially-available products. These solutions can be used by any organization deploying email services that is willing to implement certificate-based cryptographic key management and DNS Security Extensions (DNSSEC)[2]. Interoperable solutions are provided that are available from different types of sources (e.g., both commercial and open source products) and function in different operating systems environments.

This summary section describes the challenge addressed by this Volume B (Approach, Architecture, and Security Characteristics); describes the solution demonstrated to address the challenge; explains the benefits of the demonstrated solution; lists the technology partners that participated in building, demonstrating, and documenting the solution; and explains how to provide feedback on this guide. Section 2, How to Use This Guide explains how each volume of the guide may be used by business decision makers, program managers, and Information Technology (IT) professionals such as systems administrators; and Section 3, Approach provides a more detailed treatment of the scope of the project, describes the assumptions on which security platform development was based, describes the risk assessment that informed platform development, and describes the technologies and components that were provided by industry collaborators to enable platform development. Section 4, Architecture describes the usage scenarios supported by project security platforms, including Cybersecurity

[1] Request for Comments (RFC) 1591, *Domain Name System Structure and Delegation*
[2] RFC 4033, *DNS Security Introduction and Requirements*

Framework[3] functions supported by each collaborator-contributed component. Section 5, Outcome describes any changes in users' mail processing experience imposed by the additional security functionality, and summarizes changes to systems administrators' experiences with respect to integrating the new capabilities into their systems and in systems operations and maintenance. Section 6, Security Characteristic Analysis summarizes the test sequences that were employed to demonstrate security platform services, the Cybersecurity Framework functions to which each test sequence is relevant, the NIST SP 800-53-4 controls that applied to the functions being demonstrated, and an overview of platform performance in each of the two application scenarios demonstrated. Section 7, Future Build Considerations is a brief treatment of other applications that might be explored in the future in demonstrating the advantages of broader DNS security adoption. Appendices are provided for acronyms, references, and a mapping of this project to the Cybersecurity Framework Core[4] and informative security references cited in the Cybersecurity Framework Core.

1.1 Challenge

Both private industry and the government are concerned about email security and the use of email as an attack vector for cybercrime. Business operations are heavily reliant on email exchanges and need to protect the confidentiality of business information, the integrity of transactions, and privacy of individuals. Cryptographic services are used to authenticate the source of email messages, protect against undetected unauthorized alteration of messages in transit, and maintain message confidentiality. Efficiency and policies support reliance on mail servers to provide cryptographic protection for email rather than on end-to-end security operated by individual users. However, organizations need to protect their server-based email security mechanisms against intrusion and man-in-the-middle attacks during automated cryptographic service negotiation. In the absence of an appropriate combination of DNSSEC and certificate-based protections, any of these attacks can result in disclosure or modification of information by unauthorized third parties. The attacks can also enable an attacker to pose as one of the parties to an email exchange and send email that contains links to malware-ridden websites. If other content in a fraudulent message successfully motivates the user to click on the link or the user's system is configured to automatically follow some links or download content other than text, the malware will infect the user's system. Inclusion of links to malware is a major factor in most confirmed data breaches. Consequences of such breaches can range from exposing sensitive or private information, to enabling fraudulent activity by the attacker posing as the victimized user, to disabling or destroying the user's system—or that of the user's parent organization. Beyond

[3] *Framework for Improving Critical Infrastructure Cybersecurity*, Version 1.0, National Institute of Standards and Technology, February 12, 2014, https://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf

[4] https://www.nist.gov/cyberframework/

avoidance of negative consequences to users, improved email security can also serve as a marketing discriminator for email service providers.

Implementation of DNSSEC and DNS-Based Authentication of Named Entities (DANE)[5] has been impeded in the past by a shortage of easily used software libraries and by the fact that most available email applications of the protocols respond to DNSSEC failures by terminating the delivery attempt, often failing to alert the mail server that failure to deliver is based on a DNSSEC issue. The consequence of the first impediment is that, unless forced by policy to do so, IT organizations defer DNSSEC/DANE implementation pending availability of more mature software libraries. The consequence of the second is that, when DNSSEC and DANE are turned on, mail servers experience severe service degradation or crashes due to large numbers of retransmission attempts. (Note that this problem is experienced with mail servers, not DNS servers; DNS servers can handle the load.)

1.2 Solution

DNSSEC protects against unauthorized modifications to domain name information to prevent connection to spoofed or malicious hosts. The NCCoE initiated a collaborative project with industry partners to develop a proof-of-concept security platform that provides trustworthy mail server-to-mail server email exchanges across organizational boundaries. Products comprising the security platform include client mail user agents (MUAs)[6], DNS servers (authoritative and caching/recursive)[7], mail transfer agents (MTAs)[8], and X.509 cryptographic key certificate sources (components and services). The network infrastructure products are similar to those found in every enterprise and used to perform basic IT functions and handle email. The certificate utilities are needed to produce X.509 certificates[9] for mail servers and end users to support Transport Layer Security (TLS)[10] and Secure/Multipurpose

[5] RFC 6698, *The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security Protocol*: *TLSA*

[6] According to NIST SP 800-177, an MUA is a software component (or web interface) that allows an end user to compose and send messages to one or more recipients. An MUA transmits new messages to a server for further processing (either final delivery or transfer to another server).

[7] According to Section 3.2 of SP 800-177, there are two main types of name servers: authoritative name servers and caching name servers. The term **authoritative** is with respect to a zone. If a name server is an authoritative source for DNS resource records for a particular zone (or zones) of DNS addresses, it is called an **authoritative name server** for that zone (or zones). An authoritative name server for a zone provides responses to name resolution queries for resources for that zone, using the records in its own zone file. A **caching name server** (also called a resolving/recursive name server), by contrast, provides responses either through a series of queries to authoritative name servers in the hierarchy of domains found in the name resolution query or from a cache of responses built by using previous queries.

[8] Also, according to SP 800-177, mail is transmitted, in a "store and forward" fashion, across networks via MTAs. MTAs communicate using the Simple Mail Transfer Protocol (SMTP) described below and act as both client and server, depending on the situation.

[9] RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

[10] RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2*

Internet Mail Extensions (S/MIME)[11]. This project focused on Simple Mail Transfer Protocol (SMTP)[12] over TLS and S/MIME.

This project demonstrated a security platform, consistent with SP 800-177, that provides trustworthy email exchanges across organizational boundaries. The project included authentication of mail servers, digitally signing and encrypting email[13], and binding cryptographic key certificates to the servers. The software library issue was addressed in SP 1800-6C by providing installation and configuration instructions for using and maintaining existing software libraries (including installation support applications). At the same time, inclusion of software developers and vendors in the development and demonstration process revealed software and implementation guidance shortcomings that have been corrected.

1.3 Benefits

Sectors across industries, as well as the federal government, are concerned about email security and the use of email as an attack vector.[14] Both public and private sector business operations are heavily reliant on email exchanges. The need to protect the integrity of transactions containing financial and other proprietary information and to protect the privacy of employees and clients are among the factors that motivate organizations to secure their email. Whether the service desired is authentication of the source of an email message, assurance that the message has not been altered by an unauthorized party, or message confidentiality, cryptographic functions are usually employed. Economies of scale and a need for uniform implementation drive most enterprises to rely on mail servers to provide security to the members of an enterprise rather than security implemented and operated by individual users. Many server-based email security mechanisms are vulnerable to attacks involving:

- faked or fraudulent digital certificates
- otherwise invalid certificates
- failure to perform authentication process for connection

Even if there are protections in place, some attacks have been able to subvert email communication by attacking the underlying support protocols such as DNS. Attackers can spoof DNS responses to redirect email servers and alter email delivery. DNSSEC was developed to prevent this. DNSSEC protects against

[11] RFC 5751, *Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification*
[12] RFC 5321, *Simple Mail Transfer Protocol*
[13] Cryptographic protection, while voluntary for the private sector, has for a number of applications been made mandatory for federal government agencies (see Managing Information as a Strategic Resource, Office of Management and Budget (OMB) Circular A-130).
[14] "How Cybercrime Exploits Digital Certificates," Infosec Institute, *General Security*, July 28, 2014, http://resources.infosecinstitute.com/cybercrime-exploits-digital-certificates

unauthorized modifications to network management information and host IP addresses. DNSSEC can also be used to provide an alternative publication and trust infrastructure for service certificates using DANE resource records.

The business value of the security platform that results from this project includes improved privacy and security protections for users' communication, as well as improved management of DNS and email security operations. Addressing the software library and message retransmission issues, respectively, reduces the difficulty and cost of installing and maintaining DNSSEC and DANE. Mitigating the major cause of system errors resulting from faulty deployment of DNSSEC and DANE will encourage use of capabilities already present in many email systems. Demonstration and publication of these improvements encourages wider implementation of the protocols that provide Internet users with confidence that email has been protected and reaches the intended receiver in a secure manner. The demonstrated platform addresses three of the five Framework Core Functions and many requirements of relevant security standards and guidelines. Implementation of the platform will be increasingly important as a market discriminator as public awareness of email security and privacy issues grows.

2 How to Use This Guide

This NIST Cybersecurity Practice Guide demonstrates a standards-based reference design and provides users with the information they need to replicate this proof-of-concept security platform that demonstrates trustworthy email exchanges across organizational boundaries. This reference design is modular and can be deployed in whole or in parts.

This guide contains three volumes:

- NIST SP 1800-6A: *Executive Summary*
- NIST SP 1800-6B: *Approach, Architecture, and Security Characteristics* – what we built and why **(you are here)**
- NIST SP 1800-6C: *How-To Guides* – instructions for building the example solution

Depending on your role in your organization, you might use this guide in different ways:

Business decision makers, including chief security and technology officers will be interested in the *Executive Summary (NIST SP 1800-6A)*, which describes the:

- challenges enterprises may face in implementing best practices and standards to strengthen their email systems
- example solution built at the NCCoE
- benefits of adopting the example solution

Technology or security program managers who are concerned with how to identify, understand, assess, and mitigate risk will be interested in this part of the guide, *NIST SP 1800-6B,* which describes what we did and why. The following sections will be of particular interest:

- Section 3.4.3, Risk, provides a description of the risk analysis we performed
- Section 3.4.4, Cybersecurity Framework Functions, Categories, and Subcategories Addressed by the Project, maps the security characteristics of this example solution to cybersecurity standards and best practices

You might share the *Executive Summary, NIST SP 1800-6A,* with your leadership team members to help them understand the importance of adopting standards-based email security solutions.

IT professionals who want to implement an approach like this will find the whole practice guide useful. You can use the How-To portion of the guide, *NIST SP 1800-6C*, to replicate all or parts of the build created in our lab. The How-To guide provides specific product installation, configuration, and integration instructions for implementing the example solution. We do not recreate the product manufacturers' documentation, which is generally widely available. Rather, we show how we incorporated the products together in our environment to create an example solution.

This guide assumes that IT professionals have experience implementing security products within the enterprise. While we have used a suite of commercial products to address this challenge, this guide does not endorse these products. Your organization can adopt this solution or one that adheres to these guidelines in whole, or you can use this guide as a starting point for tailoring and implementing parts of the DNS-based email solution suite described herein. Your organization's security experts should identify the products that will best integrate with your existing tools and IT system infrastructure. We hope you will seek products that are congruent with applicable standards and best practices. Section 3.5, Technologies, lists the products we used and maps them to the cybersecurity controls provided by this reference solution.

A NIST Cybersecurity Practice Guide does not describe "the" solution, but a possible solution. Comments, suggestions, and success stories will help inform and improve future projects. Please contribute your thoughts to dns-email-nccoe@nist.gov.

2.1 Typographical Conventions

The following table presents typographic conventions used in this volume.

Typeface/ Symbol	Meaning	Example
Italics	filenames and pathnames references to documents that are not hyperlinks, new terms, and placeholders	For detailed definitions of terms, see the *NCCoE Glossary.*

Typeface/ Symbol	Meaning	Example
Bold	names of menus, options, command buttons and fields	Choose **File > Edit**.
Monospace	command-line input, on-screen computer output, sample code examples, status codes	`mkdir`
`Monospace Bold`	command-line user input contrasted with computer output	**`service sshd start`**
blue text	link to other parts of the document, a web URL, or an email address	All publications from NIST's National Cybersecurity Center of Excellence are available at: https://nccoe.nist.gov/

3 Approach

As stated in Section 1.1, both public and private sector business operations are heavily reliant on email exchanges. They need to protect the integrity of transactions that may include financial and other proprietary information. The privacy of employees and clients is also a factor that motivates organizations to secure their email systems. Security services such as the authentication of the source of an email message, assurance that the message has not been altered by an unauthorized party, and confidentiality of message contents require the use of cryptographic functions. A need for uniform security implementation drives most enterprises to rely on mail servers to provide security to the members of an enterprise rather than rely on end users to implement a security policy on their own. However, most current server-based email security mechanisms are vulnerable to, and have been defeated by, attacks on the integrity of the cryptographic implementations on which they depend. The consequences frequently involve unauthorized parties being able to read or modify supposedly secure information, or to use email as a vector for inserting malware into the enterprise. Improved email security can help protect organizations and individuals against these consequences and also serve as a marketing discriminator for email service providers as well as improve the trustworthiness of enterprise email exchanges.

Domain Name System Security Extensions for DNS are technical mechanisms employed by domain owners to protect against unauthorized modification to network management information. DANE is a protocol that securely associates domain names with cryptographic certificates and related security information so that clients can better authenticate network services. Despite the dangers of failure to authenticate the identities of network devices, adoption of DNSSEC has been slow. Demonstration of DANE-supported applications such as reliably secure email may support increased user demand for DNS security. Follow-on projects might include Hypertext Transfer Protocol Secure (HTTPS), the Internet of Things (IoT), Internet Protocol Security (IPsec) keys in DNS, and DNS service discovery.

This project demonstrated proof-of-concept security platforms composed of off-the-shelf components that provide trustworthy mail server-to-mail server email exchanges across organizational boundaries. The DANE protocol was used to authenticate servers and certificates in two roles: (1) by binding the X.509 certificates used for TLS to DNSSEC signed names for mail server-to-mail server communication; and (2) by binding the X.509 certificates used for S/MIME to email addresses encoded as DNS names. These bindings support trust in the use of S/MIME certificates in the end-to-end email communication. The resulting platforms encrypt email traffic between servers and allow individual email users to obtain other users' certificates in order to validate signed email or send encrypted email.[15] The project included an email sending policy consistent with a stated privacy policy that can be parsed by receiving servers so that receiving servers can apply the correct security checks.

Documentation of the resulting platform includes statements of the security and privacy policies and standards (e.g., Executive Orders, NIST standards and guidelines, Internet Engineering Task Force (IETF) RFCs). This also includes technical specifications for hardware and software, implementation requirements, and a mapping of implementation requirements to the applicable policies, standards, and best practices.

The project involved composition of a variety of components that were provided by several different technology providers. Components included MUAs, DNSSEC-capable DNS servers, MTAs, and cryptographic certificate sources. These components were used to generate and host DNSSEC signed zones and TLS-enabled mail services.

This project resulted in demonstration of support to MUAs and MTAs by four secure email platforms and this publicly available NIST Cybersecurity Practice Guide that explains how to employ the suite(s) to meet security and privacy requirements. This guide also provides platform documentation necessary to compose a DNS-based email security platform from off-the-shelf components that composed the prototype platforms.

[15] S/MIME can do this now, but DANE makes it easier to actually use.

3.1 Audience

This guide is intended for individuals responsible for implementing security solutions in organizations' IT support activities. Current IT systems, particularly in the private sector, often lack integrity protection for domain name services and email. The platforms demonstrated by this project and the implementation information provided in these Practice Guides permit integration of DNS and email integrity services and email confidentiality services with minimum changes to existing infrastructure or impact to service operations. The technical components will appeal to system administrators, IT managers, IT security managers, and others directly involved in the secure and safe operation of the business IT networks.

3.2 Scope

This project was consistent with NIST SP 800-177 and demonstrated the use of off-the-shelf TLS, DNSSEC, and DANE components to achieve trustworthy email objectives in a manner consistent with NIST SP 800-81-2.

3.2.1 Transport Layer Security (TLS)

The project used TLS to protect confidentiality of email messages exchanged between mail servers. TLS relies on public keys stored as X.509 digital certificates. These certificates can be used to authenticate the identity (server, domain or organization) of the certificate owner.

3.2.2 Domain Name System Security Extensions (DNSSEC)

The project used DNSSEC to authenticate and protect the integrity of DNS data.[16] DNSSEC uses digital signatures over DNS data to prevent an attacker from tampering with or spoofing DNS responses. Mail servers use the DNS to find the destination of email as well as storing other artifacts necessary for email security (see below).

3.2.3 DNS-Based Authentication of Named Entities (DANE)

The project used DANE, a protocol that securely associates domain names with cryptographic certificates and related security information so that they cannot be fraudulently modified or replaced to breach security. DNSSEC binds the X.509 certificates used for TLS to DNS.

[16] Note that this project addressed validation of X.509 certificates through the signing chain, not only through DANE.

3.2.4 Binding X.509 Certificates with DANE

The project also used DANE to bind the X.509 certificates used for S/MIME to email addresses encoded as DNS names verified by DNSSEC.

3.2.5 Demonstration of Digital Signature and Encryption of Email

The project demonstrated sending encrypted messages between email systems resident in different DNS domains, where the email exchanges between two organizations' email servers are carried over TLS, and the integrity of TLS key management connections is protected by DANE and DNSSEC. Signed email was sent between a message originator and a receiving party using end user applications (end-to-end) in different DNS domains, where the email exchanges between organizations were carried over TLS, the email messages were signed and verified with S/MIME on the end users' client devices, and the S/MIME key management was protected by DANE and DNSSEC. In addition, the project demonstrated that the use of DNSSEC and DANE could block an attempt by a fraudulent mail server to pose as the legitimate mail server for the receiver of the email.

3.2.6 Demonstration of End-to-End Digital Signature of Mail

The project's digital signature demonstration included sending S/MIME signed email between a message originator and a receiving party using end user applications in different DNS domains. The email exchanges between organizations are carried over TLS, the email messages are signed and verified with S/MIME on the end users' client devices, and the S/MIME certificates are stored in the DNS and protected by DNSSEC. This aspect of the project also demonstrated that use of DANE could block an attempt by a fraudulent actor to pose as the email originator.

3.3 Assumptions

The following assumptions exist for this project.

3.3.1 Security and Performance

The email platforms and DNS services demonstrated provided email integrity and confidentiality protection. An underlying assumption was that the benefits of using the demonstrated platforms outweighed any additional performance risks that may be introduced. The security of existing systems and networks was out of scope for this project. A key assumption was that all potential adopters of one of the demonstrated builds, or any of their components, already have in place some degree of network security. Therefore, we focused on what potential new system vulnerabilities were being introduced to end users if they implement this solution. The goal of this solution was to not introduce additional vulnerabilities into existing systems, but there is always inherent risk when adding systems and adding new features into an existing system.

3.3.2 Modularity

This assumption was based on one of the NCCoE core operating tenets. It is reasonably assumed that organizations already have mail client and server systems in place. Our philosophy is that a combination of certain components or a single component can improve email security for an organization; they may not need to remove or replace most existing infrastructure. This guide provides a complete top-to-bottom solution and is also intended to provide various options based on need.

3.3.3 Technical Implementation

This practice guide is written from a "how-to" perspective, and its foremost purpose is to provide details on how to install, configure, and integrate the components. The NCCoE assumes that an organization has the technical resources to implement all or parts of the build, or has access to companies that can perform the implementation on its behalf.

3.3.4 Operating System and Virtual Machine Environments

This project was conducted primarily in a VMware vCenter server version 6.0.0 Build 3018523 virtual machine environment. It is assumed that user organizations will be able to install the demonstrated applications in cloud-hosted virtual machines (VMs), local virtual machine or local native server client environments. This project uses Centos 7, Windows Server 2012R2, and Windows 10 operating systems. Operating systems were chosen based on the requirements of the software.

This project assumes, and is dependent upon, the availability of off-the shelf information security technology. Specific products and expertise on which the project is dependent include those for MUAs, MTAs, DNS servers (authoritative and recursive) and X.509 certificate utilities.

3.4 Risk Assessment

According to NIST SP 800-30, *Risk Management Guide for Information Technology Systems*, "Risk is the net negative impact of the exercise of a vulnerability, considering both the probability and the impact of occurrence. Risk management is the process of identifying risk, assessing risk, and taking steps to reduce risk to an acceptable level." The NCCoE recommends that any discussion of risk management, particularly at the enterprise level, begin with a comprehensive review of the *Framework for Improving Critical Infrastructure Cybersecurity*[17] (Cybersecurity Framework) and NIST SP 800-39, *Managing Information Security Risk: Organization, Mission, and Information System View*. The risk management

[17] *Framework for Improving Critical Infrastructure Cybersecurity*, Version 1.0, National Institute of Standards and Technology, February 12, 2014. https://www.nist.gov/sites/default/files/documents/cyberframework/cybersecurity-framework-021214.pdf

framework (RMF) and its associated references for identified security functions provide a baseline for organizing and relating to organizational objectives of:

- the risks to email and the networks it transits
- the security requirements to be met in order for the security platform to reduce these risks

While this guide does not present a full risk assessment, it does highlight the broad categories of threats and vulnerabilities associated with email.

3.4.1 Threats

Below are common threats associated with email:

- use of email as a vehicle for introducing malware
- use of email as a delivery mechanism for social engineering attacks
- theft or destruction of data communicated by email and/or its attachments due to loss or unauthorized/unintentional disposal of messages
- unauthorized access to email that results in a loss of privacy
- unauthorized modification of information communicated by email
- malicious fraudulent creation of messages or attachments attributed to third parties
- redirection or duplication of message to other than the intended recipient

3.4.2 Vulnerabilities

Vulnerabilities are commonly associated with mail client applications, mail transfer applications, and network applications that are employed in creation, delivery, and reading of email. However, vulnerabilities can be exploited at all levels in the information stack. For up-to-date information regarding vulnerabilities, this guide recommends that security professionals leverage the National Vulnerability Database (NVD). The NVD is the U.S. government repository of standards-based vulnerability management data [https://nvd.nist.gov].

3.4.2.1 Client System Vulnerabilities

Organizations are getting better at protecting network perimeters, and companies with mature security programs usually allow only certain ports through the firewall and harden internet-accessible servers to minimize the attack surface. As a result, attackers are paying closer attention to client-side vulnerabilities on internal workstations. These client-side vulnerabilities often are as simple as unpatched software on a desktop or laptop. Most client systems run at least one operating system and quite a few applications. Listing specific vulnerabilities for each is beyond the scope of this guide, but a current list of vulnerabilities and information regarding patches are available from NIST's NVD referenced above. Depending on the nature of a vulnerable application, an attacker may exploit it using

a specially crafted email attachment or by convincing the user to visit a malicious Web site. Web browsers are common targets. Other attractive targets include Adobe Acrobat[18], Macromedia Flash[19], QuickTime[20], and Java Runtime Environment[21].

3.4.2.2 Mail Server Vulnerabilities

Mail servers have many of the same vulnerabilities as client systems, but we also need to be aware of protocol-based vulnerabilities involving access to valid lists of email addresses, vulnerabilities to relay exploits for malware insertion, vulnerabilities to email header disclosures, and vulnerabilities to viruses and worms. In the case of SMTP, one way that attackers have in the past verified whether email accounts exist on a server is simply to telnet to the server on port 25 and run the VRFY command. The VRFY command makes a server check whether a specific user ID exists. Spammers often automate this method to perform a **directory harvest attack**, which is a way of gleaning valid email addresses from a server or domain for hackers to use. Scripting this attack can test thousands of email address combinations. The SMTP command EXPN may allow attackers to verify what mailing lists exist on a server. Yet another way to capture valid email addresses is to use applications such as *theHarvester* to glean addresses via Google and other search engines. In such environments, the best solution for preventing this type of email account enumeration depends on whether you need to enable commands like SMTP's VRFY and EXPN. In general, it is important to ensure that company email addresses are not posted on the web.

Protocols like SMTP relay let users send emails through external servers. Open email relays are not the problem they used to be, but they can still be sources of vulnerabilities. Spammers and hackers can use an email server to send spam or malware through email under the guise of the unsuspecting open-relay owner.

In the case of email header disclosures, email servers configured with typical defaults may be vulnerable to divulging information such as internal Internet Protocol (IP) addresses of email clients, software versions of client and email servers along with their vulnerabilities, or host names that can divulge network naming conventions.

Email systems are regularly targeted by malware such as viruses and worms. It is necessary to verify that mail servers' antivirus software is actually working. As in the case of client system vulnerabilities,

[18] See https://www.cvedetails.com/vulnerability-list/vendor_id-53/product_id-497/Adobe-Acrobat-Reader.html.
[19] See https://www.cvedetails.com/vulnerability-list/vendor_id-73/product_id-1950/version_id-8545/Macromedia-Flash-Player-6.0.29.0.html.
[20] See https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-7117.
[21] See https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-4903.

NIST's NVD (https://nvd.nist.gov) is a frequently updated source of vulnerabilities that affect mail servers.

3.4.2.3 Network Vulnerabilities

The MITRE Corporation's Common Vulnerabilities and Exposures database (CVE) lists more than 85,000 vulnerabilities that can affect web servers, Structured Query Language (SQL) servers, DNS servers, firewalls, routers, and other network components (see https://cve.mitre.org). These include vulnerabilities that can result in denial of service, code execution, overflow, cross-site scripting, directory traversal, process bypass, unauthorized gaining of information, SQL injection, file inclusion, memory corruption, cross-site request forgery, and http response splitting. Many of the vulnerabilities are operating system or application-based. Others are protocol based (e.g., vulnerabilities inherent in IP[22], TLS, DNS[23], Border Gateway Protocol (BGP)[24], SMTP and other network protocols). As in the case of client system vulnerabilities, NIST's NVD (https://nvd.nist.gov) is a frequently updated source of vulnerabilities that affect network servers.

3.4.3 Risk

Risks are examined from the point of view of consequences of vulnerabilities being exploited. Some examples of these consequences include legal liability, consequences of failure to comply with regulations, confidentiality breaches, loss of productivity, and damage to organizational reputation.

- New and existing regulations are forcing organizations to keep a record of their emails and to protect their employee and customer privacy. For example, the Health Insurance Portability and Accountability Act (HIPAA) requires health care institutions to keep a record of their email communications and secure confidentiality of information. In the new Internal Revenue Service (IRS) regulation Circular 230, the IRS requires tax advisors to add an email disclaimer to any emails including tax advice, expressly stating that the opinion cannot be relied upon for penalty purposes. The U.S. Securities and Exchange Commission and Gramm-Leach-Bliley Act impose similar duties on financial institutions. Steep penalties can apply to those organizations that do not comply with their industry's regulations. In a case lasting from 2000 until 2005, a well-known financial institution was recently forced to pay 20 million dollars in penalties by the Securities and Exchange Commission for not diligently searching for email backup tapes and overwriting multiple backup tapes.
- Most confidentiality breaches occur from within the company. These breaches can be accidental, but they can also be intentional.

[22] RFC 791, *Internet Protocol*
[23] RFC 1034, *Domain Names - Concepts and Facilities*
[24] RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

- With respect to legal liability, organizations are generally held responsible for all the information transmitted on or from their system, so inappropriate emails sent on the company network can result in multi-million dollar penalties.
- Employees sending personal emails and sifting through spam mail can cause major loss of productivity.[25]
- Even just a badly written email, or an email containing unprofessional remarks will cause the recipient to gain a bad impression of the company that the sender is representing. Fraudulent email attributable to an organization can do far more damage to an organization's reputation, both in terms of the response elicited and in terms of loss of confidence in the cybersecurity reliability of the organization.
- Another example of consequence may be damage to the perceived value of an organization, to be distinguished from organizational reputation, which is more about the habits or characteristics of a particular organization.[26]

A number of cybersecurity actions are recommended to reduce these risks. The Framework Core identified in NIST's Cybersecurity Framework is a set of cybersecurity activities, desired outcomes, and applicable references that are common across critical infrastructure sectors. The Core presents industry standards, guidelines, and practices in a manner that allows for communication of cybersecurity activities and outcomes across the organization from the executive level to the implementation/operations level. The Framework Core consists of five concurrent and continuous Functions: Identify, Protect, Detect, Respond, and Recover. When considered together, these functions provide a high-level, strategic view of the lifecycle of an organization's management of cybersecurity risk.

3.4.4 Cybersecurity Framework Functions, Categories, and Subcategories Addressed by the Project

NIST's Cybersecurity Framework provides a common language for understanding, managing, and expressing cybersecurity risk both internally and externally. It can be used to help identify and prioritize actions for reducing cybersecurity risk, and it is a tool for aligning policy, business, and technological approaches to managing that risk. It can be used to manage cybersecurity risk across entire organizations or it can be focused on the delivery of critical services within an organization. Different types of entities—including sector coordinating structures, associations, and organizations—can use the Cybersecurity Framework for different purposes, including the creation of common profiles. As stated

[25] Current spam filtering solutions consist of some sort of filtering at the network or the PC level, and they do not reveal the details of the sender without looking up the source. It takes some work for the recipient. This will always put us one step behind our adversaries. DNS provides the necessary Internet-wide scaling.

[26] Please see: https://www.wired.com/2016/10/security-news-week-verizon-reportedly-wants-1-billion-discount-yahoo-deal/ and http:/nypost.com/2016/10/06/verizon-wants-1b-discount-on-yahoo-deal-after-hacking-reports/. "The discount is being pushed because it feels Yahoo's value has been diminished," sources said.

above, the Framework Core provides a set of activities to achieve specific cybersecurity outcomes, and references examples of guidance to achieve those outcomes. The Core is not a checklist of actions to perform. It presents key cybersecurity outcomes identified by industry as helpful in managing cybersecurity risk. The Core comprises four elements: Functions, Categories, Subcategories, and Informative References.

- **Functions** organize basic cybersecurity activities at their highest level. These Functions are: Identify, Protect, Detect, Respond, and Recover. They aid an organization in expressing its management of cybersecurity risk by organizing information, enabling risk management decisions, addressing threats, and improving by learning from previous activities. The Functions also align with existing methodologies for incident management and help show the impact of investments in cybersecurity. For example, investments in planning and exercises support timely response and recovery actions, resulting in reduced impact to the delivery of services.
- **Categories** are the subdivisions of a Function into groups of cybersecurity outcomes closely tied to programmatic needs and particular activities. Examples of Categories include "Asset Management," "Access Control," and "Detection Processes."
- **Subcategories** further divide a Category into specific outcomes of technical and/or management activities. They provide a set of results that, while not exhaustive, help support achievement of the outcomes in each Category. Examples of Subcategories include "External information systems are catalogued," "Data-at-rest is protected," and "Notifications from detection systems are investigated."
- **Informative References** are specific sections of standards, guidelines, and practices common among critical infrastructure sectors that illustrate a method to achieve the outcomes associated with each Subcategory. The Informative References presented in the Framework Core are illustrative and not exhaustive. They are based upon cross-sector guidance most frequently referenced during the Framework development process.

This project supported the Cybersecurity Framework's Protect, Detect, and Respond Functions. Applicability to specific Functions, Categories, and Subcategories is described in the following paragraphs.

3.4.4.1 Protect

The Protect Function develops and implements the appropriate safeguards needed to ensure delivery of critical infrastructure services. This Function supports the ability to limit or contain the impact of a potential cybersecurity event. Examples of outcome Categories within this Function addressed by the project include: Access Control, Data Security, and Protective Technology.

1. **Access Control (PR.AC)**
    a. **PR.AC-1**

The **PR.AC-1** subcategory under Access Control supports identities and credentials being managed for authorized devices and users. The security platform resulting from the project supports effective management of the credentials associated with the addresses from which email purportedly originates and the integrity of the user identities associated with the email.

The original design of the DNS did not include security; instead, it was designed to be a scalable distributed system. DNSSEC and DANE attempt to add security, while maintaining backward compatibility with the existing DNS. DNSSEC was designed to protect applications (and caching resolvers serving those applications) from using forged or manipulated DNS data. All answers from DNSSEC protected zones are cryptographically signed (i.e., digital signature over DNS data). By checking the digital signature, a DNS resolver is able to determine whether the information is authentic (i.e., unmodified and complete) and is served on an authoritative DNS server. While protecting IP addresses is the immediate concern for many users, DNSSEC can protect any data published in the DNS, including text records or mail exchange (MX) records, and can be used to bootstrap other security systems that publish references to cryptographic certificates stored in the DNS.

All DNSSEC responses contain signed DNS data. DNSSEC signature validation allows the use of potentially untrustworthy parties if (for example) the mail server is using a self-signed certificate. The protocol permits configuration of systems to accept messages whether or not they are digitally signed. The security platform developed under the project permits email clients and transfer agents to be configured to send email messages to only servers whose DNS entries are digitally signed. At the client systems level (e.g., Outlook, Postfix, Thunderbird), digital signature of the mail messages themselves can also be applied on a user-to-user basis. In the user-to-user case, the signature provides assurance of the integrity of the identity of the sender rather than just the identity of the DNS zone(s) associated with the sender.

b. PR.AC-5

The **PR.AC-5** subcategory under Access Control supports protection of network integrity by incorporating network segregation where appropriate. The project does not specifically employ network segregation principles. However, it does support network integrity by providing operationally feasible mechanisms for preventing connections or message delivery to sources that do not implement a specified set of DNS security extensions. Rigorous adherence to a minimum security configuration can enforce effective isolation of a network from entities that do not conform to the network's security requirements. NIST SP 800-53, referred to by this subcategory, requires information systems to enforce approved authorizations for controlling the flow of

information within systems and between interconnected systems (AC-4, Information Flow Enforcement).

2. **Data Security (PR.DS)**

   The Protect Function's Data Security Category supports an outcome in which information and records (data) are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information. The project demonstrates a capability to provide source and content integrity protection by employing digital signature of messages and confidentiality protection by encrypting messages.

   a. **PR.DS-1**

      The **PR.DS-1** subcategory under Data Security supports protection of data at rest. The user-to-user digital signature capability demonstrated by the project can provide an ability to verify the source and content integrity of locally stored email messages where the digital signature is stored with the rest of the message. This supports integrity protection for data-at-rest.

   b. **PR.DS-2**

      The **PR.DS-2** subcategory under Data Security supports protection of data in transit. In addition to user-to-user digital signature of email, the project demonstrates a capability to provide source and content integrity protection to data-in-transit. The demonstration accomplishes this by employing server-to-server confidentiality protection to data-in-transit by employing server-to-server encryption.

   c. **PR.DS-6**

      The **PR.DS-6** subcategory under Data Security supports use of integrity checking mechanisms to verify software, firmware, and information integrity. The digital signature of email demonstrated by the project's security platform supports automatic integrity checking of information communicated in email messages. DNSSEC and DANE protect the integrity of address information.

3. **Protective Technology (PR.PT)**

   a. **PR.PT-4**

      The **PR.PT-4** subcategory under Protective Technology supports protection of communications and control networks. The project demonstrates a capability to provide source and content integrity protection by employing digital signature of communications and confidentiality protection by encrypting communications. The project's demonstration of DNSSEC and DANE protocols also supports communications

and control network integrity. It does this by demonstrating operationally feasible mechanisms for refusing connections to or message delivery from sources that do not implement a specified set of DNS security extensions. Rigorous adherence to a minimum security configuration can be used to enforce isolation of networks from entities that do not conform to the network's security requirements.

3.4.4.2 Detect

The Detect Function develops and implements the appropriate activities needed to identify in a timely manner the occurrence of a cybersecurity event. Examples of outcome categories within this function addressed by the project include Security Continuous Monitoring and Detection Processes.

1. **Security Continuous Monitoring (DE.CM)**

   The Security Continuous Monitoring Category supports an outcome in which information systems and assets are monitored at discrete intervals to identify cybersecurity events and to verify the effectiveness of protective measures. While not a classic example of continuous monitoring, the project's platform has the ability to automatically check all DNS responses for correct digital signatures.

   a. **DE.CM-1**

      The **DE.CM-1** subcategory under Security Continuous Monitoring supports monitoring of networks to detect potential cybersecurity events. While not a classic example of continuous monitoring, the demonstrated capability of the project's platform to automatically check all inbound DNS responses for valid digital signatures permits identification of attempts to spoof systems using bogus DNS data. Automatic signing and signature validation for email permits continuous checking for false sender identities and modification of message content. NIST SP 800-53, referred to by this subcategory, requires monitoring of inbound and outbound communications traffic for unauthorized conditions (SI-4 [4]). Validation of DNS addresses supports this requirement.

   b. **DE.CM-6**

      The **DE.CM-6** subcategory under Security Continuous Monitoring supports monitoring of external service provider activity to detect potential cybersecurity events. While not a classic example of continuous monitoring, the demonstrated capability of the project's platform to automatically check all inbound DNS responses for valid digital signatures permits detection of attempts by invalid service providers (e.g., bogus Certificate Authorities or Mail Transfer Agents) to spoof users' systems (including man-in-the-middle attacks).

2. **Detection Processes (DE.DP)**

   The Detection Processes Category supports an outcome in which detection processes and procedures are maintained and tested to ensure timely and adequate awareness of anomalous events.

   a. **DE.DP-4**

      The **DE.DP-4** subcategory under Detection Processes supports the communication of event detection information to appropriate parties. One of the shortcomings of most DNSSEC and DANE mechanisms is that they abort delivery of messages to destinations whose DNSSEC signature checks fail to validate and do not provide any indication that failure is due to an invalid signature. This usually results in numerous retransmissions and consequent performance degradation or possible crashes. The project's platform includes notifications of DNS signature failures to mail agents in its DNS resolvers in order to prevent consequent performance degradation. This communication of detection information has the potential to mitigate one of the primary impediments to private sector adoption of DNSSEC.

3.4.4.3 Respond

The Respond Function develops and implements the appropriate activities to take action regarding a detected cybersecurity event. This Function supports the ability to contain the impact of a potential cybersecurity event. Examples of outcome categories within this function addressed by the project include: Response Planning, Communications, and Mitigation.

1. **Response Planning (RS.RP)**

   The Response Planning Category supports an outcome in which response processes and procedures are executed and maintained to ensure timely response to detected cybersecurity events.

   a. **RS.RP-1**

      The **RS.RP-1** subcategory under Response Planning supports execution of a response plan during or after an event. Inclusion of DNS and email security considerations in planning for connection of systems to the Internet will necessarily include responses to detection of invalid digital signatures. This includes security flagging of connections and messages and/or refusing connections and delivery of messages. Concurrent with detection of validation failure, these responses are demonstrated by the project's platform.

2. **Communications (RS.CO)**

The Respond Communications Category requires response activities to be coordinated with internal and external stakeholders, as appropriate, to include external support from law enforcement agencies.

a. RS.CO-2

The **RS.CO-2** subcategory under Communications supports reporting of events consistent with established criteria. As stated under DE.DP-4, one of the shortcomings of most DNSSEC and DANE mechanisms is that they abort delivery of messages to destinations whose DNSSEC signature checks fail but do not provide any indication that the failure is due to an invalid signature. To prevent consequent performance degradation, the project's platform includes notifications of DNSSEC signature failures to mail agents in its DNS resolvers. This communication of detection information has the potential to mitigate one of the primary impediments to private sector adoption of DNSSEC. It also provides a mechanism that can be exploited to provide information involving failures of DNSSEC signature checks to external stakeholders.

3. Mitigation (RS.MI)

The Response Mitigation Category requires activities to be performed to prevent expansion of an event, mitigate its effects, and eradicate the incident.

a. RS.MI-1

The **RS.MI-1** subcategory under Mitigation supports containment of incidents. Implementation of the project's platform will contain the effects of incidents because any spoofing attempts or modified email will be detected and contained before they have a chance to negatively impact any organizational systems.[27]

b. RS.MI-2

The **RS.MI-2** subcategory under Mitigation supports mitigation of incidents. The project demonstrates user-to-user digital signature of messages. Retention of their digital signatures with stored messages permits later determination of whether the messages have been modified in storage. This can be a mitigating factor in the case of incidents that involve introduction of fraudulent information into email records. The project's demonstration of server-to-server encryption provides confidentiality protection for data-in-transit. This confidentiality protection can serve as a mitigating factor in the

[27] Note that if a system is subverted, a lot of assumed security goes out the window. A subverted sending MTA could still be seen as valid by receivers, for example.

case of incidents involving unauthorized access to messages captured by network devices that sit between the sender's and recipient's mail servers.

3.4.5 Cybersecurity References Directly Tied to Those Cybersecurity Framework Categories and Subcategories Addressed by the Project

The following security references were followed in accepting components for the project's platform, designing the platform, conducting demonstrations of the platform, and documenting the platform. The Framework Functions, Categories, and Subcategories addressed by these references are listed for each reference. While many of the references were written as standards and guidelines to be applied to federal government agencies, their recommendations may also be applied in the private sector as best practices that support the Cybersecurity Framework. Those Subcategories addressed by the platform are in **boldface**.

1. *Security Requirements for Cryptographic Modules*, Federal Information Processing Standard (FIPS) 140-2, May 2001. https://doi.org/10.6028/NIST.FIPS.140-2.

   FIPS 140-2 provides a standard that is required to be used by Federal organizations when these organizations specify that cryptographic-based security systems be used to provide protection for sensitive or valuable data. Protection of a cryptographic module within a security system is necessary to maintain the confidentiality and integrity of the information protected by the module. All cryptographic components employed by the Federal government outside the national security community, including NCCoE security platforms that employ cryptography, must conform to FIPS 140-2. This standard specifies the security requirements that will be satisfied by a cryptographic module. The standard provides four increasing qualitative levels of security intended to cover a wide range of potential applications and environments. The security requirements cover areas related to the secure design and implementation of a cryptographic module. These areas include cryptographic module specification; cryptographic module ports and interfaces; roles, services, and authentication; finite state model; physical security; operational environment; cryptographic key management; electromagnetic interference/electromagnetic compatibility (EMI/EMC); self-tests; design assurance; and mitigation of other attacks.

   Within the context of the Cybersecurity Framework, FIPS 140-2 provides standards for "Protection" to be provided by cryptographic modules (PR.AC-2, PR.AC-4, **PR.DS-1**, **PR.DS-2**, PR.DS-5, **PR.DS-6**, PR.IP-3, and **PR.PT-4**) and "Detection" of failures or other exception conditions that might affect the protection afforded to systems by cryptographic modules (**DE.CM-1**, DE.CM-2, and DM.DP-3).

2. *Guide for Applying the Risk Management Framework to Federal Information Systems: A security Lifecycle Approach*, NIST SP 800-37 Rev. 1, Joint Task Force Transformation Initiative; February 2010 with updates as of June 5, 2014. https://doi.org/10.6028/NIST.SP.800-37r1.

   SP 800-37 Rev. 1 provides guidelines for applying the Risk Management Framework (RMF) to federal information systems. Systems to which the RMF is to be applied include NCCoE use case and block activities. The RMF promotes the concept of near real-time risk management and ongoing information system authorization through the implementation of robust continuous monitoring processes; provides senior leaders with the necessary information to make cost-effective, risk-based decisions with regard to the organizational information systems supporting their core missions and business functions; and integrates information security into the enterprise architecture and development life cycle. Applying the RMF within enterprises links management processes at the information system level to management processes at the organization level through a risk executive (function) and establishes lines of responsibility and accountability for security controls deployed within organizational information systems and inherited by those systems (i.e., common controls).

   The six-step RMF includes security categorization, security control selection, security control implementation, security control assessment, information system authorization, and security control monitoring. With respect to the Cybersecurity Framework, SP 800-37 assumes that system components, business environment and governance structure have been identified. The risk assessment that underlies categorization is based on the assumed understanding of these factors. SP 800-37 also focuses on impacts of security incidents rather than on threats that take advantage of system vulnerabilities to create those impacts. The control selection, control implementation, and system authorization recommendations of SP 800-37 do not map directly to the Cybersecurity Framework.

   However, SP 800-37 does provide recommendations relevant to **Identify** (ID.RA-5, ID.RA-6, ID.RM 1, and ID.RM-2 in Section 3.1), **Protect** (PR.IP-3, and PR.IP-7 in Sections 3.4 and 3.6), and **Detect**, (DE.AE-5 and **DE.CM-1** in Section 3.6) elements of the Cybersecurity Framework.

3. *Guidelines on Electronic Mail Security*; NIST SP 800-45 Ver. 2; Tracy, Jansen, Scarfone, Butterfield; February 2007. https://doi.org/10.6028/NIST.SP.800-45ver2.

   SP 800-45 provides guidelines intended to assist organizations in installing, configuring, and maintaining secure mail servers and mail clients. Specifically, the publication discusses in detail:

   a. email standards and their security implications
   b. email message signing and encryption standards
   c. the planning and management of mail servers

d. securing the operating system underlying a mail server

e. mail server application security

f. email-content filtering

g. email-specific considerations in the deployment and configuration of network protection mechanisms, such as firewalls, routers, switches, and intrusion detection and intrusion prevention systems

h. securing mail clients

i. administering the mail server in a secure manner

As suggested by its 2007 publication date, SP 800-45 does not reflect the most recent developments in email security, especially the more recent IETF RFCs (e.g., S/MIME Certificate Association (SMIMEA)[28] and TLS Certificate Association (TLSA)[29]), but the recommendations it makes are still germane.

With respect to the Cybersecurity Framework's **Identify** Function and its Categories and Subcategories, SP 800-45 recommends risk management activities, but does not go into detail that maps to Subcategory references. Under the Protect Function, Subcategory references **PR.AC-1**, PR.AC-4, **PR.AC-5**, PR.AT-1, PR.AT-2, PR.AT-5, **PR.DS-2**, **PR.DS-6**, PR.IP-2, PR.IP-4, and PR.PT-1 are addressed by the guideline. Under the **Detect** Function, Subcategory references DE.DP-1 and **DE.DP-4** are addressed by the guideline. In the **Detect** Function, Subcategory references DE.AE-2, DE.CM-1, DE.CM-4, DE.CM-5, DE.CM-8, DE.DP-1, and DE.DP-4 are addressed. Under the **Respond** Function, Subcategory references RS.RP-1, **RS.CO-1**, RS.CO-2, RS.AN-1, and RS.IM-1 are addressed by the guideline. Under the **Recover** Function, Subcategory reference RC.RP-1 is addressed by the guideline.

4. *Federal S/MIME V3 Client Profile*, NIST SP 800-49, Chernick, November 2002. https://doi.org/10.6028/NIST.SP.800-49.

   SP 800-49 was developed to provide organizations with approaches to assure that S/MIME products can interoperate and meet the email security needs of federal agencies both with respect to security features and adequate cryptographic algorithms. This profile states requirements for implementing sets of cryptographic algorithm suites specified elsewhere by the standards development organizations. The profile specifies a set of email security features (e.g., encrypted email and signed receipts) that are mandatory for federal agencies. SP 800-49

[28] See *Using Secure DNS to Associate Certificates with Domain Names For S/MIME* (draft ietf-dane-smime-14).
[29] RFC 6698, *The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA*

adds specificity to the S/MIME standards, while attempting to avoid violating those standards. As its 2002 publication date suggests, SP 800-49 is even more dated with respect to protocols than SP 800-45 (e.g., recommending the now deprecated Secure Hash Algorithm 1 (SHA-1) instead of SHA-2 for hashing, and the deprecated Triple Data Encryption Standard (DES) rather than the Advanced Encryption Standard (AES) for encryption). However, it too makes security recommendations that are still germane. The SP 800-49 requirements and recommendations fall into the Cybersecurity Framework **Protect** Function. It provides guidelines that address the Subcategory references **PR.DS-2**, **PR.DS-6**, and (less precisely) **PR.PT-4**.

5. *Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations*; NIST SP 800-52 Rev. 1; Polk, McKay, Chokhani; April 2014. https://doi.org/10.6028/NIST.SP.800-52r1.

   TLS provides mechanisms to protect sensitive data during electronic dissemination across the Internet. SP 800-52 provides guidance in the selection and configuration of TLS protocol implementations, while making effective use of FIPS and NIST-recommended cryptographic algorithms. SP 800-52 requires that TLS 1.1 be configured with FIPS-based cipher suites as the minimum appropriate secure transport protocol and recommends that agencies develop migration plans to TLS 1.2 by January 1, 2015. This SP also identifies TLS extensions for which mandatory support must be provided and some other recommended extensions. Like SP 800-49, the SP 800-52 requirements and recommendations fall into the Cybersecurity Framework **Protect** Function. The guideline addresses Subcategory references **PR.DS-2**, **PR.DS-6**, and (less precisely) **PR.PT-4**.

6. *Security and Privacy Controls for Federal Information Systems and Organizations*, NIST SP 800-53 Rev. 4, Joint Task Force Transformation Initiative, April 2013. https://doi.org/10.6028/NIST.SP.800-53r4.

   SP 800-53 provides a catalog of security and privacy controls for federal information systems and organizations and a process for selecting controls to protect organizational operations (including mission, functions, image, and reputation), organizational assets, individuals, other organizations, and the nation from a diverse set of threats, including hostile cyberattacks, natural disasters, structural failures, and human errors. The controls are customizable and implemented as part of an organization-wide process that manages information security and privacy risk. The controls address a diverse set of security and privacy requirements across the federal government and critical infrastructure that are derived from legislation, Executive Orders, policies, directives, regulations, standards, and/or mission/business needs. The publication also describes how to develop specialized sets of controls, or overlays, that are tailored for specific types of missions/business functions, technologies, or environments of operation. Finally, the catalog of security controls addresses security from both a functionality perspective (the strength of security functions and mechanisms provided) and an assurance

perspective (the measures of confidence in the implemented security capability). Addressing both security functionality and security assurance ensures that information technology products and the information systems built from those products using sound systems and security engineering principles are sufficiently trustworthy.

SP 800-53 Rev. 4 addresses all Cybersecurity Framework Functions, Categories, and Subcategories. Only the RC.CO-1 (Reputation after an event is repaired) and **RC.CO-2** (Recovery activities are communicated to internal stakeholders and executive and management teams) references under the **Recover: Communications** Category are not addressed by SP 800-53.

7. *Recommendation for Key Management: Part 1 - General*, NIST SP 800-57 Part Rev.4, Barker, January 2016; *Part 2 - Best Practices for Key Management Organization*, NIST SP 800-57 Part 2, Barker, Barker, Burr, Polk, and Smid, August 2005; and *Part 3 - Application-Specific Key Management Guidance*, NIST SP 800-57 Part 3 Rev. 1, Barker and Dang, January 2015. https://doi.org/10.6028/NIST.SP.800-57pt1r4, https://doi.org/10.6028/NIST.SP.800-57p2, https://doi.org/10.6028/NIST.SP.800-57pt3r1

   NIST SP 800-57 provides cryptographic key management guidance. Part 1 provides general guidance and best practices for the management of cryptographic keying material. Part 2 provides guidance on policy and security planning requirements for U.S. government agencies. Part 3 of this SP provides guidance when using the cryptographic features of current systems that may not exhibit all the properties recommended by Part 1 of the guideline. Part 3 includes applications-specific recommendations for, among other applications, the Public Key Infrastructure (PKI), IPsec, TLS, S/MIME, and DNSSEC. All of these recommendations apply directly to this project.

   SP 800-57 addresses all of the Cybersecurity Framework Functions except **Detect**. Audit is the primary mechanism relied on in SP 800-53 for detection purposes. The Categories and Subcategory references that are addressed by the guideline include Identify (ID.AM-2, ID.BE-3, ID.BE-4, ID.BE-5, ID.GV-1, ID.GV-4, ID.RA-4, and ID.RA-5), **Protect** (**PR.AC-1**, PR.AC-2, **PR.AC-3**, PR.AC-4, PR.AT-2, PR.AT-3, PR.AT-4, **PR.DS-1**, **PR.DS-2**, PR.DS-3, PR.DS-4, **PR.DS-6**, PR.IP-2, PR.IP-3, PR.IP-4, PR.IP-5, PR.IP-6, PR.IP-9, PR.PT-1, PR.PT-2, PR.PT-3, and **PR.PT-4**); **Respond** (**RS.RP-1**, RS.CO-1, **RS.CO-2**, RS.CO-3, RS.AN-2, and **RS.MI-2**); and **Recover** (RC.RP-1).

8. *Secure Domain Name System (DNS) Deployment Guide*, NIST SP 800-81-2, Chandramouli and Rose, September 2013. https://doi.org/10.6028/NIST.SP.800-81-2.

   The DNS is a distributed database that enables access to Internet resources via user-friendly domain names, rather than IP addresses, by translating domain names to IP addresses and back. The DNS infrastructure is made up of computing and communication entities called name servers, each of which contains information about a small portion of the domain name space. The name data provided by DNS is intended to be available to any computer located anywhere

in the Internet. SP 800-81-2 provides deployment guidelines for securing DNS within an enterprise. The primary security goals for DNS are data integrity and source authentication, which are needed to ensure the authenticity of name information and maintain the integrity of name information in transit. This document provides extensive guidance on maintaining data integrity and performing source authentication. This document presents guidelines for configuring DNS deployments to prevent many redirection attacks that exploit vulnerabilities in various DNS components.

The Categories and Subcategory references that are addressed are limited to **Identify** (ID.AM-2 and ID.RA-6), **Protect** (**PR.AC-1**, **PR.AC-3**, **PR.AC-5**, PR.AT-2, **PR.DS-2**, PR.DS-5, **PR.DS-6**, PR.IP-3, PR.IP-4, PR.IP-6, and PR.IP-9), and **Detect** (**DE.CM-1** and DE.CM-7).

9. *A Framework for Designing Cryptographic Key Management Systems*; NIST SP 800-130; Barker, Branstad, Smid, Chokhani; August 2013. https://doi.org/10.6028/NIST.SP.800-130.

   SP 800-130's framework for designing Cryptographic Key Management Systems (CKMS) contains topics that should be considered by a CKMS designer when developing a CKMS design specification. For each topic, there are one or more documentation requirements that need to be addressed by the design specification. Thus, any CKMS that addresses each of these requirements would have a design specification that is compliant with this framework. A CKMS will be a part of a larger information system that executes processing applications. While the CKMS supports these applications by providing cryptographic key management services, the particular applications or particular classes of applications are beyond the scope of this framework.

   SP 800-130 addresses all the Cybersecurity Framework Functions. The Category and Subcategory references that are addressed include **Identify** (ID.BE-4, ID.GV-1, ID.GV-2, ID.GV-3, ID.GV-4, ID.RA-1, ID.RA-2, ID.RA-3, ID.RA-5, and RM-1); **Protect** (**PR.AC-1**, PR.AC-2, PR.AC-4, **PR.AC-5**, PR.AT-1, PR.AT-2, PR.AT-4, PR.AT-5, PR.DS-1, PR.DS-2, PR.DS-3, **PR.DS-6**, PR.DS-7, PR.IP-1, PR.IP-3, PR.IP-4, PR.IP-5, PR.IP-6, PR.IP-9, PR.MA-1, PR.PT-1, PR.PT-2, PR.PT-3, and **PR.PT-4**); **Detect** (DE.AE-4, **DE.CM-1**, DE.CM-4, **DE.CM-7**, DE.CM-8,DE.DP-1, DE.DP-2, DE.DP-3, and DE.DP-5); **Respond** (**RS.RP-1**, RS.CO-1, **RS.CO-2**, RS.AN-2, **RS.MI-1**, and RS.MI-2); and **Recover** (RC.RP-1).

10. *A Profile for U.S. Federal Cryptographic Key Management Systems (CKMS)*; NIST SP 800-152; Barker, Branstad, Smid; October 2015. https://doi.org/10.6028/NIST.SP.800-152.

    SP 800-152 covers major aspects of managing the cryptographic keys that protect federal information. Associated with each key is specific information (e.g., the owner identifier, its length, and acceptable uses) called metadata. The computers, software, modules, communications, and roles assumed by one or more authorized individuals when managing and using cryptographic key management services are collectively called a Cryptographic Key

Management System (CKMS). The Profile for U. S. Federal Cryptographic Key Management Systems (FCKMSs) has been prepared to assist CKMS designers and implementers in selecting the features to be provided in their "products," and to assist federal organizations and their contractors when procuring, installing, configuring, operating, and using FCKMSs.

SP 800-130 addresses all the Cybersecurity Framework Functions. The Categories and Subcategory references that are addressed include **Identify** (ID.AM-3, ID.AM-5, ID.BE-4, ID.BE-5, ID.GV-1, ID.GV-2, ID.GV-3, ID.GV-4, ID.RA-1, ID.RA-3, ID.RA-5, ID.RA-6, RM-1, and RM-2); **Protect** (**PR.AC-1**, PR.AC-2, **PR.AC-3**, PR.AC-4, **PR.AC-5**, PR.AT-1, PR.AT-2, PR.AT-4, PR.AT-5, **PR.DS-1**, **PR.DS-2**, PR.DS-3, PR.DS-4, **PR.DS-6**, PR.DS-7, PR.IP-1, PR.IP-3, PR.IP-4, PR.IP-5, PR.IP-6, PR.IP-7, PR.IP-8, PR.IP-9, PR.IP-12, PR.MA-1, PR.PT-1, PR.PT-2, PR.PT-3, and **PR.PT-4**); **Detect** (DE.AE-4, **DE.CM-1**, DE.CM-4, **DE.CM-7**, DE.CM-8, DE.DP-1, DE.DP-2, DE.DP-3, and DE.DP-5); **Respond** (**RS.RP-1**, RS.CO-1, **RS.CO-2**, RS.AN-2, **RS.MI-1**, **RS.MI-2**, RS.MI-3, and RS.IM-2); and **Recover** (RC.RP-1 and RC.IM-2).

11. *Trustworthy Email*; NIST SP 800-177; Chandramouli, Garfinkel, Nightingale, and Rose; September 2016. https://doi.org/10.6028/NIST.SP.800-177

    NIST SP 800-177 serves as a complimentary document to SP 800-45. SP 800-177 addresses email protocol security and provides descriptions, guidelines and recommendations for deploying new email security protocols such as SMTP over TLS, email supported by DANE, and other non-cryptographic authentication (e.g. Sender Policy Framework). Discussions of SMTP over TLS and S/MIME relate directly to the work on the project.

    With respect to the Cybersecurity Framework's Identify Function and its Subcategories, SP 800-177 recommends risk management activities, but does not go into detail that maps to subcategory references. Under the **Protect** Function, Subcategory references **PR.AC-1**, **PR.AC-3**, PR.AC-4, **PR.AC-5**, PR.AT-1, PR.AT-2, PR.AT-5, **PR.DS-2**, **PR.DS-6**, PR.IP-2, PR.IP-4, and PR.PT-1 are addressed by the guideline. Under the **Detect** Function, Subcategory references DE.AE-2, DE.CM-1, DE.CM-4, DE.CM-5, DE.CM-8, DE.DP-1, and **DE.DP-4** are addressed by the guideline. Under the **Respond** Function, Subcategory references RS.RP-1, RS.CO-1, **RS.CO-2**, RS.AN-1, and RS.IM-1 are addressed by the guideline. Under the **Recover** Function, Subcategory reference RC.RP-1 is addressed by the guideline.

3.4.6 Other Security References Applied in the Design and Development of the Project

The following references provided additional security and protocol standards and guidelines that were applied during design and development of the project.

1. *Systems Security Engineering: Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems*, NIST SP 800-160, November 2016. https://doi.org/10.6028/NIST.SP.800-160.

   NIST SP 800-160 defines systems security engineering processes that are tightly coupled to and fully integrated into well-established, international standards-based systems and software engineering processes. The project supports the federal cyber security strategy of "Build It Right, Continuously Monitor" and consisted of a four-phase development approach that culminated in the publication of this final systems security engineering guideline. The four phases included:

   - **Phase 1:** Development of the systems security engineering technical processes based on the technical systems and software engineering processes defined in Internet Organization for Standardization (ISO)/International Electrotechnical Commission (IEC)/Institute of Electrical and Electronics Engineers (IEEE) 15288:2008;

   - **Phase 2:** Development of the remaining supporting appendices: Information Security Risk Management (including the integration of the RMF, security controls, and other security- and risk-related concepts into the systems security engineering processes), Use Case Scenarios, Roles and Responsibilities, System Resiliency, Security and Trustworthiness, Acquisition Considerations, and the Department of Defense Systems Engineering Process;

   - **Phase 3:** Development of the systems security engineering *nontechnical processes* based on the nontechnical systems and software engineering processes (i.e., Agreement, Organizational Project-Enabling, and Project) defined in ISO/IEC/IEEE 15288:2008; and

   - **Phase 4:** Alignment of the technical and nontechnical processes based on the updated systems and software engineering processes defined in ISO/IEC/IEEE DIS 15288:201x (E).

   The full integration of the systems security engineering discipline into the systems and software engineering discipline involves fundamental changes in the traditional ways of doing business within organizations—breaking down institutional barriers that, over time, have isolated security activities from the mainstream organizational management and technical processes, including, for example, the system development life cycle, acquisition/procurement, and enterprise architecture. The integration of these interdisciplinary activities requires the strong support of senior leaders and executives, and increased levels of communication among all stakeholders who have an interest in, or are affected by, the systems being developed or enhanced.

2. *Internet X.509 Public Key Infrastructure Certificate and CRL Profile*; IETF RFC 2459; Housley, Ford, Polk, Solo; January 1999. https://datatracker.ietf.org/doc/rfc2459.

   RFC 2459 is one part of a family of standards for the X.509 PKI for the Internet, but the RFC is a standalone document; implementations of this standard proceed independent from the other parts. The RFC profiles the format and semantics of public key certificates and certificate revocation lists for the Internet. Procedures are described for the processing of certification paths in the Internet environment. Encoding rules are provided for popular cryptographic algorithms. Finally, Abstract Syntax Notation One (ASN.1) modules are provided in the appendices for all data structures defined or referenced.

3. *Threat Analysis of the Domain Name System (DNS)*, IETF RFC 3833, Atkins and Austein, August 2004. https://datatracker.ietf.org/doc/rfc3833.

   RFC 3833 attempts to document some of the known threats to the DNS, and, in doing so, measure the extent to which DNSSEC is a useful tool in defending against these threats.

4. *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*; Proposed Standard; IETF RFC 5280; Cooper, Santesson, Farrell, Boeyen, Housley, Polk; May 2008. https://datatracker.ietf.org/doc/rfc5280.

   RFC 5280 profiles the X.509 v3 certificate and X.509 v2 certificate revocation list (CRL) for use in the Internet. The RFC provides an overview and model of the specified approach, describes the X.509 v3 certificate format in detail, with additional information regarding the format and semantics of Internet name forms. Standard certificate extensions are described and two Internet-specific extensions are defined. A set of required certificate extensions is also specified, the X.509 v2 CRL format is described along with standard and Internet-specific extensions, an algorithm for X.509 certification path validation is described, and an ASN.1 module and examples are provided.

5. *Simple Mail Transfer Protocol*, IETF RFC 5321, Draft Standard, Kleinstein, October 2008. https://datatracker.ietf.org/doc/rfc5321.

   RFC 5321 is a specification of the basic protocol for Internet email transport. It covers the SMTP extension mechanisms and best practices for the contemporary Internet, but does not provide details about particular extensions. Although SMTP was designed as a mail transport and delivery protocol, this specification also contains information that is important to its use as a "mail submission" protocol for "split-UA" (User Agent) mail reading systems and mobile environments.

6. *Secure/Multipurpose Internet Mail Extensions (S/MIME)*, Version 3.2, Message Specification, Proposed Standard, IETF RFC 5751, ISSN: 2070-1721, Ramsdell and Turner, January 2010. https://datatracker.ietf.org/doc/rfc5751.

RFC 5751 defines S/MIME version 3.2. S/MIME provides a consistent way to send and receive secure MIME data. The RFC describes methods for digital signatures to provide authentication, message integrity, and non-repudiation with proof of origin; encryption to provide data confidentiality; and to reduce data size.

7. *Use Cases and Requirements for DNS-Based Authentication of Named Entities (DANE)*, IETF RFC 6394, ISSN: 2070-1721, Barnes, October 2011. https://datatracker.ietf.org/doc/rfc6394.

   Many current applications use the certificate-based authentication features in TLS to allow clients to verify that a connected server properly represents a desired domain name. Typically, this authentication has been based on PKI certificate chains rooted in well-known certificate authorities (CAs), but additional information can be provided via the DNS itself. This document describes a set of use cases in which the DNS and DNSSEC could be used to make assertions that support the TLS authentication process. The main focus of this document is TLS server authentication, but it also covers TLS client authentication for applications where TLS clients are identified by domain names.

8. *The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security Protocol: TLSA*, Proposed Standard, IETF RFC 6698, ISSN: 2070-1721, Hoffman and Schlyter, August 2012. https://datatracker.ietf.org/doc/rfc6698.

   Encrypted communication on the Internet often uses TLS, which depends on third parties to certify the keys used. RFC 6698 provides means to improve on that situation by standardizing on methods to enable the administrators of domain names to specify the keys used in that domain's TLS servers. This requires matching improvements in TLS client software, but no change in TLS server software.

9. *Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, Proposed Standard, IETF RFC 6818, ISSN: 2070- 1721, Yee, January 2013. https://datatracker.ietf.org/doc/rfc6818

   RFC 6818 updates RFC 5280, the *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. It changes the set of acceptable encoding methods for the explicit Text field of the user notice policy qualifier and clarifies the rules for converting internationalized name labels to American Standard Code for Information Interchange (ASCII). The RFC also provides some clarifications on the use of self-signed certificates, trust anchors, and some updated security considerations.

10. *SMTP security via opportunistic DANE TLS*, RFC 7672, Dukhovni and Hardaker, May 26, 2015. https://datatracker.ietf.org/doc/rfc7672

    The RFC describes a downgrade-resistant protocol for SMTP transport security between Message Transfer Agents, based on the DANE TLSA DNS record. Adoption of this protocol will

enable an incremental transition of the Internet email backbone to one using encrypted and authenticated TLS.

11. *Using Secure DNS to Associate Certificates with Domain Names For S/MIME*, RFC 8162, Hoffman and Schlyter, May 2017. https://datatracker.ietf.org/doc/rfc8162/

    The draft RFC for using secure DNS to associate certificates with domain names for S/MIME describes how to use secure DNS to associate an S/MIME user's certificate with the intended domain name; similar to the way that DANE (RFC 6698) does for TLS.

3.5 Technologies

The laboratory configuration employed for the project included components contributed by several sets of collaborating organizations. One of the component sets is Windows-based. The others are Linux-based. There were also three MUAs: Microsoft Outlook, Mozilla Thunderbird (on Linux), and a Thunderbird MUA equipped with a DANE-aware Apple Key Chain utility that were able to interact with all the mail servers via IMAP. While the Windows-based contribution used Server 2016 DNS services, the Linux-based contributions included three different implementations for DNS. One was based on NSD4 and Unbound authoritative and recursive servers, one was based on the Berkeley Internet Name Domain (BIND) DNS server, and one was based on the Secure64 DNS services. Secure 64 also contributed DNS services hosted on dedicated processors using SecureT micro operating system (OS) technology. Collaborators assisted in installation and initial configuration of products and, as necessary, in composition of components for different test cases.

Figure 3.1 below depicts, at a high level, collaborator contributions used to support the demonstration project. Elements identified in boldface are components provided or adapted by the collaborator. Other elements were incorporated into the stack to permit checking out the installed component's functionality.

Collaborator contributions identified below are organized with respect to the contributor as initially installed and checked out at the NCCoE. The architecture described in Section 4 below permits demonstration of the interconnection of components provided by different collaborators and initially checked out independently.

Figure 3.1 DNS-Based Email Security Collaborator Contributions

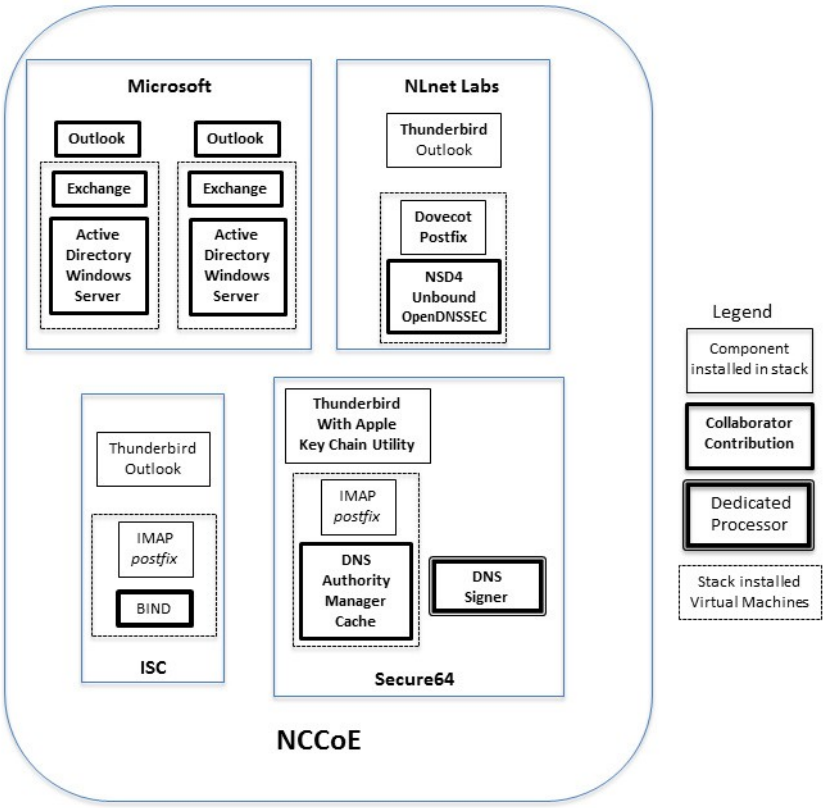

3.5.1 Microsoft

The Microsoft environments were contributed to support demonstration Scenario 1. Two environments were configured on the laboratory's VMware virtual machines (see Figure 3.1 above). Each stack included the ability to demonstrate Office Outlook[30] as an MUA, included Exchange Server 2016[31] as MTAs, and used Active Directory running on Microsoft Windows Server 2016[32] for DNS services. The Microsoft contribution included a DNSSEC-aware DNS recursive server, a DNSSEC-aware DNS authoritative server (IETF RFC 4033, 4034, and 4035), an MTA that can do SMTP over TLS (RFC 3207), management tools to configure servers and for debugging purposes, X.509 certificate sources, FIPS 140-

[30] https://en.wikipedia.org/wiki/Microsoft_Outlook
[31] https://products.office.com/en/exchange/microsoft-exchange-server-2016
[32] https://www.microsoft.com/en-us/cloud-platform/windows-server

2 validated cryptographic software, and support for multifactor authentication. The stacks were also able to be configured to demonstrate that Exchange could be used with either an Outlook or a Thunderbird MUA. Other test cases were demonstrated using Exchange with a combination of other providers' DNS implementations.

3.5.2 NLnet Labs

The NLnet Labs contribution focused on DNS services to support both demonstration scenarios. NLnet software was initially configured on the laboratory's VMware virtual machines. The components included NSD4 4.1.9[33], Unbound[34], and OpenDNSSEC[35] software for DNS services and Postfix and Dovecot for mail services. NSD4 is an authoritative only, high performance, open source name server. Unbound is a validating, recursive, caching DNS resolver. OpenDNSSEC is a set of software for signing DNS zones that are then served using NSD. While OpenDNSSEC can be configured to sign zone files or to sign zones transferred in via DNS zone transfer (AXFR), in these scenarios, it is used to sign local zone files in these scenarios. Like with the Microsoft stack above, multiple MUAs were configured to send and receive mail with the NLnet components via SMTP and IMAP.

3.5.3 Internet Systems Consortium (ISC)

The ISC contribution was focused on the BIND DNS server and supported both demonstration scenarios. BIND was initially configured on the laboratory's VMware virtual machines and included configuration for Postfix and Dovecot for email. BIND21 is open source software that is considered the reference implementation of DNS, but it is also production-grade software, suitable for use in high-volume and high-reliability applications. BIND features response rate limiting (RRL), support for FIPS 140-2 validated hardware cryptographic modules, the optional ability to retrieve zone data directly from an external database, the ability to use inline signing to automatically re-sign records as they are updated, and a scalable master/slave hierarchy. Like the other stacks, all three MUAs were able to connect and use the stack for DNS and email. BIND versions prior to BIND 9.11.0 are released under the ISC License (https://www.isc.org/downloads/software-support-policy/isc-license/).

3.5.4 Secure64

The Secure64 contributions were focused on DNS services to support both demonstration scenarios. The Secure64 environment included an automated online Secure64 DNS Signer as well as DNSSEC-capable VM images of DNS Cache, DNS Authority, and DNS Manager. DNS Manager provided centralized management of Secure64 DNS Cache software and configurations and provided network-wide monitoring of key performance indicators. DNS Manager allowed creation of groups of servers and

[33] https://www.nlnetlabs.nl/projects/nsd/
[34] http://unbound.net
[35] https://www.opendnssec.org

assignment of configurations to a group, a single server, or all servers. DNS Authority is an authoritative signer and server as a single platform. DNS Cache, DNS Authority, and DNS Manager were configured on the laboratory's VMware virtual machine; and the DNS Signer was provided as a high-assurance implementation delivered on a Secure64 dedicated appliance. Secure64 contributions were able to demonstrate Outlook, Thunderbird, or Thunderbird equipped with an Apple Key Chain utility as MUAs and use Postfix as an MTA and Dovecot to provide IMAP for clients.

4 Architecture

The Security platform architecture used for the project included combinations of components from different sources that supported two usage scenarios for DANE-enabled secure email in four different systems environments.

4.1 Usage Scenarios Supported

The scenarios supported include:

- "ordinary" email where the email exchanges between two organizations' email servers communicate over TLS with a STARTTLS extension, and relevant TLSA records are published in the receiver's DNS zone protected by DNSSEC
- end-to-end signed email, where the email exchanges between users in different organizations are carried over a channel protected by TLS (using the STARTTLS extension), and relevant artifacts used for signing and channel protection are published in a DNS zone protected by DNSSEC. Subsequently, these artifacts are used for S/MIME and TLS validation.

In both scenarios, end-entity and personal certificates were generated from CAs. Use of "well known" (i.e., installed as trust anchors in hosts), local enterprise CAs, and self-signed certificates were demonstrated.

While the second scenario demonstrated signing of emails, it does not include an end-to-end encrypted email scenario. Signing addresses the main security concerns in enterprise environments, which are the target of the project, but may neglect concerns of individual users who may also want to reduce information disclosure to their email providers. The two scenarios that are included may, however, serve as enablers for end-to-end encryption. Participation by parties having a primarily end-to-end encryption focus may succeed in generating industry support for the building blocks needed to support end-to-end encryption.

In more detail, the project's security platforms use the STARTTLS extension to include encryption of communications between two MTAs, as well as the signature of individual messages using S/MIME. The encryption and decryption with S/MIME on the end user's client was excluded from the current platform demonstration.

4.1.1 Usage Scenario 1

An individual needs to enter into an email exchange with an individual in another organization. Each individual exchanges email via the respective parent organization's mail servers. Users connect to their organization's respective mail servers within a physically protected zone of control.

In this scenario, the privacy policy of the parent organization requires encryption of the information being exchanged. The security afforded by the cryptographic process is dependent on the confidentiality of encryption keys. The mail servers are configured to use X.509 certificates to authenticate themselves during an encryption key establishment process. DNSSEC is employed to ensure that each sending mail server connects to the legitimate and authorized receiving mail server from which its X.509 certificate is obtained. DANE resource records are employed to bind the cryptographic keying material to the appropriate server name. STARTTLS is employed to negotiate the cryptographic algorithm to be employed with TLS in the email exchange in which the PII is transferred. Encryption of the email message is accomplished by the originator's email server, and decryption of the email message is accomplished by the recipient's email server.

Demonstrations of the security platform in this scenario include an attempt by a fraudulent mail server to pose as the legitimate receiver of the email and a man-in-the-middle attacker to attempt to disrupt the signal that TLS is available for the desired destination. In the latter attack, the goal is to force unencrypted transmission of the email. Both attempts should fail due to use of DNSSEC and DANE.

4.1.2 Usage Scenario 2

An individual needs to enter into an email exchange with an individual in another organization. Each individual exchanges email via the respective parent organization's mail servers. Users connect to their organization's respective mail servers within a physically protected zone of control.

The policy of the parent organization requires cryptographic digital signature of the message to provide integrity protection source authentication of the email message. S/MIME is a widely available and used protocol for digitally signing email. Each organization has therefore generated X.509 certificates for their users that include the public portion of their signature keys. These certificates are then published in the DNS using the appropriate DANE DNS Resource Record (RR) type.

DNSSEC is used to provide assurance that the originating user's mail server connects to the intended recipient's mail server. DANE records are employed to bind the cryptographic certificates to the appropriate server (for TLS) and individual user (for S/MIME), respectively. TLS is employed to provide confidentiality. Digital signature of the email message is accomplished by the originator's email client. Validating the signature (hence the integrity of the authorization provided in the email message) is accomplished by the recipient's email client.

Demonstrations of the security platform in this scenario include an attempt by a fraudulent actor to pose as the originator of the email and a man-in-the-middle attacker attempting to disrupt the validation of the S/MIME signature. Both attempts fail due to use of DNSSEC and DANE records.

4.2 Architectural Overview

The laboratory architecture for the project was designed to permit interconnection of Microsoft Outlook and Thunderbird MUAs with Microsoft Exchange and Postfix/Dovecot MTAs. It demonstrates the interconnection of either MTA with any of the DNS services contributed by collaborators. Two instantiations of each MTA type were established to demonstrate email exchanges between MTAs of the same type or different types. The various component combinations are then demonstrated with three different TLSA RR parameters: a self-signed certificate, use of local certificate authorities, and use of well-known certificate authorities.

Figure 4.1 is a deployment diagram of the architecture used for demonstrating DNS-based email security.

Figure 4.1 DNS-Based Email Security Deployment Diagram

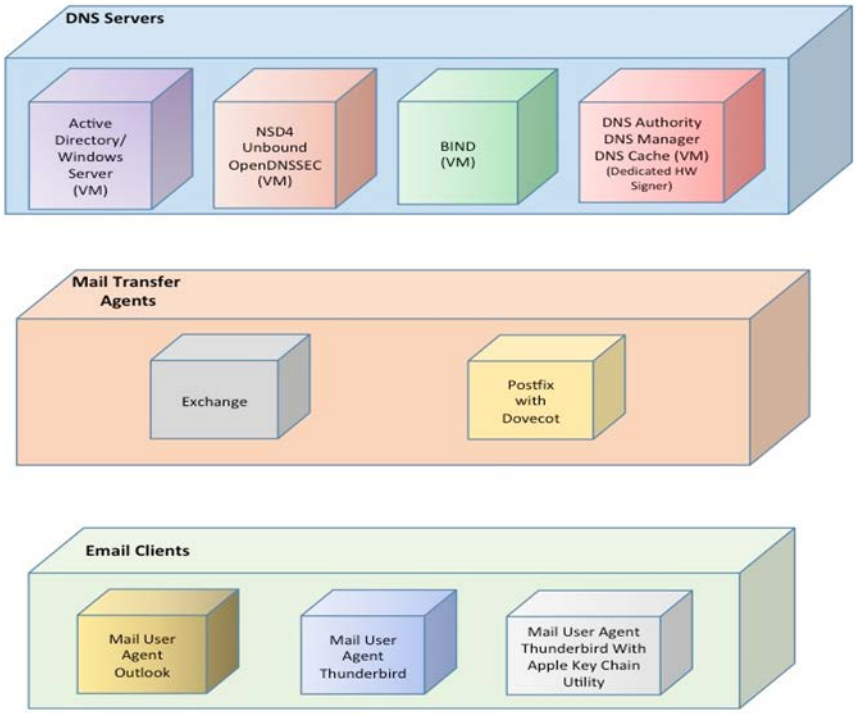

For test documentation purposes, the receiving MTA is named differently depending on the receiver's DNS service zone and the TLSA option being demonstrated. The sending MTA's implementation and DNS infrastructure can also vary for each test, but share the same basic processes.

The design of the environment permits interconnection of components provided by different collaborators (see Figure 4.2).

Figure 4.2 DNS-Based Email Security Test Setup

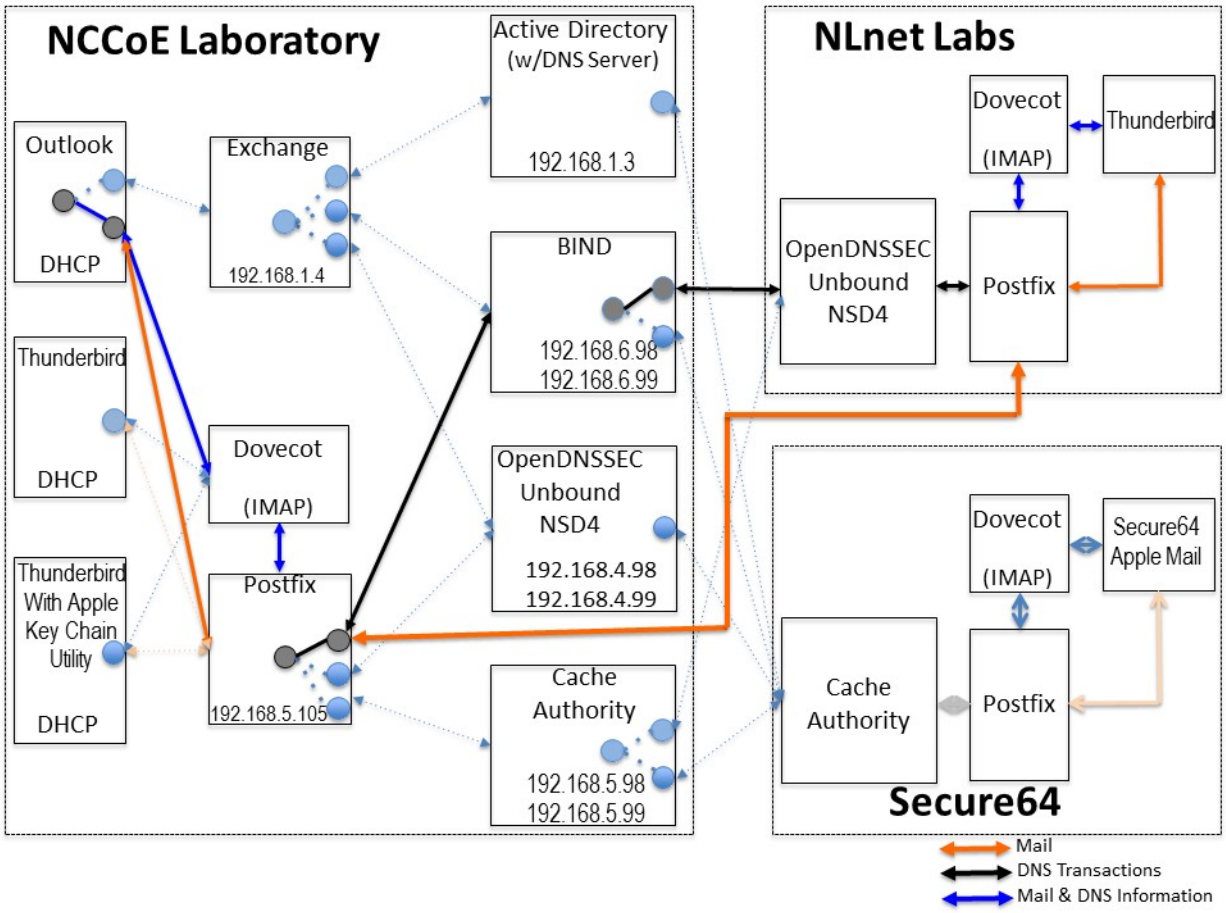

The depiction shows that the project security platform test/demonstration activity was based on three different clients, two MTAs, and four DNS service configurations in the lab at the NCCoE exchanging messages with NLnet Labs and Secure64. All messages were signed (a mail client function) and encrypted (server to server). We worked with one remote location at a time, driven by whichever is ready first. The message exchanges, including DNS activity will be logged at each end (lab and remote correspondent).

The solid connectors in the depiction illustrate one case. The dotted lines depict the other cases we want to demonstrate. A switch convention is used to reflect configuration options, but the project team actually configures each component for each option.

The orange arrows between the mail clients and the Postfix MTA reflect the fact that clients submitted email directly to the SMTP server for relay, while using Dovecot only to get mail. (The depiction in

Figure 4.2 reflects that IMAP is not used to submit mail, only retrieve it, so the MUA sent mail directly to the Postfix server, but received the reply through the Dovecot server.)

The project team demonstrated 30 different events using various combinations of MUA, MTA, and DNS Server components divided among five test sequences. In each sequence, signed and encrypted messages were sent from a sender to a recipient. Both Exchange and Postfix encrypted mail by default. Most of the exchanges employed either self-signed certificates or local CAs (see Appendix C). The BIND configuration was set up to obtain and validate certificates from the NIST Advanced Network Technologies Division's (ANTD's) DNS source (acting as a root CA). (See section 6 below for test sequence sets.)

In one test sequence, fraudulently signed S/MIME email was sent from a malicious sender to recipients using Outlook and Thunderbird MUAs configured to use Exchange and Postfix as MTAs. The Outlook/Exchange configuration used Active Directory as its DNS server. The configurations employing Postfix/Dovecot MTAs were demonstrated with each of the other three contributed DNS services. In one event, the Thunderbird MUA employed an Apple Key Chain Utility tool that allows a host to obtain X.509 certificates via of DANE RRs. All events were conducted using well-known CA and Enterprise CA-issued certificates for the impersonated sender. The fraudulent site attempted to spoof a valid sending domain belonging to a Secure64 site that was configured with DNS Authority/Cache/Signer DNS services, a Postfix/Dovecot MTA, and Thunderbird equipped with the Apple Key Chain utility. An Outlook/Exchange/ Active Directory set-up acted as the fraudulent site. The email exchange between organizations was carried over TLS, and the email message was S/MIME signed on the fraudulent user's client device. The setup for this sequence is depicted in Figure 4.3 below.

In another sequence, an NCCoE system attempted to send a TLS protected email from Exchange and Postfix MTAs (in turn) to an external Postfix MTA using DNS Authority/Cache/Signer for DNS services. The NCCoE Exchange MTA used Active Directory DNS Services, and the Postfix/Dovecot MTA used BIND and NSD4/Unbound/OpenDNSSEC DNS services. An S/MIME signed email was sent to an external Postfix MTA. Four events were conducted using Well-Known CA issued certificates, four events were conducted using Enterprise CA issued certificates (TLSA/SMIMEA RR parameter of CU=2) for TLS and S/MIME on the receiver side, and three events were conducted using self-signed certificates (TLSA/SMIMEA RR parameter of CU=3) for TLS and S/MIME on the receiver side. An Outlook/Exchange/Active Directory stack acted as a man-in-the-middle and attempted to intercept the message. Figure 4.4 depicts the configuration for a man-in-the-middle demonstration. Note that the sender is being misdirected to a malicious email server only. This is to simulate a lower-level attack where email is sent (via route hijacking or similar low-level attack) to a man-in-the-middle. Figure 4.4 depicts the configurations used with the Thunderbird/Postfix/Dovecot/Bind option selected.

Figure 4.3 Fraudulent DNS Address Spoofing Configurations

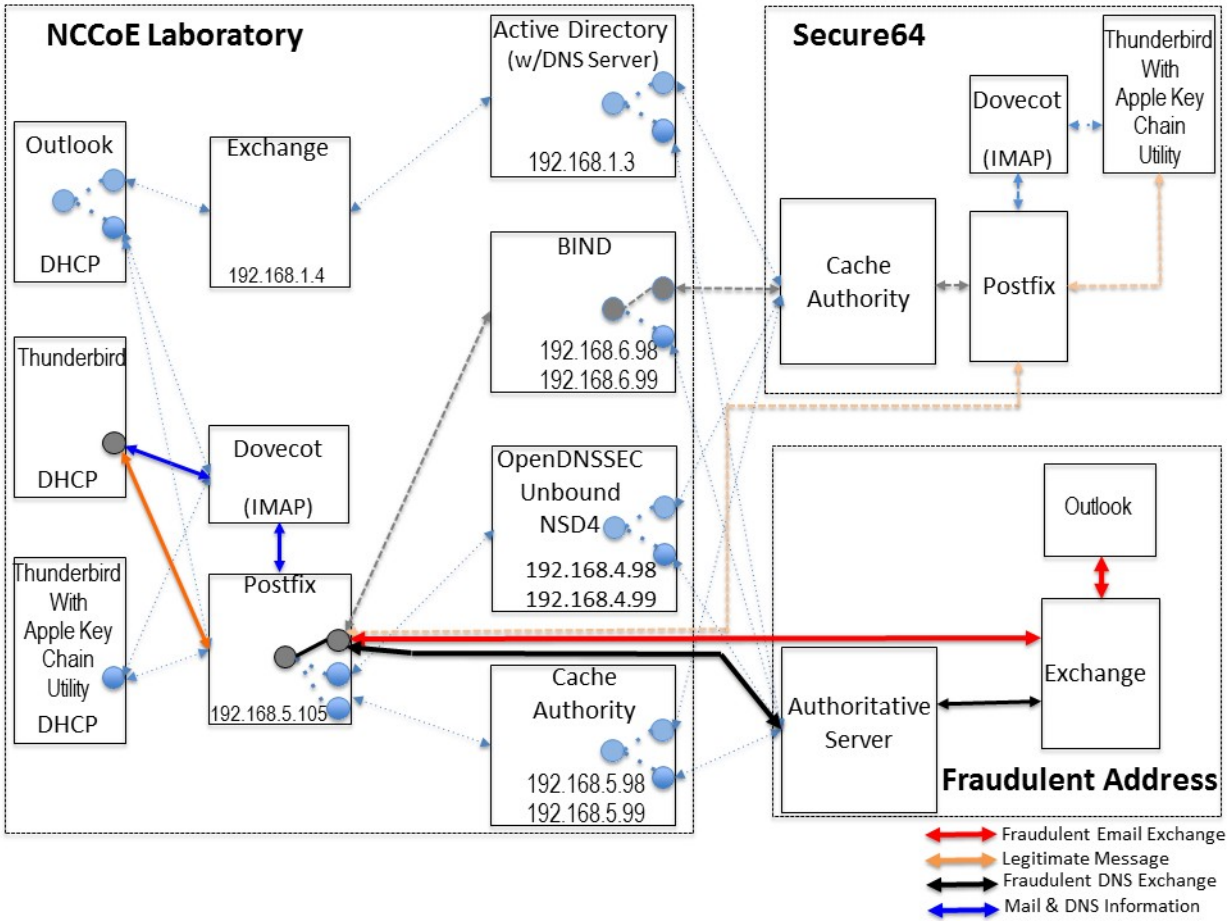

Figure 4.4 Man-In-The-Middle Event Configurations

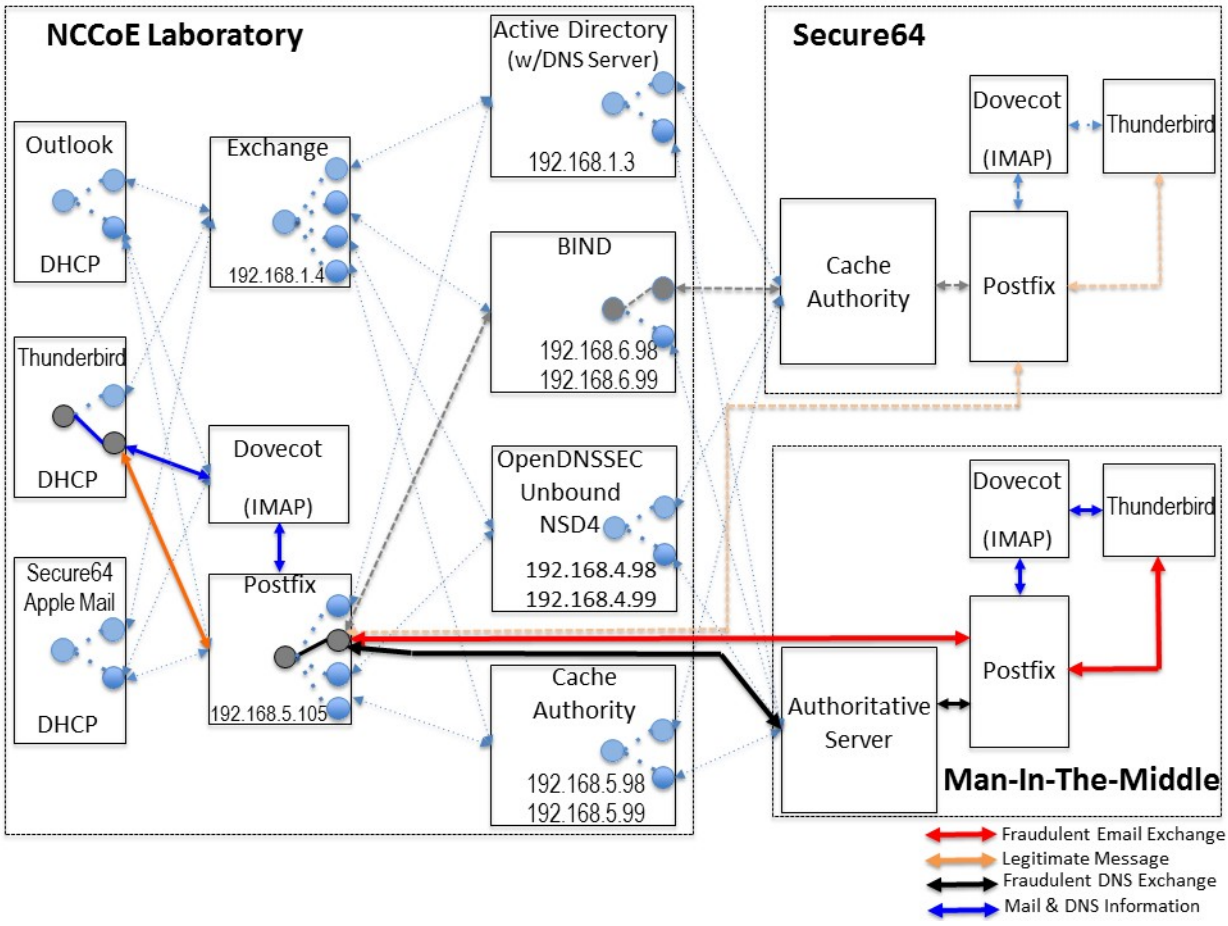

The following subsections describe the architecture's MUA, MTA, and DNS service components and Cybersecurity Framework Core Subcategories supported by those components.

4.2.1 Client Systems and MUAs

Client systems environments are Microsoft Office, Apple Mail, and open-source Linux-based Thunderbird applications. These include both commercial products and open-source software. MUA capabilities associated with the client systems are used to invoke S/MIME digital signature and signature verification for email, but user-to-user encryption is not demonstrated. Collaborators assisted in installation, integration tailoring as necessary, and testing of laboratory configurations.

Table 4.1 Client Systems

Application	Source	Collaborator Configuration Support	Cybersecurity Framework Category
Office Outlook Mail User Agent	Microsoft	Microsoft	PR.AC-1, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, RS.MI-2
Thunderbird Mail User Agent	Open (Mozilla)	NLnet Labs	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, RS.MI-2
Thunderbird with Apple Key Chain	Secure64	Secure64	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, RS.MI-2

4.2.2 Email Servers

Email servers include both Windows and Linux-based (Dovecot/Postfix) MTAs. Server-to-server encryption was demonstrated in the Postfix environments. Authentication of domain and server identity was based on DNSSEC-signed DANE records. Use of these DANE records is only supported by Postfix at the time of this project. The MTAs support each of the Cybersecurity Framework Functions, Categories, and Subcategories identified in section 3.4.4 above. The servers were demonstrated in different DNS environments and different TLSA RR usage scenarios. To demonstrate representative TLSA parameters, the demonstrations used self-signed certificates, end-entity certificates generated by well-known CAs and end-entities generated by enterprise local CAs.

Table 4.2 Mail Transfer Agents

Application	Source	Collaborator Configuration Support	Cybersecurity Framework Category
Exchange 2016[36] Mail Transfer Agent TLS Capable	Microsoft	Microsoft	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, PR.CM-1, DE.CM-6, DE.DP-4, DE.RP-1, RS.CO-2, RS.MI-1, RS-MI-2
Postfix Mail Transfer Agent TLS Capable DANE Capable	Open (postfix.com)	NLnet Labs Fraunhofer Secure64	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, PR.CM-1, DE.CM-6, DE.DP-4, DE.RP-1, RS.CO-2, RS.MI-1, RS-MI-2

4.2.3 DNS Servers

Both Windows and Linux-based DNS server and support components were contributed. DNS services provided include DNSSEC validating DNS resolvers (stub and recursive) and authoritative DNS servers for DNSSEC signed zones. Support for SMIMEA and TLSA records was demonstrated. The DNS server components support each of the Cybersecurity Framework Functions, Categories, and Subcategories identified in section 3.4.4 above with the exception of PR.DS-1 (protection of data-at-rest).

Application	Source	Collaborator Configuration Support	Cybersecurity Framework Category
Active Directory and Windows Server 2016 ■ Supports DNSSEC	Microsoft	Microsoft	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, PR.CM-1, DE.CM-6, DE.DP-4, DE.RP-1, RS.CO-2, RS.MI-1, RS-MI-2

[36] Exchange provided integrity protection only for PR.DS-1, PR.DS-2, and PR.PT-4 (Scenario 2).

Application	Source	Collaborator Configuration Support	Cybersecurity Framework Category
BIND ■ Supports DNSSEC ■ Supports DANE	Open (ISC)	Internet Systems Consortium (ISC)	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, PR.CM-1, DE.CM-6, DE.DP-4, DE.RP-1, RS.CO-2, RS.MI-1, RS-MI-2
NSD4 ■ Supports DNSSEC ■ Supports DANE Unbound ■ Supports DNSSEC OpenDNSSEC	Open (NLnet Labs)	Open (NLnet Labs)	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, PR.CM-1, DE.CM-6, DE.DP-4, DE.RP-1, RS.CO-2, RS.MI-1, RS-MI-2
DNS AUTHORITY DNS MANAGER ■ Supports DNSSEC ■ Supports DANE (Caching authority is labeled DNS CACHE, and signer runs on a dedicated processor)	Secure64	Secure64	PR.AC-1, PR.AC-5, PR.DS-1, PR.DS-2, PR.DS-6, PR.PT-4, PR.CM-1, DE.CM-6, DE.DP-4, DE.RP-1, RS.CO-2, RS.MI-1, RS-MI-2

5 Outcome

This section discusses the security platform from the perspective of the user and the system administrator. We define system administrator as a person within the organization who has elevated privileges on the management systems in the build. System administration functions include identification of system components, system installation, system integration, system configuration, configuration monitoring, identification of exception conditions, system maintenance, and status reporting to management.

5.1 The User's Experience

The user's experience varies from relatively minimal additional impact in enterprise environments with established system administration and support to a significant impact in the case of individual self-supported users. Where the enterprise offers systems administration and support services, the user's experience with respect to DNS services is essentially unchanged. One exception is that, where DNSSEC authentication fails, email messages sent to or by a user will not be delivered. This should be an uncommon experience for correspondents but it is up to the enterprise DNS administrator to prevent this happening. Errors due to DNSSEC validation failures are not sent back to the end user and may not be logged at the sending MTA, but at the validating recursive resolver which detected the error.

Similarly, for server-to-server encryption, the security protection features should be essentially transparent to the user.

For user-to-user digital signature, the user must first have a certificate installed in their MUA. This may be included in digital identity credentials, or it may be provided by the system administrator in the process of provisioning the user's computer. Otherwise, the procedure required would be similar to that followed in section 3.2 of SP 1800-6C. The steps required vary from platform to platform (e.g., Windows, Linux, Mac), user agent to user agent (e.g., Outlook vs Thunderbird) and how the private key is stored (on the system, smart cards, etc.) Representative user requirements are described below (in this case for Outlook running on MacBook and Thunderbird running on Linux.)

5.1.1 User's Digital Signature Experience with Outlook on MacBook

To use digital signatures and encryption, both the sender and recipient must have a mail application that supports the S/MIME standard (e.g., Outlook).

Note: Before this procedure is started, a certificate must be added to the keychain on the computer. For information about how to request a digital certificate from a certification authority, see macOS Help or click on "Help" on the Outlook tool bar.

1. On the **Tools** menu, click **Accounts**.
2. Click the account that is to be used to send a digitally signed message, click **Advanced**, and then click the **Security** tab.
3. Under **Digital signing**, on the **Certificate** pop-up menu, click the certificate that is to be used.

   *Note: The **Certificate** pop-up menu only displays certificates that are valid for digital signing or encryption that have already been added to the keychain for the macOS user account. To learn more about how to add certificates to a keychain, see macOS Help.*
4. Do any of the following:

To	Do this
Make sure that the digitally signed messages can be opened by all recipients, even if they do not have an S/MIME mail application and cannot verify the certificate	Select the **Send digitally signed messages as clear text** check box.
Allow the recipients to send encrypted messages to you	Make sure that signing and encryption certificates have been selected on this screen, and then select the **Include my certificates in signed messages** check box.

5. Click **OK**, and then close the **Accounts** dialog box.
6. In an email message, on the **Options** tab, click **Security**, and then click **Digitally Sign Message**.
7. Finish composing the message, and then click **Send**.

5.1.2 User's Digital Signature Experience with Thunderbird

For purposes of illustration, the description of the user experience with Thunderbird also included certificate management requirements. The example here shows both S/MIME and PGP examples of certificate management. The S/MIME approach is recommended. Note that when using OpenPGP, a FIPS 140-conformant version should always be used.

5.1.2.1 S/MIME Certificate Management

S/MIME certificates are used for digitally signed and (optionally) encrypted email messages. For information about getting or creating S/MIME certificates, see: http://kb.mozillazine.org/Getting_an_SMIME_certificate.

Installing an S/MIME certificate

Note: Before a user can create or import his or her own certificate and private key, he or she must first set a master password if this has not already been done. The master password is needed so that imported certificates are stored securely. See http://kb.mozillazine.org/Master_password for instructions for setting a master password. The user may have his or her own personal certificate and private key in a .p12 or .pfx file, and may wish to import it into Thunderbird. Once a Master Password has been set, the user can import/install a personal S/MIME certificate from a .p12 or .pfx file by doing the following steps.

1. Open the Certificate Manager by going to **Tools -> Options... -> Advanced -> Certificates -> Manage Certificates....**

2. Go to the tab named **Your Certificates**.
3. Click on **Import**.
4. Select the **PKCS12** certificate file (.pfx or .p12).
5. It will ask the user for the master password for the software security device. The user enters his or her master password and clicks **OK**.
6. Next, it will ask the user for the password protecting his or her personal certificate. If the user's .p12 or .pfx file has a password, he or she enters it here, otherwise leave this field empty. Then click **OK**.

The S/MIME certificate should now have been imported. If the certificate was not trusted, consult the instructions at http://kb.mozillazine.org/Thunderbird_:_FAQs_:_Import_CA_Certificate.

Configuring Thunderbird for using the certificate to sign email

Go to **Tools -> Account Settings...** in Thunderbird. Then find the account with the email address that matches the email address in the certificate that has just been installed. Choose **Security** under that account and select the certificate that has just been installed. The rest of the options should be self-explanatory. When the user selects a certificate in Account Settings, that selection only applies to the account's default identity or identities. There is no user interface for specifying certificates for an account's other identities. If desired, this can be worked around by editing the settings manually, copying the settings from an account's default identity to some other identity. The settings have names ending in: signing_cert_name, sign_mail, encryption_cert_name, and encryptionpolicy.

User installation of a self-signed S/MIME certificate

If the S/MIME certificate in a user's .p12 or .pfx file is a self-signed certificate for the user's own identity, then before that file can be installed into the tab named **Your Certificates**, the user must first install that certificate as a certificate authority in the **Authorities** tab. The PKCS12 certificate file will not install into the **Authorities** tab. The user will need a copy of a self-signed certificate that does not contain the user's private key. This is usually in the form of a .cer file. One way to obtain the .cer form of a certificate from the .p12 file is to use the Firefox Add-on Key Manager to extract the .cer certificate from the .p12 file. With that Add-on installed in Thunderbird, the user goes to **Tools -> Key Manager Toolbox -> Key Manager -> Your Keys**, select his or her key, selects **Export** and chooses **X.509** as file format.

1. Go to **Tools -> Options... -> Advanced -> Certificates -> Manage Certificates....**
2. Go to the **Authorities** tab.
3. Click on **Import**.

4. Select the **.cer** file.
5. It will ask the user for what purposes he or she wants to trust the certificate. Select **Trust this CA to identify email users**.
6. Click **OK** to complete the import.

*Note: Thunderbird automatically adds other people's S/MIME certificates to the **Other People's** tab of a user's Certificate Manager when he or she receives from them a digitally signed message with a valid signature and with an S/MIME certificate issued by a recognized and trusted CA. CA certificates that appear in Thunderbird's Authorities tab are recognized, and may also be trusted. CA certificates that do not appear in that tab are considered **unrecognized**. An S/MIME certificate that was issued by an unrecognized CA will not be automatically added to the **Other People's** tab of the user's Certificate Manager. If the user attempts to manually import an S/MIME certificate that was issued by an unrecognized CA, nothing will happen--literally. Thunderbird will not even display an error dialog. It will just not import the S/MIME certificate. This is generally not a problem when receiving an S/MIME certificate that was issued by a trusted CA, but could be a problem for a certificate that was issued by an unrecognized or untrusted CA, or for a certificate that is self-signed (i.e., it has no CA other than itself). So, before a user can import an S/MIME certificate that is issued by an unrecognized CA or is self-signed, he or she must first acquire and import the certificate for the issuing CA. In the case of a self-signed certificate, a .cer file needs to be acquired from the individual whose certificate the user wishes to add.*

5.1.2.2 Sending a Digitally Signed Email

1. Compose the message as usual.
2. To digitally sign a message, select **OpenPGP** from the Thunderbird menu and enable the **Sign Message** option.

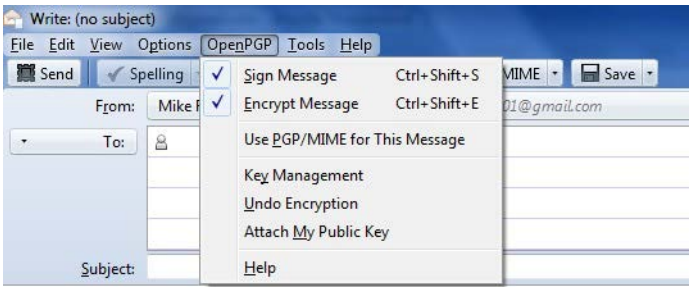

3. If the email address is associated with a cryptographic certificate, the message will be signed with the key contained in that certificate. If the email address is not associated with a cryptographic certificate, a certificate must be selected from a list.
4. Send the message as usual.

5.1.2.3 Reading a Digitally Signed Email

When a signed message is received, and If Thunderbird recognizes the signature, a green bar (as shown below) appears above the message. To determine whether or not the incoming message has been signed, look at the information bar above the message body.[37]

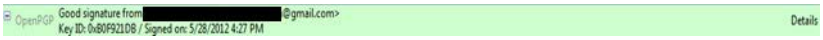

If the message has been signed, the green bar also displays the text, "Signed message". A message that has not been signed could be from someone trying to impersonate someone else.

5.2 The System Administrator's Experience

The system administrator(s) will generally be responsible for configuring the MUAs, MTA, and DNS servers. Specific installation and configuration instructions and examples are provided in Section 2, Section 3, Appendix F, Appendix G, and Appendix H of the How-To Guides, SP 1800-6C. Configuration includes setting up and publishing certificates in the DNS as TLSA and SMIMEA RRs. Certificate management using Well-Known CA-issued certificates or Enterprise CA-issued certificates is required for federal government applications and is strongly recommended in other applications. While instructions for configuration for DNSSEC are provided for environments described in SP 1800-6C, this more secure set of configuration options are not generally invoked by default. Therefore, more effort and expertise are needed on the part of the DNS administrator.

Configuring and activation of mail servers (MTAs) for channel encryption by default is described in section 3.3 of SP 1800-6C. Summary information is provided here and in links for illustration purposes for Microsoft Office 365 Exchange and Postfix.

In general, the bulk of the system administrator's effort is in acquiring and publishing the necessary certificates. Maintenance of the security functions, once they've been set up, is a relatively routine system administration activity.

5.2.1 Microsoft Exchange

Only Microsoft Exchange for Office 365 encrypts users' data while it is on Microsoft servers and while it is being transmitted between the MTAs. Exchange for Office 365 does provide controls for end users and administrators to fine tune what kind of encryption is desired to protect files and email communications.

[37] If the message is also encrypted on a user-to-user basis, Thunderbird will also ask for the entry of a secret passphrase to decrypt the message.

5.2.2 Postfix

Postfix TLS support is described at http://www.postfix.org/TLS_README.html. Postfix can be configured to always use TLS when offered by receivers.[38]

6 Security Characteristic Analysis

6.1 Assumptions and Limitations

This security characteristic evaluation has the following limitations:

- It is not a comprehensive test of all security components, nor is it a red team exercise.
- It cannot identify all weaknesses.
- It does not include the lab infrastructure. It is assumed that its devices are hardened. Testing these devices would reveal only weaknesses in implementation that would not be relevant to those adopting this reference architecture.

6.2 Build Testing

The evaluation included analysis of the security platforms to identify weaknesses and to discuss mitigations. The focus of this portion of the evaluation was hands-on testing of the laboratory build and examination of product manuals and documentation. Our objective was to evaluate the building block and not specific products. The presence of four primary OSs for domains tested (Linux, macOS, SourceT Micro OS, and Windows) made complete product-independent hands-on testing unrealistic.

Table 6.1 describes the goals of each sequence of test cases. For each sequence, the Cybersecurity Framework (CSF) Subcategories and associated SP 800-53 control(s), the test environment(s) involved, and evaluation objective of the test are identified. The results of the tests are provided in NIST SP 1800-6C.

In all test sequences, the sending MTA attempted to establish a TLS protected channel to deliver the email message to the receiver. In the attack scenarios, a malicious actor attempts to disrupt this transfer. In all test sequences, the sending MUA signed the message, and the receiving MUA, checked the signature. Exchange was used only for Scenario 2.[39] In all test sequences, the sending MTA attempted to verify the correctness of all DNS responses via DNSSEC validation. In most scenarios, alice@<somedomain> sent an email to bob@<receivername>. Both senders and receivers had their own (separate) DNS infrastructures consisting of both authoritative and recursive servers. The Exchange

[38] "Setting Postfix to encrypt all traffic when talking to other mail servers," *Snapdragon Tech Blog*, August 9, 2013. http://blog.snapdragon.cc/2013/07/07/setting-postfix-to-encrypt-all-traffic-when-talking-to-other-mailservers/

[39] Exchange MTAs did not attempt to encrypt or decrypt MTA-to-MTA message exchanges.

as Sender tests were conducted for completeness and for examples of SMTP over TLS without DANE support—what it looked like and how well it worked.

Table 6.1 Tests Performed

Test Sequence	CSF Subcategories	SP 800-53 Controls	Configuration	Evaluation Objective
Sequence 1	PR.AC-1 PR.AC-5 PR.DS-1 PR.DS-2 PR.DS-6 RS.MI-2	AC-2, AC-17, AC-19, AC-20, IA Family, IR-4, SC-8, SC-28, SI-7	An Outlook MUA, interfacing with an Exchange MTA, was configured to use Active Directory and BIND DNS services in turn. Each of the six configurations exchanged email with ■ a Secure64 MUA/MTA/DNS service stack that included a Postfix MTA and a Thunderbird MUA running on a Mac OS system ■ an NLnet Labs MUA/MTA/DNS service stack that included a Postfix MTA and a Thunderbird MUA running on Linux The events include those showing use of Well-Known CAs (Certificate Usage Type 1 (CU=1)), Enterprise CAs (CU=2), and Self-Signed Certificates (CU=3) for TLS and S/MIME-enabled mail receivers and S/MIME. Figure 4.2 above depicts the set-up for laboratory support for the Secure64 destination variant of this test sequence.[40]	Email messages between Postfix MTAs were encrypted and successfully decrypted via TLS (Scenario 1). Signature was logged. All messages were S/MIME signed. Outlook attempted to verify received messages (Scenario 2). Signature verification results were noted. DNS name verification results were noted.

[40] The connections depicted in the figure are actually for the Secure64 variant of the first Sequence 2 configuration. Capabilities for Sequence 1 support are shown as dotted lines.

Test Sequence	CSF Subcategories	SP 800-53 Controls	Configuration	Evaluation Objective
Sequence 2	PR.AC-1 PR.AC-5 PR.DS-1 PR.DS-2 PR.DS-6 RS.MI-2	AC-2, AC-17, AC-19, AC-20, IA Family, IR-4, SC-8, SC-28, SI-7	Outlook and Thunderbird MUAs, configured to use a Postfix MTA with Dovecot IMAP support, were configured in turn to use BIND and Secure64's DNS Authority, DNS Cache, and DNS Signer implementations. Each of the six configurations exchanged email with a Secure64 MUA/MTA/DNS service stack that included a Thunderbird MUA, Postfix/Dovecot MTA, and DNS Signer/DNS Cache/DNS Authority services for processing received messages; and an NLnet Labs MUA/MTA/DNS service stack that included a Thunderbird MUA, Postfix/Dovecot MTA, and NSD4, Unbound, and OpenDNSSEC DNS services. The test events include using Well-Known CA issued (TLSA/SMIMEA CU=1), Enterprise CA issued (CU=2), and Self-Signed Certificates (CU=3). Figure 4.2 above depicts the setup for laboratory support for this test sequence.	Email messages between MTAs were encrypted and successfully decrypted (Scenario 1). Signature and encryption were logged. All messages were S/MIME signed. Outlook attempted to verify received messages (Scenario 2). Signature verification results were noted. DNS name verification results were noted.

Test Sequence	CSF Subcategories	SP 800-53 Controls	Configuration	Evaluation Objective
Sequence 3	PR.AC-1 PR.AC-5 PR.DS-2 RS.MI-1	AC-2, AC-4, AC-17, AC-19, AC-20, IA Family, IR-4, SC-7, SC-8	Fraudulently S/MIME-signed email was sent from a malicious sender to recipients using Outlook and Thunderbird MUAs configured to use Exchange and Postfix as MTAs. The Outlook/Exchange configuration used Active Directory as its DNS server. The configurations employing Postfix/Dovecot MTAs were demonstrated with each of the other three contributed DNS Services. In one event, the Thunderbird MUA employed an Apple Key Chain Utility tool that allows a host to obtain X.509 certificates via of DANE RRs. All events were conducted using well-known CA and Enterprise CA-issued certificates for the impersonated sender. The setup for this sequence is depicted in Figure 4.3 above.	The fraudulent site attempted to spoof a valid sending domain belonging to a Secure64 site. An Outlook/Exchange/ Active Directory setup acted as the fraudulent site. The email exchange between organizations was carried over TLS, and the email message was S/MIME signed on the fraudulent users' client device. Where Well-Known CA-issued certificates or Enterprise CA-issued certificates were used, and the MTA was DANE aware, the MUA using a SMIMEA utility was able to detect the fraudulent email and mark the email as not validated.
Sequence 4	PR.AC-1 PR.AC-5 PR.DS-2 PR.DS-6 RS.MI-1 RS.MI-2	AC-2, AC-4, AC-17, AC-19, AC-20, IA Family, IR-4, SC-7, SC-8, SI-7	The sender used an Outlook MUA sending mail through a Postfix/Dovecot MTA and using (in turn): Active Directory and DNS Server, BIND DNS Server, and NLnet Labs DNS Services. Self-signed certificates were used on the legitimate receiver side (TLSA RR parameter CU=3) for TLS. Each of the three configurations attempted to initiate an email exchange with an external Secure64 site. The setup for this sequence is depicted in Figure 4.4 above.	The Outlook/Exchange/Active Directory stack attempted to intercept the email from the NCCoE Laboratory Configuration by acting as a man-in-the-middle. The email and DNS transactions were logged in each case, and the results are provided in Volume C Appendix C. Where the MTA was DANE-aware, spoofing was detected. The mail connection to the MTA was established but closed the connection before the mail was transferred. Otherwise, the MTA failed to detect the man-in-the-middle and sent the email.

Test Sequence	CSF Subcategories	SP 800-53 Controls	Configuration	Evaluation Objective
Sequence 5	PR.AC-1 PR.DS-6 DE.CM-1 DE.DP-4 RS.CO-2	AC-4, IR-5, SC-5, SC-20, SC-21, SC-23, SI-4, SI-13	A DANE-enabled Postfix MTA sent message traffic to four MTAs with one Authoritative Server serving all four zones. An NSD4 Authoritative DNS server and Unbound recursive server were provided for the Postfix sending MTA, and a Secure64 DNS Authority and Signer provided the DNS services for the recipient zones. We reviewed the log files. One of the recipient MTAs did not employ TLSA, one employed a valid TLSA with the CU set to 3, one employed a TLSA with a certificate usage field of 1, but with an incomplete (i.e., bad) PKI certification path (PKI X.509 [PKIX] failure), and one employed mismatched server cert/TLSA with the certificate usage field set to 3 (DANE validation failure).	A large number of email messages are generated in the Postfix server device using a Python script, and the Postfix MTA sends the messages to each of four recipient MTAs in different zones. In the recipient MTA running without TLSA and that running with a valid matching TLSA and certificate usage field set to 3, all messages should be accepted. In the recipient MTA with a TLSA RR using certificate usage of 1, but with an incomplete PKIX validation path, and the recipient MTA with a mismatched certificate/TLSA (cert usage 3), the sender should close the connection without sending the message. Logwatch running on the sending Postfix server device logged the instances of failure to deliver due to certificate expiration or bad certificate path.

6.3 Scenarios and Findings

One aspect of our security evaluation involved assessing how well the reference design addresses the objectives of the scenario it was intended to support.

6.3.1 Scenario 1

Scenario 1 involved the ordinary exchange of email between two organizations' email servers carried over TLS, where the TLS key management was protected by DANE and DNSSEC. Private certificates were generated by either well-known CAs, enterprise local CAs or self-signed. User connections to their organizations' respective mail servers were established and maintained within a physically protected zone, and email was encrypted between mail servers using TLS. The confidentiality of encryption keys was maintained such that no unauthorized third party had access to the keys. The mail servers used X.509 certificates to store and transport public keys to establish the TLS channel. DNSSEC ensured that each sending mail server receives the IP address to the legitimate and authorized receiving mail server and (if applicable) validate its X.509 certificate. DANE bound the cryptographic keying material to the appropriate server. TLS was used to protect the confidentiality of the email exchange. Encryption of the email message was accomplished by the originator's email server, and decryption of the email message was accomplished by the recipient's email server using standard server libraries.

The tests included an attempt by a fraudulent mail server to pose as the legitimate mail receiver for a domain. The tests also include a man-in-the-middle attack to attempt to disrupt the TLS connection with the objective of achieving an unencrypted transmission of the email. Both attempts failed due to use of DNSSEC and DANE. In both cases, an indication was made available to the sending email server when the DNSSEC signature associated with the domain data is determined to be invalid.

6.3.2 Scenario 2

Scenario 2 involved end-to-end signed email, where the email exchanges between organizations were carried over TLS as in (1), the email messages were signed and verified with S/MIME on the end users' client devices, and the S/MIME key management was protected by DANE and DNSSEC. Private certificates were generated by well-known and enterprise local CAs. Self-signed certificates were not used. Individuals established connections to their domains' respective mail servers within a physically protected zone of control. Cryptographic digital signatures were applied to messages to provide authentication and integrity protection for the email. S/MIME was the protocol used for the digital signing. These certificates were then encoded in the DNS using the appropriate DANE DNS record type. DNSSEC ensured that each originating user's mail server connects to the intended recipient's mail server. DANE bound the cryptographic keying material to the appropriate server and individual user digital signature certificates. TLS was employed to protect the confidentiality of the email. Digital signing of email messages was accomplished by the originator's MUA, and checking the validity of the signature

(hence the integrity of the authorization provided in the email message) was accomplished by the recipient's MUA.

The tests in this scenario included an attempt by a fraudulent actor to pose as an originator of the email. This attempt failed due to use of DNSSEC and DANE. The receiving MUA, using a third party SMIMEA tool, was able to fetch the sender's real S/MIME certificate from the DNS and confirm that the fraudulent email was signed using a different certificate.

6.3.3 Effects of DANE Errors

In addition to the scenarios described above, a DANE-enabled Postfix MTA sent message traffic to four other postfix MTAs. A single BIND instance was set up to serve the TLSA and A RRs for the four receivers. One of the receiving MTAs did not employ DANE. The second employed DANE with a valid TLSA with the certificate usage field2 set to 3. The third employed a TLSA with a certificate usage field of 2, but with an incomplete (i.e. bad) PKI certification path (generating a PKIX validation failure). The TLSA contained a local enterprise trust anchor, but the server did not have the full certificate chain (missing intermediate certificate). The final one employed DANE with a TLSA RR using Certificate Usage of 3, but there was a mismatch between the server cert and TLSA RR (generating a DANE validation failure).

Little or nothing appeared in the sender's logs for messages sent to either the MTA not employing TLS or the employing a valid TLSA. The growth rates for logs for the MTA that employed a TLSA with a certificate usage field of 1, but with a PKIX failure and the one that employed mismatched server cert/TLSA (i.e., DANE validation failure) were measured.

When the sender was configured to never use TLS, the mail was sent in plaintext regardless of the TLS/DANE configuration of the receiver. When the sender was configured to use TLS opportunistically, it used TLS regardless of the status of the certificate, or TLSA. In fact, the sender did not issue a query to find TLSA RRs even if published. When the sender used opportunistic DANE, it used TLS when available regardless of the DANE validations results. If validation failed, the mail was still sent and the result was logged as an "Untrusted" or "Anonymous" TLS connection, depending on the presence of a TLSA RR.

Of the four options used in the lab, "dane-only" is the most rigorous in what a sender would accept before sending mail. When the receiver did not offer the STARTTLS option, or lacked a TLSA RR, mail was not sent. Likewise, if a TLSA RR was present, but there was an error in validation (either the TLSA RR itself had an error, or PKIX failed), the mail was not sent. Therefore, use of this option is not recommended for general use as this will result in the majority of email being deferred. It should only be used in scenarios where senders and receivers are coordinated and maintain a stable DANE deployment.

7 Future Build Considerations

Both public sector and private sector enterprises are heavily dependent on web-based technology other than email for e-commerce and other public-facing applications. Fraudulent web sites pose at least as great a security and privacy problem as fraudulent email. Further, as email becomes a more difficult medium for malicious entities to use as a penetration vector, other web-based media will be more intensively exploited. Already, emerging communications trends appear to be replacing email exchanges among individuals with other social media (e.g., Baidu, Facebook, Facebook Messenger, Google+, Instagram, LinkedIn, Pinterest, Snapchat, Tieba, Tumblr, Twitter, Viber, WhatsApp, and YouTube). Therefore, an extension of the current project that focuses on use of improved DNSSEC applications such as DANE for web applications other than mail may be justified.

Additionally, the test scenarios did not include the Exchange for Office 365 MTA to demonstrate Scenario 1. Future builds might be considered to demonstrate this capability.

Finally, utilities are currently under development that would provide improved support for SMIMEA and improved system notification of failed DNSSEC signature validation events. Future builds might be considered to demonstrate these capabilities as well.

Appendix A List of Acronyms

AES	Advanced Encryption Standard
ANTD	Advanced Network Technologies Division
ASCII	American Standard Code for Information Interchange
ASN.1	Abstract Syntax Notation One
AXFR	DNS Full Zone Transfer Query Type
BGP	Border Gateway Protocol
BIND	Berkeley Internet Name Domain
CA	Certificate Authority
CKMS	Cryptographic Key Management System
CRADA	Cooperative Research and Development Agreement
CRL	Certificate Revocation List
CSF	Cybersecurity Framework
CU	Certificate Usage Type
CVE	Common Vulnerabilities and Exposures
DANE	DNS-Based Authentication of Named Entities
DARPA	Defense Advanced Research Projects Agency
DES	Data Encryption Standard
DNS	Domain Name System
DNSSEC	DNS Security Extensions
Email	Electronic Mail
EMC	Electromagnetic Compatibility
EMI	Electromagnetic Interference
FCKMS	Federal Cryptographic Key Management System
FIPS	Federal Information Processing Standard
FOIA	Freedom of Information Act

HIPAA	Health Insurance Portability and Accountability Act
HTTPS	Hypertext Transfer Protocol Secure
IDIQ	Indefinite Delivery/Indefinite Quantity
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IoT	Internet of Things
IP	Internet Protocol
IPsec	Internet Protocol Security
IRS	Internal Revenue Service
ISC	Internet Systems Consortium
ISO	Internet Organization for Standardization
IT	Information Technology
ITL	Information Technology Laboratory
MIME	Multipurpose Internet Mail Extension
MTA	Mail Transfer Agent
MUA	Mail User Agent
MX	Mail Exchange (Resource Record)
NCCoE	National Cybersecurity Center of Excellence
NIST	National Institute of Standards and Technology
NVD	National Vulnerability Database
OMB	Office of Management and Budget
OS	Operating System
PKI	Public Key Infrastructure
PKIX	Public Key Infrastructure X.509
RFC	Request for Comments
RMF	Risk Management Framework

RR	Resource Record
RRL	Response Rate Limiting
S/MIME	Secure/Multipurpose Internet Mail Extensions
SHA	Secure Hash Algorithm
SMIMEA	S/MIME Certificate Association (Resource Record)
SMTP	Simple Mail Transfer Protocol
SP	Special Publication
SQL	Structured Query Language
TLS	Transport Layer Security
TLSA	TLS Certificate Association (Resource Record)
UA	User Agent
VM	Virtual Machine

Appendix B References

Securing the Federal Government's Domain Name System Infrastructure, Executive Office of the President, Office of Management and Budget, M-08-23, August 22, 2008. https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2008/m08-23.pdf

Enhancing the Security of Federal Information and Information Systems, Executive Office of the President, Office of Management and Budget, M-14-03, November 18, 2013. https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2014/m-14-03.pdf

Improving Critical Infrastructure Cybersecurity, Executive Office of the President, Executive Order 13636, February 12, 2013. https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity

Federal Information Security Management Act, United States Congress, Public Law 107-347, December 17, 2002. https://www.govtrack.us/congress/bills/107/hr2458

Gramm-Leach-Bliley Act, United States Congress, Public Law 104-191, August 21, 1996. https://www.gpo.gov/fdsys/pkg/PLAW-106publ102/html/PLAW-106publ102.htm

Health Insurance Portability and Accountability Act, United States Congress, Public Law 106-102, November 12, 1999. https://aspe.hhs.gov/report/health-insurance-portability-and-accountability-act-1996

Managing Information as a Strategic Resource, OMB Circular A-130, Executive Office of the President, Office of Management and Budget, July 28, 2016. https://www.federalregister.gov/documents/2016/07/28/2016-17872/revision-of-omb-circular-no-a-130-managing-information-as-a-strategic-resource

Rules Governing Practice before the Internal Revenue Service, Internal Revenue Service, Circular Number 230, Revised June 2014. https://www.irs.gov/tax-professionals/circular-230-tax-professionals

Security Requirements for Cryptographic Modules, Federal Information Processing Standard (FIPS) 140-2, May 2001 (including change notices as of 12-03-2002). https://doi.org/10.6028/NIST.FIPS.140-2

Guide for Conducting Risk Assessments, NIST Special Publication 800-30 Revision 1, Joint Transformation Initiative, September 2012. https://doi.org/10.6028/NIST.SP.800-30r1

Guide for Applying the Risk Management Framework to Federal Information Systems: A security Lifecycle Approach, NIST Special Publication 800-37 Rev. 1, Joint Task Force Transformation Initiative; February 2010 with updates as of June 5, 2014. https://doi.org/10.6028/NIST.SP.800-37r1

Managing Information Security Risk: Organization, Mission, and Information System View, NIST Special Publication 800-39, Joint Task Force Transformation Initiative, March 2011. https://doi.org/10.6028/NIST.SP.800-39

Guidelines on Electronic Mail Security; NIST Special Publication 800-45 Ver. 2; Tracy, Jansen, Scarfone, Butterfield; February 2007. https://doi.org/10.6028/NIST.SP.800-45ver2

Federal S/MIME V3 Client Profile, NIST Special Publication 800-49, Chernick, November 2002. https://doi.org/10.6028/NIST.SP.800-49

Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations; NIST Special Publication 800-52 Rev. 1; Polk, McKay, Chokhani; April 2014. https://doi.org/10.6028/NIST.SP.800-52r1

Security and Privacy Controls for Federal Information Systems and Organizations, NIST Special Publication 800-53 Rev. 4, Joint Task Force Transformation Initiative, April 2013. https://doi.org/10.6028/NIST.SP.800-53r4

Recommendation for Key Management: Part 1 - General, NIST Special Publication 800-57 Part 1 Rev.4, Barker, January 2016. https://doi.org/10.6028/NIST.SP.800-57pt1r4

Recommendation for Key Management: Part 2 - Best Practices for Key Management Organization, NIST Special Publication 800-57 Part 2, Barker, Barker, Burr, Polk, and Smid, August 2005. https://doi.org/10.6028/NIST.SP.800-57p2

Recommendation for Key Management: Part 3: Application-Specific Key Management Guidance, NIST Special Publication, SP 800-57 Part 3 Rev. 1, Barker and Dang, January 2015. https://doi.org/10.6028/NIST.SP.800-57pt3r1

Electronic Authentication Guideline; NIST Special Publication 800-63-2; Burr, Dodson, Newton, Perlner, Polk, Gupta, Nabbus; August 2013. doi:10.6028/NIST.SP.800-63-2 [Direct Link]

Digital Identity Guidelines; NIST Special Publication 800-63-3; Burr, Choong, Danker, Grassi, Garcia, Greene, Fenton, Lefkovitz, Nadeau, Netwon, Perlner, Regenscheid, Richer, Squire, Theofanos; June 2017. doi:10.6028/NIST.SP.800-63-3 https://pages.nist.gov/800-63-3/

Secure Domain Name System (DNS) Deployment Guide, NIST Special Publication 800-81-2, Chandramouli and Rose, September 2013. https://doi.org/10.6028/NIST.SP.800-81-2

A Framework for Designing Cryptographic Key Management Systems; NIST Special Publication 800-130; Barker, Branstad, Smid, Chokhani; August 2013. https://doi.org/10.6028/NIST.SP.800-130

A Profile for U.S. Federal Cryptographic Key Management Systems (CKMS); NIST Special Publication 800-152; Barker, Smid, Branstad; October 2015. https://doi.org/10.6028/NIST.SP.800-152

Systems Security Engineering: Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems, NIST Special Publication 800-160, Ross, McEvilley, Oren, November 2016. https://doi.org/10.6028/NIST.SP.800-160.

Trustworthy Email; NIST Special Publication 800-177; Chandramouli, Garfinkel, Nightingale and Rose; September 2016. https://doi.org/10.6028/NIST.SP.800-177

"Internet of Things: Standards and Guidance from the IETF", *IETF Journal*, Keränen and Bormann, April 2016. https://www.ietfjournal.org/internet-of-things-standards-and-guidance-from-the-ietf/

X.509 Certificate Policy for the U.S. Federal PKI Common Policy Framework, Version 1.24. https://www.idmanagement.gov/wp-content/uploads/sites/1171/uploads/Common-Policy-Framework.pdf

Internet Protocol, RFC 791, Defense Advanced Research Projects Agency (DARPA), September 1981. https://datatracker.ietf.org/doc/rfc791

Domain Names - Concepts and Facilities, RFC 1034, Mockapetris, November 1987. https://datatracker.ietf.org/doc/rfc1034

Domain Name System Structure and Delegation, RFC 1591, Postel, March 1994. https://datatracker.ietf.org/doc/rfc1591

Internet X.509 Public Key Infrastructure Certificate and CRL Profile, RFC 2459, Housley, Ford, Polk, and Solo, January 1999. https://datatracker.ietf.org/doc/rfc2459

The Secure HyperText Transfer Protocol, RFC 2660, Rescorla and Schiffman, August 1999. https://datatracker.ietf.org/doc/rfc2660

Threat Analysis of the Domain Name System (DNS), RFC 3833, Atkins and Austein, August 2004. https://datatracker.ietf.org/doc/rfc3833

A Method for Storing IPsec Keying Material in DNS, RFC 4025, Richardson, February 2005. https://datatracker.ietf.org/doc/rfc4025

DNS Security Introduction and Requirements, RFC 4033, Arends, Austein, Larson, Massey, and Rose, March 2005. https://datatracker.ietf.org/doc/rfc4033

A Border Gateway Protocol 4 (BGP-4), RFC 4271, Rekhter, Li, and Hares, January 2006. https://datatracker.ietf.org/doc/rfc4271

The Transport Layer Security (TLS) Protocol Version 1.2, RFC 5246, Dierks and Rescorla, August 2008. https://datatracker.ietf.org/doc/rfc5246

Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, Proposed Standard, RFC 5280, Cooper, Santesson, Farrell, Boeyen (Entrust), Housley, and Polk, May 2008. https://datatracker.ietf.org/doc/rfc5280/

Simple Mail Transfer Protocol, RFC 5321, Draft Standard, Kleinstein, October 2008. https://datatracker.ietf.org/doc/rfc5321

Secure/Multipurpose Internet Mail Extensions (S/MIME), Version 3.2, Message Specification, Proposed Standard, RFC 5751, Ramsdell and Turner, January 2010. https://datatracker.ietf.org/doc/rfc5751

Use Cases and Requirements for DNS-Based Authentication of Named Entities (DANE), RFC 6394, Barnes, October 2011. https://datatracker.ietf.org/doc/rfc6394

The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security Protocol: TLSA, Proposed Standard, RFC 6698, Hoffman and Schlyter, August 2012. https://datatracker.ietf.org/doc/rfc6698

DNS-Based Service Discovery, RFC 6763, Cheshire and Krotchmal, February 2013. https://datatracker.ietf.org/doc/rfc6763

Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, Proposed Standard, RFC 6818, Yee, January 2013. https://datatracker.ietf.org/doc/rfc6818

SMTP Security via Opportunistic DNS-Based Authentication Named Entities (DANE) Transport Layer Security (TLS), RFC 7672, Dukhovni and Hardaker, May 26, 2015. https://datatracker.ietf.org/doc/rfc7672

Using Secure DNS to Associate Certificates with Domain Names for S/MIME, RFC 8162, Hoffman and Schlyter, May 2017. https://datatracker.ietf.org/doc/rfc8162/

Domain Name System-Based Security for Electronic Mail, Barker, National Institute of Standards and Technology's Dakota Consulting Indefinite Delivery/Indefinite Quantity (IDIQ) Contract SB1341-12-CQ-0011, Task Order 15- 421 Task 3 Report #2, December 17, 2016. https://nccoe.nist.gov/library/dns-based-secured-email

Task 2: Report #1 on Standards Review and Support for NCCoE Project Activities, Barker, National Institute of Standards and Technology's Dakota Consulting IDIQ Contract SB1341-12-CQ-0011, Task Order 15-421 Task 2 Report #1, November 30, 2015.

Task 3: Report #1 on Standards Review and Support for NCCoE Project Activities, Barker, National Institute of Standards and Technology's Dakota Consulting IDIQ Contract SB1341-12-CQ-0011, Task Order 15-421 Task 3 Report #1, November 30, 2015.

Appendix C Project Mapping to the Framework Core and Informative References

The following tables map informative NIST and consensus security references to Framework Core Subcategories that are addressed by the project's platform set. The references do not include protocol specifications that are implemented by the individual products that comprise the demonstrated security platforms. While some of the references provide general guidance that informs implementation of referenced Framework Core functions, the NIST Special Publication references provide specific recommendations that should be considered when composing and configuring security platforms from DNS and email components, implementing DNSSEC and mail security platforms, and operating email systems securely.

Table C.1 PROTECT (PR)

Category	Subcategory	Informative References
Data Security (PR.DS): Information and records (data) are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information.	**PR.DS-1**: Data-at-rest is protected	**FIPS 140-2** Sec. 4 **NIST SP 800-53 Rev. 4** SC-28 **NIST SP 800-57 Part 1 Rev. 4** 4.2.5, 5.1.1, 5.2.1, 5.3.4, 5.3.5, 5.3.6, 6.2.2.3 **NIST SP 800-57** Part 2 2.2, 2.4, 3.2, 4.3, 5.3.3, 5.3.4, A.1.2, A.2.1, A.3.2 **NIST SP 800-130** 1, 2.1, 2.2, 2.9, 6.1, 6.2, 6.5 **NIST SP 800-152** 2.2, 4.3, 4.6, 4.7, 6.1.3, 6.4.14, 6.4.29 **CCS CSC** 17 **COBIT 5** APO01.06, BAI02.01, BAI06.01, DSS06.06 **ISA 62443-3-3:2013** SR 3.4, SR 4.1 **ISO/IEC 27001:2013** A.8.2.3
	PR.DS-2: Data-in-transit is protected	**FIPS 140-2** Sec. 4

Category	Subcategory	Informative References
		NIST SP 800-45 Ver. 2 All **NIST SP 800-49** 2 **NIST SP 800-52 Rev. 1** 3, 4, D1.4 **NIST SP 800-53 Rev. 4** SC-8 **NIST SP 800-57 Part 1 Rev. 4** 4.2.5, 5.1.1, 5.2.1, 5.3.4, 5.3.5, 5.3.6, 6.2.1.3 **NIST SP 800-57** Part 2 2.2, 5.3.3, A.2, A.3.1, A.3.2 **NIST SP 800-81-2** All **NIST SP 800-130** 1, 2.1, 2.2, 2.9, 6.1, 6.2, 6.4, 6.7.2 **NIST SP 800-152** 6.1.2, 6.2.1 **NIST SP 800-177** All **CCS CSC** 17 **COBIT 5** APO01.06, DSS06.06 **ISA 62443-3-3:2013** SR 3.1, SR 3.8, SR 4.1, SR 4.2 **ISO/IEC 27001:2013** A.8.2.3, A.13.1.1, A.13.2.1, A.13.2.3, A.14.1.2, A.14.1.3
	PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and information integrity	**FIPS 140-2** Sec. 4 **NIST SP 800-45 Ver. 2** 2.4.2, 3, 4.2.3, 4.3, 5.1, 6.1, 7.2.2, 8.2, 9.2 **NIST SP 800-49** 2.2.1, 2.3.2, 3.4 **NIST SP 800-52 Rev. 1** 3, 4, D1.4 **NIST SP 800-53 Rev. 4** SI-7 **NIST SP 800-57 Part 1 Rev. 4** 5.5, 6.1, 8.1.5.1, B.3.2, B.5

Category	Subcategory	Informative References
		NIST SP 800-57 Part 2 1, 3.1.2.1.2, 4.1, 4.2, 4.3, A.2.2, A.3.2, C.2.2 **NIST SP 800-81-2** All **NIST SP 800-130** 2.2, 4.3, 6.2.1, 63, 6.4, 6.5, 6.6.1 **NIST SP 800-152 6.**1.3, 6.2.1, 8.2.1, 8.2.4, 9.4 **NIST SP 800-177** 2.2, 4.1, 4.4, 4,5, 4,7, 5.2, 5.3 **ISA 62443-3-3:2013** SR 3.1, SR 3.3, SR 3.4, SR 3.8 **ISO/IEC 27001:2013** A.12.2.1, A.12.5.1, A.14.1.2, A.14.1.3
Protective Technology (PR.PT): Technical security solutions are managed to ensure the security and resilience of systems and assets, consistent with related policies, procedures, and agreements.	**PR.PT-4**: Communications and control networks are protected	**OMB M-08-23** FIPS 140-2 Sec. 4 **NIST SP 800-49** 2.4.3, 2.4.4 **NIST SP 800-52 Rev. 1** 3, 4 **NIST SP 800-53 Rev. 4** AC-4, AC-17, AC-18, CP-8, SC-7 NIST SP 800-57 Part 1 Rev. 4 5.3.1, 6.2.2 **NIST SP 800-130** 8.3 **NIST SP 800-15**2 4.7, 4.11.1, 6.8.6, 8.3 **CCS CSC** 7 **COBIT 5** DSS05.02, APO13.01 **ISA 62443-3-3:2013** SR 3.1, SR 3.5, SR 3.8, SR 4.1, SR 4.3, SR 5.1, SR 5.2, SR 5.3, SR 7.1, SR 7.6 **ISO/IEC 27001:2013** A.13.1.1, A.13.2.1

Table C.2 DETECT (DE)

Category	Subcategory	Informative References
Security Continuous Monitoring (DE.CM): The information system and assets are monitored at discrete intervals to identify cybersecurity events and verify the effectiveness of protective measures.	**DE.CM-1**: The network is monitored to detect potential cybersecurity events	**FIPS 140-2** Sec. 4 **SP 800-37 Rev. 1** 3.6 **NIST SP 800-45 Ver. 2** 4.1, 5.1.1, 5.1.5, 6.2.1, 6.2.2, 7.2.2 **NIST SP 800-53 Rev. 4** AC-2, AU-12, CA-7, CM-3, SC-5, SC-7, SI-4 **NIST SP 800-81-2** 2, 9, 12, 13 **NIST SP 800-130** 5, 6.8.5, 8.2.4, 9.8.4 **NIST SP 800-152** 6.8.5, 8.2.3, 8.2.4, 8.3, 8.5 **NIST SP 800-177 3.1**.1 **CCS CSC** 14, 16 **COBIT 5** DSS05.07 **ISA 62443-3-3:2013** SR 6.2
	DE.CM-6: External service provider activity is monitored to detect potential cybersecurity events	**NIST SP 800-53 Rev. 4** CA-7, PS-7, SA-4, SA-9, SI-4 **NIST SP 800-81-2** 2, 9, 12, 13 **NIST SP 800-130** 6.8.5, 8.2.4, 9.8.4, 12 **NIST SP 800-152** 6.8.5, 8.2.3, 8.2.4, 8.3, 8.5 **ISO/IEC 27001:2013** A.14.2.7, A.15.2.1
Detection Process (DE.DP): Detection processes and procedures are maintained and tested to ensure timely and	**DE.DP-4**: Event detection information is communicated to appropriate parties	**NIST SP 800-45** Ver. 2 9.3 **NIST SP 800-53 Rev. 4** AU-6, CA-2, CA-7, RA-5, SI-4 **NIST SP 800-177** 4.6 **COBIT 5** APO12.06

Category	Subcategory	Informative References
adequate awareness of anomalous events.		**ISA 62443-2-1:2009** 4.3.4.5.9 **ISA 62443-3-3:2013** SR 6.1 **ISO/IEC 27001:2013** A.16.1.2

Table C.3 RESPOND (RS)

Category	Subcategory	Informative References
Response Planning (RS.RP): Response processes and procedures are executed and maintained, to ensure timely response to detected cybersecurity events.	**RS.RP-1**: Response plan is executed during or after an event	**NIST SP 800-45 Ver. 2** 9.3 **NIST SP 800-53 Rev. 4** CP-2, CP-10, IR-4, IR-8 **NIST SP 800-57 Part 1 Rev. 4** **NIST SP 800-57 Part 2** 3.1.2.1.3, 3.2.2.6 **NIST SP 800-130** 6.2.1, 6.4.5, 6.4.6, 6.8, 10.1 **NIST SP 800-152** 6.8, 10 **NIST SP 800-177** 4.6 **COBIT 5** BAI01.10 **CCS CSC** 18 **ISA 62443-2-1:2009** 4.3.4.5.1 **ISO/IEC 27001:2013** A.16.1.5
Communications (RS.CO): Response activities are coordinated with internal and external stakeholders, as appropriate, to include external	**RS.CO-2**: Events are reported consistent with established criteria	**NIST SP 800-45 Ver. 2** 9.3 **NIST SP 800-53 Rev. 4** AU-6, IR-6, IR-8 **NIST SP 800-57 Part 1** Rev. 4 8.3.5, 9.3.4, 10.2.9

Category	Subcategory	Informative References
support from law enforcement agencies.		**NIST SP 800-57 Part 2** 3.1.2.1.2, 3.2.2.10, 3.2.2.14, 3.2.2.15, A.1.1, A.1.4, C.2.2.12 **NIST SP 800-130** 6.8 **NIST SP 800-152** 6.8 **NIST SP 800-177** 4.6 **ISA 62443-2-1:2009** 4.3.4.5.5 **ISO/IEC 27001:2013** A.6.1.3, A.16.1.2
Mitigation (RS.MI): Activities are performed to prevent expansion of an event, mitigate its effects, and eradicate the incident.	**RS.MI-1**: Incidents are contained	**NIST SP 800-53 Rev. 4** IR-4 **NIST SP 800-130** 6.8.1 **NIST SP 800-152** 6.8 **ISA 62443-2-1:2009** 4.3.4.5.6 **ISA 62443-3-3:2013** SR 5.1, SR 5.2, SR 5.4 **ISO/IEC 27001:2013** A.16.1.5
	RS.MI-2: Incidents are mitigated	**NIST SP 800-53 Rev. 4** IR-4 **NIST SP 800-57 Part 1 Rev. 4** 5.3, 5.4, 5.5, 8.3.4, 8.3.5 **NIST SP 800-57 Part 2** 5.3.7, 5.3.8 **NIST SP 800-130** 4.9.3, 6.8, 9.5, 12 **NIST SP 800-152** 3.4.2, 4.5, 6.8, 9.5, 9.8, 12 **ISA 62443-2-1:2009** 4.3.4.5.6, 4.3.4.5.10 **ISO/IEC 27001:2013** A.12.2.1, A.16.1.5