

Securing Web Transactions

TLS Server Certificate Management

Volume D:
How-To Guides

Murugiah Souppaya
NIST

Paul Turner
Venafi

Mehwish Akram
Brandon Everhart
Brian Johnson
Brett Pleasant
Susan Symington
The MITRE Corporation

Clint Wilson
DigiCert

Dung Lam
F5

Alexandros Kapsouris
Symantec

William C. Barker
Strativia

Rob Clatterbuck
Jane Gilbert
Thales Trusted Cyber Technologies

June 2020

This publication is available free of charge from:
<http://doi.org/10.6028/NIST.SP.1800-16>

The first draft of this publication is available free of charge from:
<https://www.nccoe.nist.gov/projects/building-blocks/tls-server-certificate-management>

DISCLAIMER

Certain commercial entities, equipment, products, or materials may be identified by name or company logo or other insignia in order to acknowledge their participation in this collaboration or to describe an experimental procedure or concept adequately. Such identification is not intended to imply special status or relationship with NIST or recommendation or endorsement by NIST or NCCoE; neither is it intended to imply that the entities, equipment, products, or materials are necessarily the best available for the purpose.

National Institute of Standards and Technology Special Publication 1800-16D Natl. Inst. Stand. Technol. Spec. Publ. 1800-16D, 223 pages, (June 2020), CODEN: NSPUE2

FEEDBACK

As a private-public partnership, we are always seeking feedback on our practice guides. We are particularly interested in seeing how businesses apply NCCoE reference designs in the real world. If you have implemented the reference design, or have questions about applying it in your environment, please email us at tls-cert-mgmt-nccoe@nist.gov.

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

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 information technology 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 Cybersecurity Framework and details the steps needed for another entity to re-create the example solution. The NCCoE was established in 2012 by NIST in partnership with the State of Maryland and Montgomery County, Maryland.

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

NIST CYBERSECURITY PRACTICE GUIDES

NIST Cybersecurity Practice Guides (Special Publication 1800 series) 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

Transport Layer Security (TLS) server certificates [4][5] are critical to the security of both internet-facing and private web services. A large- or medium-scale enterprise may have thousands or even tens of thousands of such certificates, each identifying a specific server in its environment. Despite the critical importance of these certificates, many organizations lack a formal TLS certificate management program, and the ability to centrally monitor and manage their certificates. Instead, certificate management tends to be spread across each of the different groups responsible for the various servers and systems in an organization. Central security teams struggle to ensure certificates are being properly managed by each of these disparate groups. Where there is no central certificate management service, the organization is

at risk, because once certificates are deployed, current inventories must be maintained to support regular monitoring and certificate maintenance. Organizations that do not properly manage their certificates face significant risks to their core operations, including:

- application outages caused by expired TLS server certificates
- hidden intrusion, exfiltration, disclosure of sensitive data, or other attacks resulting from encrypted threats or server impersonation
- disaster-recovery risk that requires rapid replacement of large numbers of certificates and private keys in response to either certificate authority compromise or discovery of vulnerabilities in cryptographic algorithms or libraries

Despite the mission-critical nature of TLS server certificates, many organizations have not defined the clear policies, processes, roles, and responsibilities needed for effective certificate management. Moreover, many organizations do not leverage available automation tools to support effective management of the ever-growing numbers of certificates. The consequence is continuing susceptibility to security incidents.

This NIST Cybersecurity Practice Guide shows large and medium enterprises how to employ a formal TLS certificate management program to address certificate-based risks and challenges. It describes the TLS certificate management challenges faced by organizations; provides recommended best practices for large-scale TLS server certificate management; describes an automated proof-of-concept implementation that demonstrates how to prevent, detect, and recover from certificate-related incidents; and provides a mapping of the demonstrated capabilities to the recommended best practices and to NIST security guidelines and frameworks.

The solutions and architectures presented in this practice guide are built upon standards-based, commercially available, and open-source products. These solutions can be used by any organization managing TLS server certificates. Interoperable solutions are provided that are available from different types of sources (e.g., both commercial and open-source products).

KEYWORDS

Authentication; certificate; cryptography; identity; key; key management; PKI; private key; public key; public key infrastructure; server; signature; TLS; Transport Layer Security

DOCUMENT CONVENTIONS

The terms “shall” and “shall not” indicate requirements to be followed strictly in order to conform to the publication and from which no deviation is permitted.

The terms “should” and “should not” indicate that among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others, or that a certain course of action is

preferred but not necessarily required, or that (in the negative form) a certain possibility or course of action is discouraged but not prohibited.

The terms “may” and “need not” indicate a course of action permissible within the limits of the publication.

The terms “can” and “cannot” indicate a possibility and capability, whether material, physical, or causal.

CALL FOR PATENT CLAIMS

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory [ITL] publication). Such guidance and/or requirements may be directly stated in this ITL Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this ITL publication and of any relevant unexpired U.S. or foreign patents.

ITL may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

- a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or
- b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL publication either:
 - i) under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or
 - ii) without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to tls-cert-mgmt-nccoe@nist.gov.

ACKNOWLEDGMENTS

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

Name	Organization
Dean Coclin	DigiCert
Tim Hollebeek	DigiCert
Robert Smith	F5
Nancy Correll	The MITRE Corporation
Mary Raguso	The MITRE Corporation
Aaron Aubrecht	Venafi
Justin Hansen	Venafi

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
DigiCert	External Certificate Authority and CertCentral console
F5	BIG-IP Local Traffic Manager load balancer
Thales TCT	Luna SA 1700 Hardware Security Module
Symantec	SSL Visibility Appliance for TLS interception and inspection

Technology Partner/Collaborator	Build Involvement
Venafi	Trust Protection Platform (TLS certificate manager, log server, and scanning tool)

Contents

1	Introduction	1
1.1	Practice Guide Structure	1
1.2	Build Overview	3
1.2.1	Usage Scenarios	3
1.2.2	Logical Architecture	6
1.3	Build Architecture Summary	8
1.4	Typographic Conventions.....	11
1.5	Supporting Infrastructure.....	11
1.5.1	Lab Backbone	12
1.5.2	Supporting Infrastructure Operating Systems.....	13
1.5.3	Supporting Infrastructure Component Services	17
1.5.4	Database Services	28
1.5.5	TLS Web Services	30
1.5.6	DevOps Services.....	41
2	Product Installation and Configuration Guides.....	42
2.1	Product Installation Sequence (Example Build).....	42
2.2	Thales TCT Luna SA 1700 Hardware Security Module	43
2.2.1	Day 0: Product Installation and Standard Configuration	44
2.2.2	Day 1: Product Integration Configuration.....	55
2.2.3	Day N: Ongoing Security Management and Maintenance	89
2.3	DigiCert Certificate Authority.....	91
2.3.1	Day 0: Installation and Standard Configuration.....	91
2.3.2	Day 1: Integration Configuration	97
2.3.3	Day N: Ongoing Security Management and Maintenance	103
2.4	F5 BIG-IP Local Traffic Manager (LTM).....	109
2.4.1	Day 0: Installation and Standard Configuration.....	110
2.4.2	Day 1: Product Integration Configuration.....	123
2.4.3	Day N: Ongoing Security Management and Maintenance	127

- 2.5 Symantec SSL Visibility Appliance 136
 - 2.5.1 Day-0: Install and Standard Configuration.....136
 - 2.5.2 Day 1: Product Integration Configuration.....146
 - 2.5.3 Day N: Ongoing Security Management and Maintenance154
- 2.6 Venafi Trust Protection Platform (TPP)..... 155
 - 2.6.1 Prerequisites155
 - 2.6.2 Installation155
 - 2.6.3 CA Integration163
 - 2.6.4 Folder Creation164
 - 2.6.5 Custom Fields.....165
 - 2.6.6 Assigning Certificate Owners166
 - 2.6.7 Setting Policies167
 - 2.6.8 Establishing a Domain Allowlist169
 - 2.6.9 Workflow – RA Reviews170
 - 2.6.10 CA Import171
 - 2.6.11 Network Discovery.....173
 - 2.6.12 Identify Certificate Risks/Vulnerabilities.....173
 - 2.6.13 Automate Management174
 - 2.6.14 Continuous Monitoring.....192
- Appendix A Passive Inspection 197**
- Appendix B Hardening Guidance..... 200**
- Appendix C Venafi Underlying Concepts 202**
 - C.1 Venafi TPP Object Model 204
 - C.2 Certificate Metadata in Venafi TPP 205
 - C.3 Custom Fields 207
 - C.3.1 Organizing Certificate Inventory.....207
 - C.3.2 Policy Enforcement208
 - C.4 The Domain Allowlist..... 208
 - C.4.1 Certificate Owner Assignment208
 - C.4.2 Permissions208

C.4.3	Contacts	209
Appendix D	List of Acronyms	210
Appendix E	Glossary	214
Appendix F	References	222
Appendix G	Supplemental Architecture Configurations.....	223
G.1	Mail Server Configuration Files	223

List of Figures

Figure 1-1	TLS Server Certificate Management Example Implementation: Logical Architecture.....	7
Figure 1-2	TLS Server Certificate Management Example Implementation: Laboratory Configuration	9
Figure 1-3	TLS Lab Logging Infrastructure	39
Figure 2-1	Overview of Dependencies Among Components Deployed for the Example Build	43
Figure 2-2	Venafi Dashboard Expiration Widget showing the Certificate Expiration Profile.....	174

List of Tables

Table 1-1	Naming and Addressing Information for all Microsoft Windows Servers	14
Table 1-2	Naming and Addressing Information for all Microsoft Windows 10 Workstations	15
Table 1-3	Naming and Addressing Information for All Fedora-Based Systems	16
Table 1-4	Naming and Addressing Information for All CentOS Servers	17

1 Introduction

Organizations that improperly manage their Transport Layer Security (TLS) server certificates [4][5] risk system outages and security breaches, which can result in revenue loss, harm to reputation, and exposure of confidential data to attackers. TLS is the most widely used protocol for securing web transactions and other communications on internal networks and the internet. TLS certificates are central to the operation and security of internet-facing and private web services. Some organizations have tens of thousands of TLS certificates and keys requiring ongoing maintenance and management.

The National Cybersecurity Center of Excellence (NCCoE) at the National Institute of Standards and Technology (NIST) built a laboratory environment to demonstrate how large and medium enterprises can better manage TLS server certificates in the following ways:

- defining operational and security policies and identifying roles and responsibilities
- establishing comprehensive certificate inventories and ownership tracking
- conducting continuous monitoring of the certificate operation and security status
- automating certificate management to minimize human error and maximize efficiency on a large scale
- enabling rapid migration to new certificates and keys as needed in response to certificate authority (CA) compromise or discovery of vulnerabilities in cryptographic algorithms or libraries

The following volumes of this guide show information technology (IT) professionals and security engineers how we implemented this example solution. We cover all the products employed in this reference design. We do not re-create the product manufacturers' documentation, which is presumed to be widely available. Rather, these volumes show how we incorporated the products together in our environment.

Note: These are not comprehensive tutorials. There are many possible service and security configurations for these products that are out of scope for this reference design.

1.1 Practice Guide Structure

This National Institute of Standards and Technology (NIST) Cybersecurity Practice Guide demonstrates a standards-based reference design and provides users with the information they need to replicate automated management of TLS server certificates. This reference design is modular and can be deployed in whole or in part.

This guide contains four volumes:

- NIST SP 1800-16A: *Executive Summary*
- NIST SP 1800-16B: *Security Risks and Recommended Best Practices*
- NIST SP 1800-16C: *Approach, Architecture, and Security Characteristics*—what we built and why

- NIST SP 1800-16D: *How-To Guides*—instructions for building the example solution (**you are here**)

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 1800-16A, which describes the following topics:

- recommendations for TLS server certificate management
- challenges that enterprises face in proper deployment, management, and use of TLS
- example solution built at the NCCoE

You might share the *Executive Summary*, NIST SP 1800-16A, with your leadership team members to help them understand the importance of adopting standards-based TLS server certificate management.

Senior information technology and security officers will be informed by NIST SP 1800-16B, which describes the:

- TLS server certificate infrastructure and management processes
- risks associated with mismanagement of certificates
- organizational challenges associated with server certificate management
- recommended best practices for server certificate management
- recommendations for implementing a successful certificate management program
- mapping of best practices for TLS server certificate management to the NIST Framework for Improving Critical Infrastructure Cybersecurity (Cybersecurity Framework) [8]
- application of specific controls defined within NIST Special Publication (SP) 800-53 [4] to the TLS server certificate management recommended best practices

Technology or security program managers who are concerned with how to identify, understand, assess, and mitigate risk will be interested in NIST SP 1800-16C, which describes what we did and why. The following sections will be of particular interest:

- Section 3.4.1, Threats, Vulnerabilities and Risks, provides a description of the risk analysis we performed.
- Section 3.4.2, Security Categorization and SP 800-53 Controls [4], lists the security controls assigned to address TLS server certificate risks.
- Section 3.4.3, Security Control Map, maps the security characteristics of this example solution to cybersecurity standards and best practices.

IT professionals who want to implement such an approach will find this whole practice guide useful. You can use this How-To portion of the guide, NIST SP 1800-16D, to replicate all or parts of the build created in our lab. This How-To portion of the guide provides specific product installation, configuration, and

integration instructions for implementing the example solution. We do not re-create 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 and open source products to address this challenge, this guide does not endorse these particular 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 providing automation support for TLS server certificate management. Your organization's security experts should identify the products that will best integrate with your existing tools and IT system infrastructure. We hope that you will seek products that are congruent with applicable standards and best practices. [Section 1.3](#), Build Architecture Summary, lists the products that 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. We seek feedback on its contents and welcome your input. Comments, suggestions, and success stories will improve subsequent versions of this guide. Please contribute your thoughts to tls-cert-mgmt-nccoe@nist.gov.

1.2 Build Overview

This NIST Cybersecurity Practice Guide addresses the use of commercially available technologies to develop an example implementation for managing TLS server certificates. This project focuses on certificate management in medium and large enterprises that rely on TLS to secure customer-facing and internal applications. The example implementation developed in this project demonstrates how to manage TLS server certificates to reduce outages, improve security, and enable disaster recovery activities. It shows how to establish, assign, change, and track an inventory of TLS certificates; automate management of TLS certificates; perform continuous monitoring of TLS certificates; perform large-scale replacement of certificates that are not trusted; log all certificate and private-key management operations; manage certificates and keys on proxy servers, load balancers, and inspection appliances; and use a Hardware Security Module (HSM). The HSM can securely generate, store, manage, and use private keys corresponding to TLS server certificates, the signing keys of internal certificate authorities (CAs), and symmetric keys that must be kept secret.

1.2.1 Usage Scenarios

The example implementation fulfills the following use cases:

- building and maintaining inventory of the enterprise's deployed TLS server certificates

- automating management of those certificates, including use of an external CA and protection of private keys and other secrets by using an HSM
- continuously monitoring the certificates for validity
- supporting disaster recovery by quickly replacing a large number of certificates
- logging all certificate and private-key management operations
- for those enterprises with a policy to perform passive inspection, copying private keys from several different TLS servers to the TLS inspection appliance

1.2.1.1 *Building the Inventory*

The example implementation demonstrates the ability to establish and maintain a systematized inventory of certificates (and keys) in use on the network. It enables a user to discover certificates not currently being managed by the inventory, efficiently enroll and provision new certificates (and keys), store relevant information with those certificates, and discover the absence of an expected certificate from a machine where it should be installed. It also enables certificates to be revoked and to change the owner associated with a certificate, as needed.

1.2.1.2 *Automation*

The example implementation demonstrates the ability to automatically enroll and provision a new certificate and can replace a certificate approaching expiration. Automated certificate management is demonstrated on various enterprise systems, including load balancers acting as TLS proxies that use remote agentless management, web servers with remote agentless management, web servers using the Automatic Certificate Management Environment (ACME) protocol, and servers that are deployed via development operations (DevOps) technologies by using a certificate management plug-in to the DevOps framework. In conjunction with the demonstration of ACME, HSM is used to securely generate, store, manage, and process the cryptographic key pairs for one TLS server. Remote agentless management was used to automate management of the certificates and keys for this system. In the current effort the NCCoE undertook only a limited demonstration. This limited demonstration employed Kubernetes in a cloud environment where DevOps frameworks are commonly used.

1.2.1.3 *Continuous Monitoring*

The example implementation demonstrates the ability to continuously monitor TLS certificates (and keys) managed by the inventory system and can act upon the status of any certificate (e.g., report the status of or replace a certificate that has expired, is about to expire, or does not conform to policy). It can send periodic expiration reports to certificate owners to show which of their certificates are nearing expiration, and a variety of notifications and escalating alerts if a certificate's expiration date approaches. Continuous monitoring also includes periodic network scans to ensure any unaccounted-for certificates are discovered and added to the inventory.

1.2.1.4 *Disaster Recovery*

The example implementation demonstrates how to quickly replace large numbers of certificates that are located across multiple networks and that are on a variety of server types, because the certificates are no longer trusted. It can replace certificates that:

- were issued by a given CA (which would require replacement if the issuing-CA were either compromised or untrusted)
- have associated keys dependent on a specific cryptographic algorithm (which would need replacement, e.g., if the algorithm they depend on is no longer considered secure)
- have associated keys generated by a specific cryptographic library after a specific date (which would need replacement, e.g., if a bug invaded a library on that date)

The example implementation can also track and report on replacement of large numbers of certificates, so the progress of the large-scale certificate replacement effort can be monitored.

1.2.1.5 *Logging*

The example implementation demonstrates how to log all certificate and private-key management operations, including certificate creation, installation and revocation key pair generation, certificate requests and request approvals, certificate and key copying, and certificate and key replacement.

1.2.1.6 *Passive Inspection*

The example implementation demonstrates how to perform passive inspection of encrypted TLS connections. The decision to perform this inspection is complex, because it involves important trade-offs between traffic security and traffic visibility that each organization should weigh for itself. Some organizations have determined that the security risks posed by inspection of internal TLS traffic are not worth the potential benefits of visibility into the encrypted traffic. Other organizations have concluded that the visibility into their internal traffic provided by TLS inspection is worth the trade-off of the weaker encryption and other risks that come with such inspection. For these organizations, TLS inspection may be considered standard practice and may represent a critical component of their threat detection and service assurance strategies.

Organizations that perform TLS traffic inspections can use the example implementation to securely copy private keys from several different TLS servers to the TLS inspection appliance, securely replace expiring keys on servers, and immediately copy those keys to the inspection appliance before expiration—manually and via standardized automated certificate installation. See Appendix A for more detail on passive inspection, including a scenario.

1.2.2 Logical Architecture

Figure 1-1 depicts the example implementation's logical architecture, which provides a network structure and components that enable various types of TLS server certificate management operations to function. Figure 1-1 illustrates the logical architecture of the TLS server certificate management example implementation—consisting of an external and an internal portion. The external portion contains an external CA that is used to issue TLS certificates for some TLS servers in the example implementation. The internal portion of the network is logically organized into three zones that roughly model a defense-in-depth strategy of grouping components on subnetworks that require increasing levels of security as one moves inward from the perimeter of the organization. The zones comprise a demilitarized zone (DMZ) that sits between the internet and the rest of the enterprise; a data center hosting applications and services widely used across the enterprise; and a more secure data center hosting critical security and infrastructure components, including certificate management components.

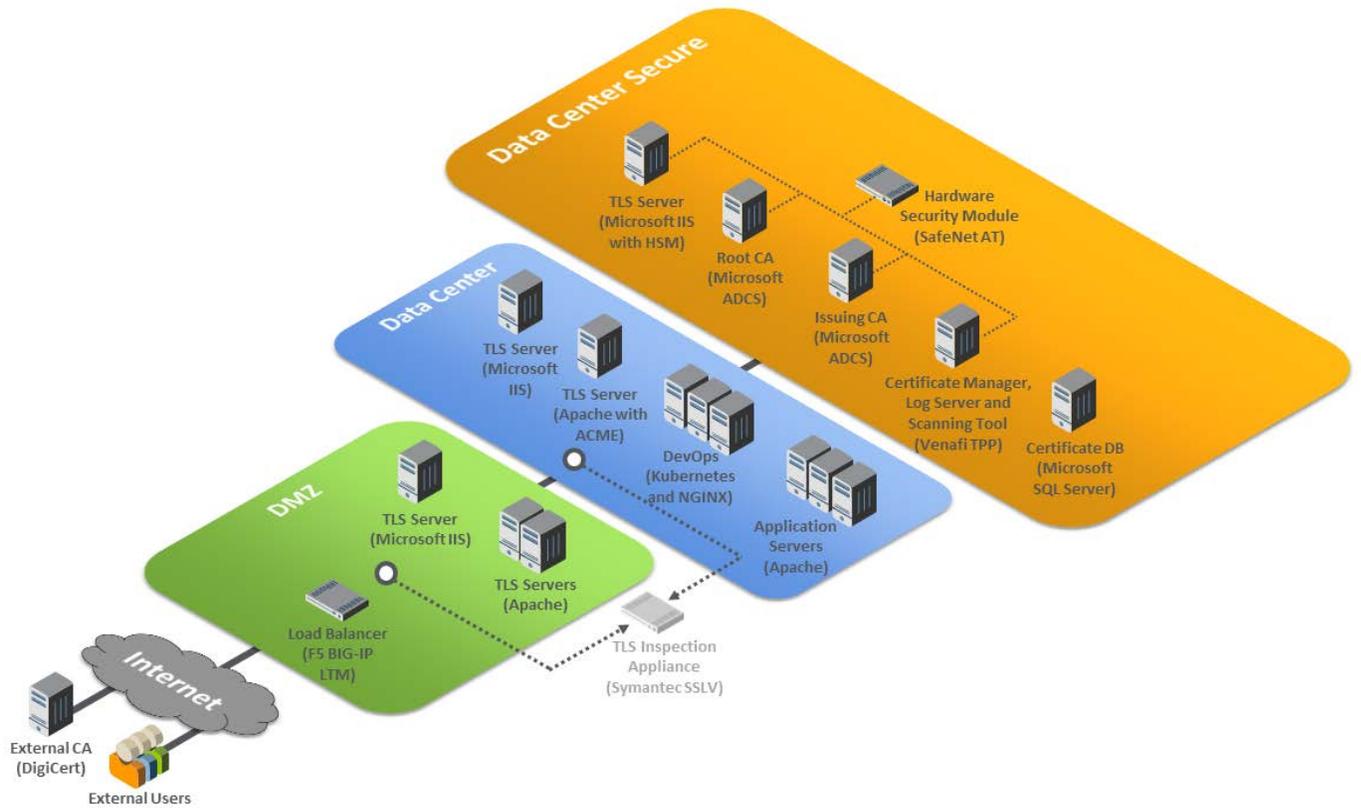
At the ingress from the internet within the DMZ, a load balancer acts as a TLS proxy and distributes the traffic it receives from external users across three TLS servers behind it—all serving up the same application: two Apache servers and one Microsoft Internet Information Services (IIS) server. (Note: To maintain the diagram's simplicity in depicting this network, the connections between individual components are not shown. In the actual network architecture, the load balancer's network connection to all three TLS servers is shown behind it.) TLS certificate management demonstrates how to enroll and provision new certificates to the load balancer and servers in the DMZ and how to perform overall certificate management on these devices, including automatically replacing a certificate that is nearing expiration.

Within the data center zone of the logical architecture sit various types of web servers, application servers, and a DevOps framework—all act as TLS servers. These components demonstrate the ability to automatically enroll and provision a new certificate and can automatically replace a certificate that is nearing expiration on these different systems. Various types of certificate management are also demonstrated, including remote agentless management, the ACME protocol, and the DevOps certificate management plug-in.

Within the DMZ and the data center zones, taps (depicted as white dots) are used on the network connections between the load balancer and the servers behind it, and on the network connections between the DMZ servers and the second-tier servers in the data center behind them. Taps enable all traffic on the encrypted TLS connections to travel to a TLS inspection appliance for passive decryption. Figure 1-1 depicts this TLS inspection appliance as a faded icon to convey that some organizations, as a matter of policy, may not want to include it as part of their network architecture. However, organizations that consider passive inspection as part of their security assurance strategy can use the certificate manager depicted in the architecture to securely copy private keys from several different TLS servers to the TLS inspection appliance, and to securely replace expiring keys on those servers and

immediately copy those keys to the decryption device before expiration—manually and via standardized automated certificate installation.

Figure 1-1 TLS Server Certificate Management Example Implementation: Logical Architecture



Within the data center secure zone of the logical architecture sit the components that perform TLS server certificate management. These components include internal root and issuing CAs, a certificate manager, a certificate log server, a certificate network scanning tool, a certificate database, and an HSM. For demonstration purposes, a TLS server connected to an HSM is also present in this zone.

The certificate manager can be used in conjunction with the certificate database and the various types of servers in the architecture to demonstrate how to establish and maintain a systematized inventory of certificates (and keys) used on the network. The certificate manager can also continuously monitor TLS certificates (and keys) managed by the inventory system and act upon the status of any certificate (e.g., report a certificate that is expired, about to expire, or does not conform to policy, or it can replace an expired certificate). It can also send expiration reports and notifications to certificate owners and can support disaster recovery by quickly replacing a large number of certificates located throughout the network architecture.

The certificate manager can be used in conjunction with the CAs to enroll and provision certificates (and keys), store attributes with those certificates, and discover the absence of an expected certificate from a machine where it should be installed. The certificate manager can revoke certificates and change the owner associated with that certificate.

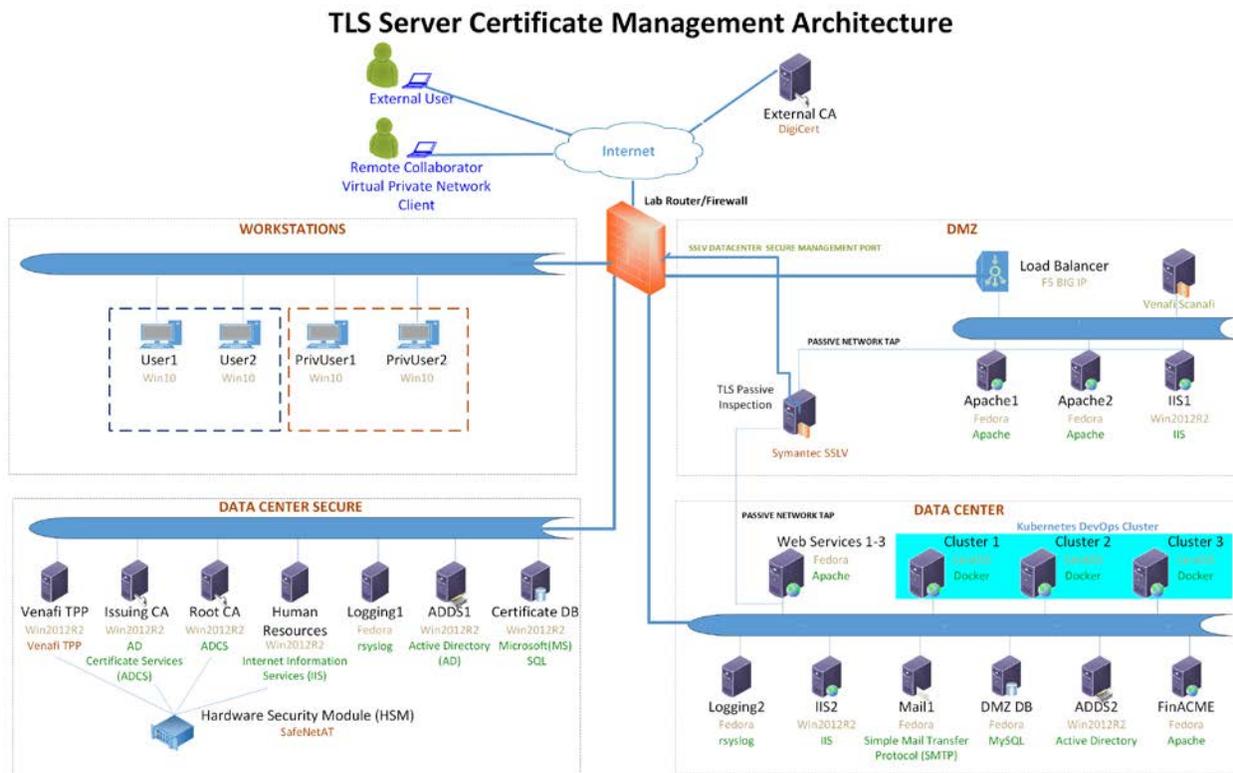
The certificate network scanning tool can discover certificates not being managed by the inventory. The certificate log server can record all certificate and private-key management operations, including certificate creation, installation, and revocation; key pair generation; certificate requests and request approvals; certificate and key copying; and certificate and key replacement.

All components in this portion of the architecture—except for the certificate database—are configured to use the HSM, which can securely generate, store, manage, and process the private key corresponding to the TLS server’s certificate. The HSM is capable of storing and protecting the symmetric keys that secure sensitive data in the certificate database, and can generate, store, manage, and process internal CAs’ signing keys.

1.3 Build Architecture Summary

Figure 1-2 depicts the physical architecture of the example implementation deployed in the NCCoE laboratory.

Figure 1-2 TLS Server Certificate Management Example Implementation: Laboratory Configuration



The NCCoE laboratory environment provided the following supporting infrastructure for the example implementation:

- firewall-protected connection to the internet where an external CA resides
- Windows 2012 server with remote desktop manager, which acts as a jump box to facilitate installation, deployment, and management of server software for collaborative projects
- segmented laboratory network backbone that models the separation typically existent between subnetworks belonging to different parts of a medium-to-large-scale enterprise—for example, a DMZ, a data center hosting widely used applications and services, a more secure data center hosting critical security infrastructure components, and a segment containing user workstations
- virtual machine and network infrastructure
- Windows 2012 server serving as a Microsoft Active Directory (AD) primary domain controller
- the Windows 2012 server running AD Certificate Services, including
 - an internal Root CA that can issue and self-sign its own TLS certificate

- an internal issuing CA that:
 - issues TLS certificates to servers that request them (issue CAs are subordinate to and certified by the root CA)
 - manages the life cycle of certificates (including request, issuance, enrollment, publication, maintenance, revocation, and expiration)
- Microsoft structured query language (SQL) Server hosting the database of TLS certificates and keys, and corresponding configuration data
- DevOps automation framework, including Kubernetes, Docker, and Jetstack, that demonstrates automated certificate management when performing open-source container orchestration
- Apache, Microsoft IIS, and NGINX servers, which demonstrate various ways of managing TLS server certificates, including remote agentless certificate management, management via the ACME protocol (via the Certbot utility), and management via DevOps
- Apache servers used to demonstrate certificate management on second-tier internal application servers

The following collaborator-supplied components were integrated into the above supporting infrastructure to yield the TLS server certificate management example implementation:

- Venafi Trust Protection Platform (TPP), which maintains the certificate inventory, performs automated TLS server certificate and private-key management, including monitoring, remediation, and rapid replacement of TLS certificates and keys; TLS certificate and key policy enforcement; automated certificate requests and renewals; automated network scanning for TLS certificates; and logging of certificate and private-key management operations
- Symantec SSL Visibility (SSLV), a visibility appliance used to inspect intercepted traffic on encrypted TLS connections
- Thales Trusted Cyber Technologies (Thales TCT) Luna SA 1700 HSM, used to securely generate, store, manage, and process the cryptographic key pair; also uses it to sign TLS certificates within a hardened, tamper-resistant physical appliance. It is also used to store other keys, such as the database encryption key and the TLS certificate keys for the key manager component (Venafi TPP) and the CAs
- DigiCert external CA, which issues and renews TLS certificates
- F5 Networks BIG-IP Local Traffic Manager load balancer, which acts as a TLS proxy and distributes received traffic across a number of other TLS servers

The remainder of this volume describes in detail the installation, configuration, and integration of the above supporting infrastructure and collaborator components.

1.4 Typographic Conventions

The following table presents typographic conventions used in this volume.

Typeface/Symbol	Meaning	Example
<i>Italics</i>	file names and path names; references to documents that are not hyperlinks; new terms; and placeholders	For detailed definitions of terms, see the <i>NCCoE Glossary</i> .
Bold	names of menus, options, command buttons, and fields	Choose File > Edit .
Monospace	command-line input, onscreen computer output, sample code examples, and status codes	<code>mkdir</code>
Monospace Bold	command-line user input contrasted with computer output	<code>service sshd start</code>
blue text	link to other parts of the document, a web URL, or an email address	All publications from NIST's NCCoE are available at https://www.nccoe.nist.gov .

1.5 Supporting Infrastructure

This section is the first in a series of how-to guidance offered in this guide. It contains step-by-step instructions and points to specific, well-known, and trusted information for installing, configuring, and securely maintaining the supporting infrastructure components outlined in previous sections of this document.

All supporting infrastructure components in the following how-to subsections are high-level examples of services and functions that may reside on any network. For example, the Microsoft suite of AD, CA services, domain name server (DNS), web, and database services would typically reside on most organizational networks. Each section follows the other in building the prerequisites. This section on supporting infrastructure is the basis for the subsequent how-to sections on collaborator capabilities.

The lab backbone is the fundamental component of the architecture and forms the basis to develop the implementers' understanding of the simulated build experience. Guidance is provided for each operating system (OS) installation, with specific instructions on the necessary security and system

configurations. Finally, specific ancillary services, installation and security configurations for database services, web services, etc. are provided.

1.5.1 Lab Backbone

The NCCoE has a specific implementation of its supporting lab network infrastructure or lab backbone. Although implementors using this document may possess some or most of the components in the TLS lab backbone, they may encounter slight but significant differences in their lab build. These differences are attributed to how we configured our lab backbone to suit the needs of the TLS lab and the larger multitiered lab community within the NCCoE.

The components and configuration approaches listed below may help clarify what basic capabilities are needed at a minimum to simulate the TLS lab infrastructure backbone.

- network topology—designed to provide strict separation of system and workstation duties:
 - Data Center Secure Network—provides physical and logically secure separation of critical security services from nonprivileged or privileged users without specific security responsibilities
 - Data Center Network—provides less privileged users with access to security maintenance services that do not require special access to critical security management services
 - Workstations Network—provides secure, controlled, and monitored access to nonprivileged authorized users to perform organizational business
 - DMZ—provides secure separation and mitigation of risk to the rest of the critical network services from public access to public-facing services
- multiple virtual local area networks (VLANs) and separate subnets—customized naming convention for VLAN names and subnets can be used, or follow the TLS lab approach below:
 - VLAN 2198 services the Data Center Secure Network 192.168.1.0/24
 - VLAN 2199 services the Data Center Network 192.168.3.0/24
 - VLAN 2200 services the Workstations Network 192.168.2.0/24
 - VLAN 2197 services the DMZ Network 192.168.4.0/24
 - VLAN 2196 services connections between the F5 load balancer and lab firewall 192.168.5.0/24
 - VLAN 2202 services wide area network connections between the internet and the firewall; the address used here should mirror whatever is currently used for what the internet provider gave in a subnet address
- One or more managed layer three switches must be capable of:

- traffic separation for six VLANs with multiple devices on each VLAN (see the architecture diagram for more)
- switched port analyzer (SPAN) or port mirroring functions
- VLAN trunk ports when using multiple switches
- One or more manageable advanced firewalls:
 - must be capable of accepting at least six Ethernet port connections for all VLANs if using one firewall
 - must be capable of network address translation (NAT) (port forwarding, hide NAT, and static NAT)
 - should at least be stateful
 - should support deep packet inspection for every possible subnet where feasible and financially practical

1.5.2 Supporting Infrastructure Operating Systems

1.5.2.1 *Microsoft Windows*

Microsoft Windows and Windows Server are within a group of OSs designed by Microsoft to efficiently manage enterprise needs for data storage, applications, networking, and communications. In addition to the standard OSs used, additional ancillary Microsoft services were installed. These are native components of the OS and critical to the TLS lab design. Guidance on configuration of these ancillary services will be discussed later in this document in the Supporting Infrastructure Component Services section.

- AD Services
- DNS Services
- CA Services

1.5.2.1.1 Microsoft Windows and Server Prerequisites

Both Microsoft Windows servers and workstations have minimal hardware prerequisites, listed directly below this paragraph. In addition, TLS lab host configuration information is provided in Table 1-1 and [Table 1-2](#) below. While it is not imperative that an implementer uses the TLS lab host naming convention and internet protocol (IP) addressing schemes, the tables below may prove useful with informing an organization of the servers and workstations needed should there be customizations to the TLS lab approach.

While the hardware requirements listed below represent the minimum, most business applications of this effort may have higher but differing requirements. All the applications in this TLS build will greatly

benefit from adding more than the minimum resources that Microsoft requires, as shown below, in a production environment.

Microsoft’s Minimum Hardware Requirements:

- Microsoft Windows Servers 2012
 - 1 gigahertz (GHz) 64-bit processor
 - 512 megabyte (MB) random access memory (RAM)
 - 32 gigabytes (GB) disk space
- Microsoft Windows Workstations 2010
 - 1 GHz 64-bit processor
 - 2 GB RAM
 - 20 GB disk space

1.5.2.1.2 Microsoft Windows Server 2012 Installation

- For instructions regarding downloading the Microsoft Windows Server 2012, refer to the download and deployment guidance at: <https://www.microsoft.com/en-us/evalcenter/evaluate-windows-server-2012-r2>.

Given that AD and domain services are critical to the adds1 and adds2 installation process, refer to the **Microsoft Active Directory and Domain Services Installation and Configuration** section, [1.5.3.1](#), of this document for full instructions after initial basic installation of the OS.

Please use the table below to name and assign IP addresses to all Microsoft Windows Servers used in the TLS lab build. The Windows Server version used in most cases is Windows 2012 version R2.

Table 1-1 Naming and Addressing Information for all Microsoft Windows Servers

Host Name	IP Address	Subnet	Gateway	Software Selection
iis1.ext-nccoe.org	192.168.4.4	255.255.255.0	192.168.4.1	Win2012 R2
adds1.int-nccoe.org	192.168.1.6	255.255.255.0	192.168.1.1	Win2012 R2
HSMrootca.int-nccoe.org	192.168.1.10	255.255.255.0	192.168.1.1	Win2012 R2
BaseSubCA.int-nccoe.org	192.168.1.41	255.255.255.0	192.168.1.1	Win2012 R2
HRhsm	192.168.1.16	255.255.255.0	192.168.1.1	Win2012 R2
Venafi1	192.168.1.81	255.255.255.0	192.168.1.1	Win2012 R2
VTPPTrustDB	192.168.1.89	255.255.255.0	192.168.1.1	Win2012 R2
iis2.int-nccoe.org	192.168.3.5	255.255.255.0	192.168.3.1	Win2012 R2

Host Name	IP Address	Subnet	Gateway	Software Selection
adds2.int-nccoe.org	192.168.3.7	255.255.255.0	192.168.3.1	Win2012 R2
dmzdc.ext-nccoe.org	192.168.3.8	255.255.255.0	192.168.3.1	Win2012 R2

1.5.2.1.3 Microsoft Windows 10 Workstations Installation

- For instructions regarding download of the Microsoft Windows 10 workstation used in this TLS lab build, refer to the guidance at <https://www.microsoft.com/en-us/software-download/windows10>.

Please use the table below to name and assign IP addresses to all Microsoft Windows 10 workstations used in the TLS lab build. The Windows 10 version used in most cases is Windows 10 Pro.

Table 1-2 Naming and Addressing Information for all Microsoft Windows 10 Workstations

Host Name	IP Address	Subnet	Gateway	Software Selection
win10-1.int-nccoe.org	192.168.2.11	255.255.255.0	192.168.2.1	Win10_Pro
win10-2.int-nccoe.org	192.168.2.2	255.255.255.0	192.168.2.1	Win10_Pro
privuser1.int-nccoe.org	192.168.2.3	255.255.255.0	192.168.2.1	Win10_Pro
privuser2.int-nccoe.org	192.168.2.4	255.255.255.0	192.168.2.1	Win10_Pro

1.5.2.2 Linux

Linux is a family of free and open-source OSs based on the Linux kernel, an OS kernel first released on September 17, 1991, by Linus Torvalds. Fedora Server is a Red Hat Corporation-supported, short life-cycle, and fully community-supported server OS. Fedora enables system administrators of any skill to freely (in most cases) make use of the very latest technologies available in the open-source community.

The CentOS Linux distribution is no different in its ability to allow mostly free use of world-class security and general IT capabilities. CentOS is a manageable and reproducible platform derived from the sources of Red Hat Enterprise Linux (RHEL) by an open-source community of volunteers.

1.5.2.2.1 Linux Prerequisites

[Table 1-3](#) and [Table 1-4](#) include the host names and IPs used in the TLS lab for all Linux machines. The recommended minimum hardware requirements for the default installations of Fedora and CentOS have been noted below. An organization's requirements may differ. However, it is highly recommended that the maximum optimal configuration (in accordance with the organization's available resources) for each system be applied, as all the applications used in this TLS lab build will benefit from more than the minimum resources in a production environment.

- 1 GHz or faster processor
- 1 GB system memory
- 10 GB unallocated drive space
- 1 VMXNET 3 network adapter

1.5.2.2.2 Fedora and CentOS Installation

The OS installation process for the TLS lab Linux machines did not deviate from the standard installation instructions that exist for each Linux distributor. The links below provide standard guidance for the Fedora and CentOS installations.

When running through the installation process, in some cases, a standard Fedora installation for software selection will not suffice. Should this occur, use Table 1-3. If the Software Selection column includes Fedora Server/Basic Web Server, select Fedora Server for Base Environment, then select Basic Web Server installation for add-ons, and when prompted, select software packages during the installation.

The CentOS Software Selection column includes Basic Web Server—select this as the software package to install when prompted during the installation process for CentOS.

- <https://docs.fedoraproject.org/en-US/fedora/f28/install-guide/>
- <https://docs.centos.org/en-US/centos/install-guide/>

Please use Table 1-3 for IP, host name, and other installation-specific options for all Fedora-based systems in the TLS lab build.

Table 1-3 Naming and Addressing Information for All Fedora-Based Systems

Host Name	IP Address	Subnet	Gateway	Software Selection
syslog2.int-nccoe.org	192.168.3.12	255.255.255.0	192.168.3.1	Fedora Server
finacme.int-nccoe.org	192.168.3.61	255.255.255.0	192.168.3.1	Fedora Server/ Basic Web Server
mail1.int-nccoe.org	192.168.3.25	255.255.255.0	192.168.3.1	Fedora Server
dmzdb.ext-nccoe.org	192.168.3.6	255.255.255.0	192.168.3.1	Fedora Server
syslog1.int-nccoe.org	192.168.1.12	255.255.255.0	192.168.1.1	Fedora Server
apache1.ext-nccoe.org	192.168.4.2	255.255.255.0	192.168.4.1	Fedora Server/ Basic Web Server
apache2.ext-nccoe.org	192.168.4.3	255.255.255.0	192.168.4.1	Fedora Server/ Basic Web Server

Host Name	IP Address	Subnet	Gateway	Software Selection
ws1.int-nccoe.org	192.168.3.87	255.255.255.0	192.168.3.1	Fedora Server/ Basic Web Server
ws2.int-nccoe.org	192.168.3.88	255.255.255.0	192.168.3.1	Fedora Server/ Basic Web Server
ws3.int-nccoe.org	192.168.3.89	255.255.255.0	192.168.3.1	Fedora Server/ Basic Web Server

Please use Table 1-4 for IP, host name, and other installation-specific options for all CentOS servers used in the TLS lab build.

Table 1-4 Naming and Addressing Information for All CentOS Servers

Host Name	IP Address	Netmask	Gateway	Software Selection
scanafi.ext-nccoe.org	192.168.4.107	255.255.255.0	192.168.4.1	Infrastructure Server
cluster1.int-nccoe.org	192.168.3.103	255.255.255.0	192.168.3.1	Basic Web Server
cluster2.int-nccoe.org	192.168.3.104	255.255.255.0	192.168.3.1	Basic Web Server
cluster3.int-nccoe.org	192.168.3.105	255.255.255.0	192.168.3.1	Basic Web Server

1.5.3 Supporting Infrastructure Component Services

1.5.3.1 *Microsoft Active Directory and Domain Services Installation and Configuration*

Active Directory Services (ADS) and DNS work together to store directory data and make those resources available to administrators and users. For example, ADS stores information about user accounts such as names and passwords. Security is integrated with ADS through log-on authentication and enforced access control for user, file, directory, and other system objects in the directory of services.

Administrators are able to manage directory data and organization roles across the enterprise. They can assign permissions to users, which allows users to access resources anywhere on the network. ADS authenticates and authorizes all users and computers in a Windows domain network. ADS works in conjunction with Group Policies Objects (GPOs) in assigning and enforcing security policies for all computers.

A DNS is a protocol for how computers translate domain names. It manages a database used to resolve domain names to IP addresses, allowing computers to identify each other on the network. DNS is the primary locator service for AD. ADS is highly dependent on the DNS in most cases, and as a result, most implementations—including the TLS lab—opt to install the DNS service on the same server as the ADS.

1.5.3.1.1 ADS and DNS Prerequisites

Below are the minimum recommended tools, services, and configurations needed to install ADS and DNS.

- The adds1 and adds2 hosts should be built with the Windows Server 2012 OS installed. As described in Section [1.5.2.1.2](#) of this document, there are two ADS and DNS servers. The TLS lab ADS and DNS server names used are adds1.int-nccoe.org and adds2.int-nccoe.org. (Note: The DNS server may be run locally on the same Active Directory Domain Services [ADDS] server.)
- local network configurations—all of the local network VLANs, IP addresses, and proper routes
- familiarity with Server Manager

Server Manager is a Windows Server management console that allows administrators to install, configure, and manage server roles and features. Administrators can manage local and remote servers without having physical access to them. The ADS and DNS installation process is integrated with Server Manager, which can be used when installing other server roles.

1.5.3.2 *ADS and DNS Installation*

For instructions on deploying ADS and DNS on a Windows 2012 server, refer to the guidance at one of the links below:

- **Graphical User Interface (GUI)-Based Installation:** <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/deploy/ad-ds-installation-and-removal-wizard-page-descriptions>
- **Command Line-Based Installation:** <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/deploy/install-active-directory-domain-services--level-100->

1.5.3.3 *Certificate Authority Services*

In an organization where public key infrastructure (PKI) has been implemented, a CA is responsible for validating the identity of users and computers. The CA assigns a trusted credential for use in authenticating user and system identities, by issuing a digitally signed and trusted certificate. The CA can also assist in managing revocation and renewal of its signed certificates.

The first CA built and implemented in a PKI environment is often referred to as the root CA. As the originator and root of trust, the root CA authorizes all subsequent CAs, called subordinates or issuing CAs. Subordinate CAs can also designate their own subsidiaries as defined by the root CA, which results in a certificate hierarchy. The metadata supplied in all certificates issued to CAs lower in the hierarchy from the root CA contain a trace path back to the root.

A compromised root CA will cripple any organization that depends on the integrity of its issued PKI certificates, even in lightweight transactions. With full control or significant unauthorized access to the root CA, a malicious actor may fully infiltrate any transaction that relies on the integrity of the trust chain where that root CA presides as the anchor. It is recommended all organizations—size notwithstanding—implement an enterprise stand-alone offline root CA and separate issuing subordinate

CA(s) topology wherever possible. Doing so mitigates many of the risks associated with compromised root CAs.

The TLS lab followed Microsoft's guidance to develop a highly secure offline stand-alone root CA coupled with an enterprise online issuing CA. The following CA installation and configuration how-to guidance aligns with that goal.

1.5.3.3.1 CA Prerequisites

The prerequisite steps to configure the CA(s) include:

- Build HSMrootca.int-nccoe.org and BaseSubCA.int-nccoe.org in accordance with the OS installation and configuration instructions in Section 1.5.2.1.2.
- Join BaseSubCA.int-nccoe.org to the already created int-nccoe.org domain.
- HSMrootca.int-nccoe.org and BaseSubCA.int-nccoe.org should have network connections to all the TLS lab subnets needed for CA certificate issuance.

1.5.3.3.2 Installation of Offline Root and Issuing CA

In this implementation scenario, the offline root CA is built, configured, and established as the root of the trust chain. The root CA is then configured to securely sign and issue certificates for all of its subordinates. Afterward, it is taken completely offline. Being taken offline includes complete power-down and highly secures physical storage of the root CA device (specifically the hard drive if possible).

Installation of the root CA through the Server Manager console can be done by installing Active Directory Certificate Services (ADCS). ADCS is used to create CAs and configure their role to issue and manage certificates. For instructions on installing ADCS on the root CA and issuing CA server, refer to the steps below:

1. In the **Server Manager**, select **Manage** > click on **Add Roles and Features**.
2. Follow the Add Roles and Features wizard > in **Select Installation Types**, select **Role-Based or feature installation**.
3. In **Select destination server**, confirm **Select a server from the server pool** is selected > select your local computer.
4. In **Select server roles** > under **Roles**, select **Active Directory Certificate Services** > click **Add Features**.
5. In **Select features** > click **Next**.
6. In **Active Directory Certificate Services** > click **Next**.
7. In **Select role services** > in **Roles**, select **Certification Authority**.
8. In **Confirm installation records** > click **Install**.
9. When installation is complete, click **Close**.

1.5.3.3.3 Offline Root CA Configuration

After installing ADCS, refer to the steps below to configure and specify cryptographic options for the root CA:

1. Run **Post-deployment Configuration** wizard > click on **Configure Active Directory Services** link.
2. In **Credentials**, read the credentials information. If needed, provide administrator credentials.
3. In **Role Services** > select **Certification Authority**.
4. In **Setup Type** > select **Standalone CA**.
5. In **CA Type** > select **Root CA**.
6. In **Private Key** > select **Create a new private key** to specify type of private key.
7. In **Cryptography for CA**:
 - Select a cryptographic provider: **RSA#SafeNet Key Storage Provider**.
 - Key Length = **2048**
 - Select the hash algorithm for signing certificates issued by this CA: **SHA256**.
8. In **CA Name** > specify the name of CA > **RootCA**.
9. For **Validity Period** > select **2 Years**.
10. Specify the database location > *C:\Window\system32\CertLog*.
11. Review the CA configuration and click **Configure**.
12. Click **Close** when the confirmation message appears.

To configure the CRL Distribution Point (CDP) and Authority Information Access (AIA) extensions on the root CA, follow the steps below:

1. In **Server Manager**, go to **Tools** > select **Certification Authority**.
2. Right-click **RootCA** > click **Properties**.
3. Click the **Extensions** tab. Ensure **Select Extension** is set to **CDP**.
4. In the **Specify locations from which users can obtain a certificate revocation list (CRL)**, do the following:
 - a. Select the entry
file://<ServerDNSName>/CertEnroll/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl and then click **Remove**. In **Confirm removal**, click **Yes**.

- b. Select the entry

http://<ServerDNSName>/CertEnroll/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl and then click **Remove**. In **Confirm removal**, click **Yes**.

5. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, click **Add**.
6. In **Add Location**, in **Location**, type *http://BaseSubCA/CertEnroll/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl* and then click **OK**. This returns to the CA properties dialog box.
7. On the **Extensions tab**, select the following checkboxes:
 - **Include in CRLs. Clients use this to find the Delta CRL locations.**
 - **Include in the CDP extension of issued certificates.**
8. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, select the entry that starts with **ldap://CN=CATruncatedName>,CRLNameSuffix>,CN=<ServerShortName>**.
9. On the **Extensions tab**, select the following checkbox:
 - **Include in all CRLs. Specifies where to publish in the Active Directory when publishing manually.**
 - In **Specify locations, users can obtain a certificate revocation list (CRL)**. Select the entry **C:\\Windows\\system32\\CertSrv\\CertEnroll\\<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**.
10. On the **Extensions tab**, select the following checkboxes:
 - **Publish CRLs to this location.**
 - **Publish Delta CRLs to this location.**
11. Change **Select extension** to **Authority Information Access (AIA)**.
12. In the **Specify locations, users can obtain a certificate revocation list (CRL)** do the following:
 - a. Select the entry *http://<ServerDNSName>/CertEnroll/<ServerDNSName>_<CaName><CertificateName>.crt* and then click **Remove**. In **Confirm removal**, click **Yes**.

- b. Select the entry
file://<ServerDNSName>/CertEnroll/<ServerDNSName>_<CaName><CertificateName>.crt and then click **Remove**. In **Confirm removal**, click **Yes**.
13. In **Specify locations, users can obtain a CRL**, click **Add**.
14. In **Add Location**, in **Location**, type
http://BaseSubCA/CertEnroll/<ServerDNSName>_<CaName><CertificateName>.crt and then click **OK**. This returns to the CA properties dialogue box.
15. On the **Extensions** tab, select the following checkbox:
 - **Include in the AIA of issued certificates.**
16. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, select the entry that starts with **ldap://CN=CATruncatedName>,CN=AIA,CN=PublicKeyServices**.
17. On the **Extensions** tab, select the following checkbox:
 - **Include in the AIA extension of issued certificates.**
18. In **Specify locations, users can obtain a certificate revocation list CRL**. Select the entry
C:\\Windows\\system32\\CertSrv\\CertEnroll\\<ServerDNSName>_<CaName><CertificateName>.crt.
19. On the **Extensions** tab, ensure **AIA extension of issued certificates** is not selected.
20. When prompted to restart Active Directory Certificate Services, click **No**. Restart that service later.
21. Go back to **RootCA** and expand folders to right-click on **Revoked Certificates** > select **All Tasks** > click **Publish**.
22. When prompted to Publish CRL, select **New CRL** > click **OK**.
23. To configure the Registry Settings, run cmd as an administrator and type the following commands:

```
certutil -setreg CA\\ValidityPeriod "Years"  
certutil -setreg CA\\ValidityPeriodUnits 2
```

```

Administrator: Command Prompt
Microsoft Windows [Version 6.0.6002]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Windows\system32>certutil -setreg CA\ValidityPeriod "Years"
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\ROOTCA1-CA\ValidityPeriod:
Old Value:
    ValidityPeriod REG_SZ = Years
New Value:
    ValidityPeriod REG_SZ = Years
CertUtil: -setreg command completed successfully.
The CertSvc service may need to be restarted for changes to take effect.

C:\Windows\system32>certutil -setreg CA\ValidityPeriodUnits 2
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\ROOTCA1-CA\ValidityPeriodUnits:
Old Value:
    ValidityPeriodUnits REG_DWORD = 1
New Value:
    ValidityPeriodUnits REG_DWORD = 2
CertUtil: -setreg command completed successfully.
The CertSvc service may need to be restarted for changes to take effect.

C:\Windows\system32>_

```

certutil -setreg CA\DSConfigDN "CN=Configuration,DC=int-nccoe,DC=org"

```

Administrator: Command Prompt

C:\Windows\system32>certutil -setreg CA\DSConfigDN "CN=Configuration,DC=int-nccoe,DC=org"
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\ROOTCA1-CA\DSConfigDN:
New Value:
    DSConfigDN REG_SZ = CN=Configuration,DC=int-nccoe,DC=org
CertUtil: -setreg command completed successfully.
The CertSvc service may need to be restarted for changes to take effect.

```

certutil -setreg CA\DSDomainDN "DC=int-nccoe,DC=org"

```

Administrator: Command Prompt

C:\Windows\system32>certutil -setreg CA\DSDomainDN "DC=int-nccoe,DC=org"
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\ROOTCA1-CA\DSDomainDN:
New Value:
    DSDomainDN REG_SZ = DC=int-nccoe,DC=org
CertUtil: -setreg command completed successfully.
The CertSvc service may need to be restarted for changes to take effect.

C:\Windows\system32>

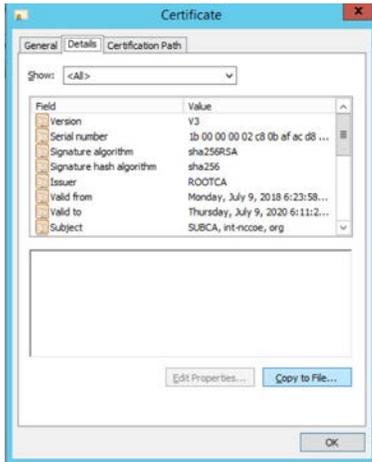
```

24. For it to accept the new values, restart services > go to **Administrative Tools** > double-click **Certification Authority**.
25. Select the **RootCA** > right-click to select **All Tasks** > click **Start Service**.
26. Go back to **RootCA** to expand folders > right-click on **Revoked Certificates** > select **All Tasks** > click **Publish** to publish revoked certificates.

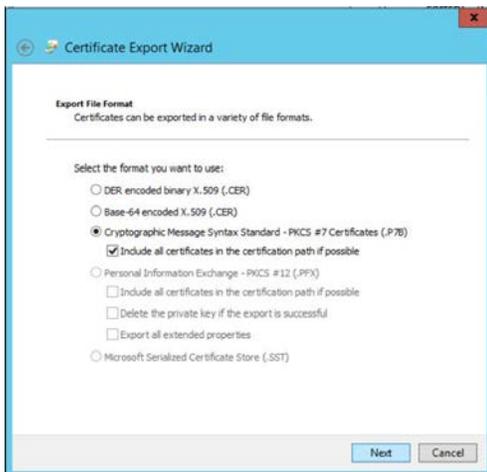
1.5.3.3.4 Enterprise Subordinate/Issuing CA Configuration

After installing ADCS, follow the steps below to configure and specify cryptographic options for the issuing CA:

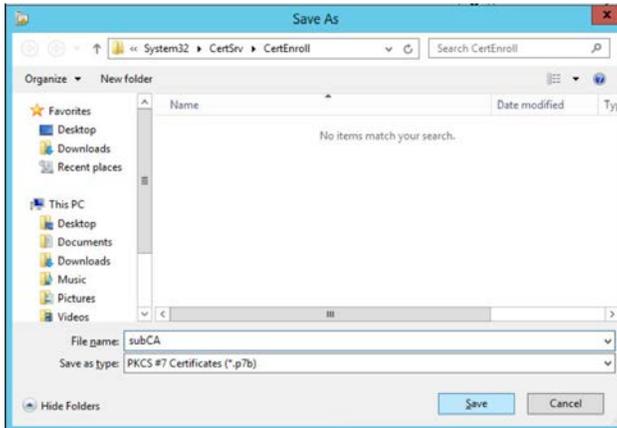
1. Run **Post-deployment Configuration** wizard > click on **Configure Active Directory Services** link.
2. In **Credentials**, read the credentials information. If needed, provide administrator credentials.
3. In **Role Services** > select **Certification Authority**.
4. In **Setup Type** > select **Enterprise CA**.
5. In **CA Type** > select **Subordinate CA**.
6. In **Private Key** > select **Create a new private key** to specify type of private key.
7. In **Cryptography for CA**:
 - Select a cryptographic provider: **RSA#SafeNet Key Storage Provider**.
 - Key Length = **2048**
 - Select the hash algorithm for signing certificates issued by this CA: **SHA256**.
8. In **CA Name** > specify the name of the CA > **BaseSubCA**.
9. In **Certificate Request** > select **Save a certificate request to file on the target machine** > specify folder location > *C:\BaseSubCA.int-nccoe.org_int-nccoe-BASESUBCA-CA.req*.
10. In **CA Database** > specify the folder location for the certification database > **C:\Windows\system32\CertLog**.
11. In **Confirmation** > confirm configurations and select **Configure** > click **Close**.
12. Copy the BaseSubCA request file from the BaseSubCA server to the RootCA server at **C:\Windows\System32\CertServ\CertEnroll**.
13. Copy *rootCA.crl* and *rootCA.crt* to the BaseSubCA server at **C:\Windows\System32\CertServ\CertEnroll**.
14. To issue a certificate to the BaseSubCA server from the RootCA server, go to **Administrative Tools** > double-click **Certification Authority**.
15. Select **BaseSubCA** > right-click to select **All Tasks** > click **Submit new request**.
16. Select and open the request file in the dialogue box.
17. Go back to the **Certification Authority** > select **BaseSubCA** and expand folders > click on **Pending Requests**.
18. Right-click the pending certificate > right-click to select **All Tasks** > click **Issue**.
19. Go to **Issued Certificates** to view the issued certificate.
20. Double-click on the issued certificate.
21. Go to the **Details** tab > click **Copy to File**.



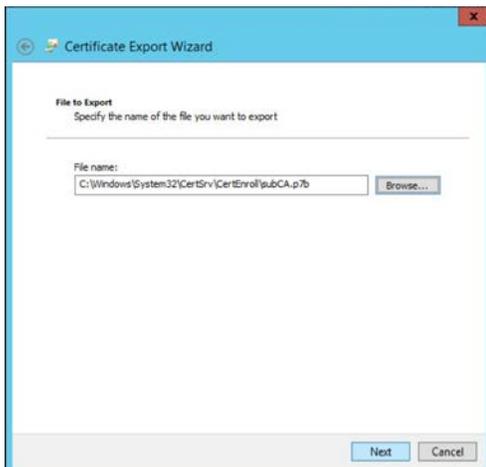
22. Follow the Certificate Export wizard and select the desired format:



23. Save the file as **subCA** > file type is **PKCS #7 Certificates (*.p7b)**.



24. Specify the file name to export:



25. Complete the Certificate Export Wizard by confirming settings > click **Finish**.

26. In **Export was successful** > click **OK**.

27. Copy **subCA.p7b** from the RootCA server at **C:\Windows\System32\CerServ\CertEnroll** to the BaseSubCA server at **C:\Windows\System32\CerServ\CertEnroll**.

28. On the BaseSubCA server > shift right-click > open the command prompt.

29. Publish the CA Root certificate into Directory Services with the following command:

```
certutil -dspublish -f (tab to rootCA.crt file) RootCA
```

```

Administrator: C:\Windows\system32\cmd.exe
C:\Windows\System32\CertSvc\CertEnroll>certutil -dsPublish -f rootCA_ROOTCA-CA.c
rt RootCA
ldap:///CN=ROOTCA-CA,CN=Certification Authorities,CN=Public Key Services,CN=Serv
ices,CN=Configuration,DC=int-nccoe,DC=org?cACertificate
Certificate added to DS store.
ldap:///CN=ROOTCA-CA,CN=ATA,CN=Public Key Services,CN=Services,CN=Configuration,
DC=int-nccoe,DC=org?cACertificate
Certificate added to DS store.
CertUtil: -dsPublish command completed successfully.
C:\Windows\System32\CertSvc\CertEnroll>

```

30. To publish the crl file, type the following command:
certutil -dsPublish -f (tab to .crl file)

```

Administrator: C:\Windows\system32\cmd.exe
C:\Windows\System32\CertSvc\CertEnroll>certutil -dsPublish -f rootCA_ROOTCA-CA.c
rt RootCA
ldap:///CN=ROOTCA-CA,CN=Certification Authorities,CN=Public Key Services,CN=Serv
ices,CN=Configuration,DC=int-nccoe,DC=org?cACertificate
Certificate added to DS store.
ldap:///CN=ROOTCA-CA,CN=ATA,CN=Public Key Services,CN=Services,CN=Configuration,
DC=int-nccoe,DC=org?cACertificate
Certificate added to DS store.
CertUtil: -dsPublish command completed successfully.
C:\Windows\System32\CertSvc\CertEnroll>certutil -dsPublish -f ROOTCA-CA.crl
ldap:///CN=ROOTCA-CA,CN=RootCA,CN=CDP,CN=Public Key Services,CN=Services,CN=Conf
iguration,DC=int-nccoe,DC=org?certificateRevocationList?base?objectClass=cRLDist
ribution?int?certificateRevocationList
Base CRL added to DS store.
CertUtil: -dsPublish command completed successfully.
C:\Windows\System32\CertSvc\CertEnroll>

```

31. Set the **Domain Policy** to make the RootCA trusted by all domain computers.
32. Install the certificate in the subCA server > go to **Administrative Tools** > double-click **Certification Authority**.
33. Select the CA > right-click to select **All Tasks** > click **Install CA Certificate**.
34. Select the *.p7b* file to complete the CA installation.
35. A warning message will be received that the revocation server is offline > click **OK** to ignore the message.
36. Power down the RootCA server.
37. Go to **Administrative Tools** > right-click the CA > select **All Tasks** > click **Start Service** to start services.
38. Install *.crt* files on the Default Domain Policy.
39. Go to the domain controller (DC).
40. Go to **Administrative Tools** > open **Group Policy Management** console.
41. Go to the organization's domain > right-click the **Default Domain Policy** folder > select **Edit**.

42. Navigate to **Computer Configuration**, go to **Policies > Window Settings > Security Settings > Public Key Policies** > right-click **Intermediate Certification Authorities** > select **Import**.
43. Follow the **Certificate Import Wizard** > click **Next**.
44. Select the **subCA.crt** file to import > click **Next** to import file.
45. Confirm details > click **Finish**.
46. A dialogue box will pop up to confirm **The import was successful**.
47. Go to **Trusted Root Certification Authority** folder and right-click> select **Import**.
48. Follow the **Certificate Import Wizard** > click **Next**.
49. Select the **rootCA.crt** file to import > click **Next** to import file.
50. Confirm details > click **Finish**.
51. A dialogue box will appear to confirm **The import was successful**.

1.5.4 Database Services

1.5.4.1 Microsoft SQL Database Services

Microsoft SQL (MSQL) Server is a relational database management system developed by Microsoft. As a database server and a software product, its primary function is to store and retrieve data as requested by other software applications. MSQL can operate on the same or another computer across a network.

1.5.4.1.1 Prerequisites for MSQL Database Services

The information below is Microsoft's recommended minimum for default installation of MSQL. An organization's requirements may differ. However, all applications can benefit from more than the minimum resources in a production environment.

- 1.4 GHz 64-bit processor
- 1 GB RAM
- 6 GB disk space
- administration privileges (local installations must run Setup as an administrator)

One MSQL database was used for the TLS lab build to support the Venafi TPP server. This guide installs only the basic MSQL application on a server. This prepares the specific configurations that are discussed in the Venafi TPP How -To guidance section. As a prerequisite, see the OS installation instructions in Section [1.5.2.1.2](#) to build the VTPPTrustDB.int-nccoe.org server.

1.5.4.1.2 Installation of MSQL Database Services

To install MSQL on a Windows 2016 Server, follow the Microsoft steps in the link below:

- Download here: https://www.microsoft.com/en-us/sql-server/sql-server-downloads?&OCID=AID739534_SEM_at7DarBF&MarinID=sat7DarBF_340829462634_microsoft%20sql%20download_e_c_68045082145_kwd-343189224165
- Install and configure here: <https://docs.microsoft.com/en-us/sql/database-engine/install-windows/install-sql-server-from-the-installation-wizard-setup?view=sql-server-2017>
- Install MSQL as a stand-alone server.
- Specify the Database Engineer Configuration in step 15 by selecting SQL Server Administrators.

1.5.4.2 *MariaDB Database Services*

The original inventors of MySQL developed the MariaDB server, which is highly compatible with MySQL. This allows a drop-in replacement capability with library binary parity and exact matching with MySQL's application programming interfaces and commands.

Like MySQL, the open-source version of MariaDB can scale and performs as well as most enterprise database servers. The TLS lab uses the MariaDB to serve its public-facing (DMZ) web-based TLS services described in this document.

1.5.4.2.1 Prerequisites for MariaDB Database Services

The host named dmzdb.ext-nccoe.org should have already been set up within the Fedora OS how-to guidance of Section [1.5.2.2.2](#). Complete this setup prior to installing the MariaDB server.

1.5.4.2.2 Installation of MariaDB Database Services

- To download and install MariaDB, please refer to the [fedoraproject.org](https://fedoraproject.org/wiki/MariaDB) guidance at <https://fedoraproject.org/wiki/MariaDB>

1.5.4.2.3 Configuration of MariaDB Database Services

MariaDB is used to serve dynamic web content with the Drupal application. All three web servers used in the DMZ must be configured via Drupal to point to one database. As a result, the database must be configured to accept connections from the Drupal web servers. MariaDB can be configured by using the Fedora Linux command line. To start, first set up a secure password for the root and any other administrative accounts (see the MariaDB setup instructions on how to specify other accounts). Log in to the dmzdb.int-nccoe.org by using the local command line shell or secure remote administration client (ssh, putty, openssh). Once logged into the system, use the following command to launch MariaDB from the Fedora Linux:

```
[root@dmzdb ~]# mysql -p
```

Note: Although the root account is displayed here as the login account, configuring MariaDB with the root user in a production environment is not recommended.

Configure the database to allow remote connections from either the IP addresses or host names used in the TLS lab. If the IP addresses and host names were customized (apache1: 192.168.4.2, apache2: 192.168.4.3, iis1: 192.168.4.4), please double-check and change the IP addresses in the database by using the commands below. If custom host names were used in place of the IP addresses, the database DNS or host resolution is set to properly resolve to the right IP addresses.

```
[root@dmzdb ~]# mysql -p

Enter password:

Welcome to the MariaDB monitor. Commands end with ; or \g.
Your MariaDB connection id is 1012018
Server version: 10.2.16-MariaDB MariaDB Server

Copyright (c) 2000, 2018, Oracle, MariaDB Corporation Ab and others.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

MariaDB [(none)]> create database EXT_NCCOE_DB;

MariaDB [(none)]> grant all privileges on EXT_NCCOE_DB.* to
'EXTADMIN'@'192.168.4.2' IDENTIFIED BY 'YOUR PASSWORD';

MariaDB [(none)]> grant all privileges on EXT_NCCOE_DB.* to
'EXTADMIN'@'192.168.4.3' IDENTIFIED BY 'YOUR PASSWORD';

MariaDB [(none)]> grant all privileges on EXT_NCCOE_DB.* to
'EXTADMIN'@'192.168.4.4' IDENTIFIED BY 'YOUR PASSWORD';

MariaDB [(none)]> quit;
```

Add rules to the local Linux firewall to allow database traffic inbound. Please use the following commands to allow database traffic to inbound ports on the MariaDB server:

- Type the following command to allow database connections to Apache:

```
iptables-I INPUT -p tcp -dport 3306 -mstate --state related, ESTABLISHED, new -
j ACCEPT
```

1.5.5 TLS Web Services

1.5.5.1 Microsoft Internet Information Services

The web server (IIS) role in Windows Server 2012 provides a means for hosting websites, services, and applications. IIS information can be shared with users on the internet, an intranet, or an extranet. IIS is a unified web platform that integrates IIS, ASP.NET, File Transfer Protocol services, Personal Home Page Hypertext Preprocessor (PHP), and Windows Communication Foundation.

The TLS lab utilized the IIS server as a public-facing member of a load balance web cluster for public-facing internet services. It was also used as an intranet server to simulate an employee web-based knowledge management system that is internal to an organization.

1.5.5.1.1 IIS Prerequisites

Complete the following prerequisite steps prior to installing and configuring IIS:

- Server iis2.int-nccoe.org should ideally be a member of the domain for more streamlined TLS certificate management.
- The IIS administrator must have Request Certificates permission on the issuing CA.
- The iis1.int-nccoe.org and iss2.int-nccoe.org servers should be set up per Section [1.5.2.1.2](#).
- Server iis1.int-nccoe.org should be used for the public-facing web-based cluster.
- Server iis2.int-nccoe.org should be used as the internal intranet server.

1.5.5.2 IIS Installation

IIS is the topic of this section, however, the PHP is a key component of the IIS installation for the TLS lab implementation of the iis1.int-nccoe.org internet-facing server. PHP is a script language and interpreter and a server-side language that assists IIS and Drupal in serving dynamic web content.

Please follow the instructions in the link below to install IIS and PHP. The iis2.int-nccoe.org server can be set up without PHP installed. Please follow the same instructions below for the iis2 server—skip the PHP part of the installation process.

- <https://docs.microsoft.com/en-us/iis/application-frameworks/scenario-build-a-php-website-on-iis/configuring-step-1-install-iis-and-php>

Windows 2012 Server provides several methods for enrolling certificates: two of these are the Certificate Enrollment Policy (CEP) and Certificate Enrollment Service (CES). The CEP web service enables users and computers to obtain certificate enrollment policy information. This information includes what types of certificates can be requested and what CAs can issue them. CES provides another web service that allows users and computers to perform certificate enrollment by using the hypertext transfer protocol secure (https). To separate traffic, the CES can be installed on a computer that is separate from the CA. Together with the CEP web service, CES enables policy-based certificate enrollment when the client computer is not a member of a domain or when a domain member is not connected to the domain. CEP/CES also enables cross-forest, policy-based certificate enrollment.

For the purpose of the lab, the IIS configuration option selected for authentication type for the CES is **Windows integrated authentication**. This option provides Kerberos authentication for devices connected to the internal network and joined to a domain. The service account selected is the **Use the built-in application pool identity**.

To configure the SSL protocol to encrypt network traffic, obtain a certificate for IIS, and configure https on the default website, please refer to the link below.

- <https://social.technet.microsoft.com/wiki/contents/articles/12485.configure-ssl-tls-on-a-web-site-in-the-domain-with-an-enterprise-ca.aspx>

1.5.5.3 Apache Web Services

The Apache HTTP Server is a free and open-source cross-platform web server software, released under the terms of Apache License 2.0. Apache is developed and maintained by an open community of developers under the Apache Software Foundation.

1.5.5.3.1 Apache Web Services Prerequisites

The Apache web server was used extensively throughout the TLS lab architecture to demonstrate the various means of automated and manual management of TLS certificates. The following servers should be built in accordance with the instructions in Section [1.5.2.2.2](#).

- *apache1.ext-nccoe.org*
- *apache2.ext-nccoe.org*
- *ws1.int-nccoe.org*
- *ws2.int-nccoe.org*
- *ws3.int-nccoe.org*

1.5.5.3.2 Apache Installation

PHP is a key component of the Apache installation for the TLS lab implementation of all of the above web servers. PHP assists Apache and Drupal in serving dynamic web content. Please follow the instructions below for installing Apache and PHP.

For the Apache web server installation, please refer to this guidance: https://docs.fedoraproject.org/en-US/fedora/f28/system-administrators-guide/servers/Web_Servers/

All Drupal installations have dependencies on the base PHP application and its supplemental modules. In addition to the base PHP installation, also install the additional modules by using the following command.

- ```
dnf install drush php php-mysqli php-json php-mbstring php-gd php-dom php-xml php-simplexml php-cli php-fpm php-mysqlnd php-pdop-gd php-dom php-xml php-simplexml php
```

#### 1.5.5.3.3 Apache Web Services Configuration

The TLS lab enabled https on the Apache web servers. For instructions on setting up OpenSSL, refer to the “Using mod\_ssl” section from the following link: <https://docs.fedoraproject.org/en-US/quick-docs/getting-started-with-apache-http-server/>

To allow http and https connections through the local Fedora firewall to Apache, perform the following steps:

- Type the following command to allow http connections to Apache:

```
iptables-I INPUT -p tcp -dport 80 -mstate --state related, ESTABLISHED, new -j ACCEPT
```

- Type the following command to allow https connections to apache:

```
iptables-I INPUT -p tcp -dport 443 -mstate --state related, ESTABLISHED, new -j ACCEPT
```

Save the newly created firewall rules with the following command: `iptables-save`

#### 1.5.5.4 *Drupal Web Content Management Services*

Drupal is a scalable, open platform for web content management. Drupal can be installed on multiple OSs, including, Fedora, CentOS, and IIS. The TLS lab utilized Drupal to serve web pages on all three of the load balanced web servers in the public-facing DMZ.

##### 1.5.5.4.1 *Drupal Prerequisites*

- PHP 5.5.9 or higher
- MySQL 5.5.3 or MariaDB 5.5.20
- Apache or IIS web server

##### 1.5.5.4.2 *Drupal Web Content Management System Download and Installation*

One server should run throughout the setup process, including the database setup. The remaining two servers should be set up to point to the existing database once the first server has been set up. All web servers should be set up to use MariaDB, **not MSQl**. Use the guidance below for download, installation, and configuration of Drupal to simulate the TLS lab architecture:

- download: <https://www.drupal.org/download>
- Apache installation and configuration: <https://www.drupal.org/docs/7/install>
- IIS installation and configuration: <https://www.drupal.org/docs/develop/local-server-setup/windows-development-environment/installing-on-windows-server>

##### 1.5.5.4.3 *Web Services Drupal Configuration*

A web service is a software system designed to support machine-to-machine interaction over a network. A web service is normally accessed over a network and then executed on a remote system hosting the requested services. Web services protocols normally use application programming interfaces (APIs) based on RESTful, simple object access protocol (SOAP), and extensible markup language (XML)

protocols. It is a best practice to execute web services that carry critical personally identifiable information and other sensitive information by using TLS-based encrypted communication channels.

The TLS lab tested implementation of passive monitoring for TLS-enabled web services traffic. The rationale behind this approach is covered in the Symantec How-To guide section of this document. In [Appendix A](#), Passive Inspection, see the full description of how the passive monitoring network was configured.

The web services servers are configured to test the basic passive TLS monitoring capability and are not typical of a fully operational web services implementation. The RESTful, SOAP, and XML protocols are not used in the TLS Lab. Rudimentary machine-to-machine communication over a secured TLS network is configured within each DMZ web server by using JavaScript, PHP, and Drupal's in-line What-You-See-Is-What-You-Get (also known as WYSIWYG) hypertext markup language (HTML) content creation editor. A simple PHP script that was created for each web service prompted each of the three web services servers to retrieve and push its current times to the main web server. The JavaScript included in the Drupal-based DMZ servers was set to grab updates of the time each second by using https connectivity. Use the steps below to re-create this setup.

### Part 1: Drupal DMZ Servers Configuration

1. Log in to Drupal by using the content administrator with enough rights to create a basic page.
2. Navigate to the following administrative menu item (top of the page on the left side, then use the links within the Content administration page itself to navigate to the remaining sections):  
**Content > Add Content > Basic Page**
3. Verify that a page is displayed that allows entry of data by using a **Title** and **Body** HTML form.
4. Give this page any title.
5. Before populating the body section of the page, ensure that the **Text Format** is set to **Full Html and PHP**. If that selection is not present, enable the **PHP Filter** module in the Drupal **Modules** section of Drupal, and try again.
6. Upon completing step 5, paste the following code into the body of the new document:

```
<div id="timeid"></div>

<?php

$serveraddress = $_SERVER['SERVER_ADDR'];

$javagetime = <<<EOFF
<script>
mydata = "TEST";
function ExportValues(mydata) {
```

```

 var xmlhttp;
 if (window.XMLHttpRequest) {
 // code for modern browsers
 xmlhttp = new XMLHttpRequest();
 } else {
 // code for IE6, IE5
 xmlhttp = new ActiveXObject("Microsoft.XMLHTTP");
 }
 xmlhttp.onreadystatechange = function() {
 if (this.readyState == 4 && this.status == 200) {
 document.getElementById("timeid").innerHTML =
this.responseText;
 }
 };

 xmlhttp.open("GET", "https://$serveraddress/PHPTIME.php", true);
 xmlhttp.send();
 }

 ExportValues(mydata);
 setInterval(function(){ ExportValues(mydata); }, 1000);
</script>

EOFF;
echo $javagetime;

?>

```

7. Click on the **Publishing options** tab below, then make sure that **Published** and **Promoted to front page** are selected as options.
8. **Save** the page.
9. Repeat these steps for each web services server.

## Part II: Drupal DMZ Servers Configuration

The code above in Part I instructs the DMZ web server to connect to itself and execute the script *PHPTIME.php* within its own Drupal directory. This file will be created here in Part II. The *PHPTIME.php* file uses a curl script to simulate secure TLS server-to-server communication between the DMZ web server and its designated web services server. Follow the steps below to create this file on *all* the DMZ web servers.

1. Log in to the local web administration account for each of the three DMZ-based web servers. Navigate to the local Drupal stored file system where Drupal is served to the public. On Apache servers, this will be `/var/www/html/<DRUPAL DIRECTORY NAME USED>`. On IIS servers, this will be the Drupal document root for the website instantiation.

2. Launch a text editor (notepad++ or notepad for Windows or VIM or VI editor for Linux), then paste the following into that file:

```
<?php
 header("Access-Control-Allow-Origin: *");
 $ch = curl_init();

 curl_setopt($ch, CURLOPT_URL, 'https://ws2.int-nccoe.org');
 curl_setopt($ch, CURLOPT_RETURNTRANSFER, 1);
 curl_setopt($ch, CURLOPT_SSL_VERIFYHOST, false);
 curl_setopt($ch, CURLOPT_SSL_VERIFYPEER, false);

 $result = curl_exec($ch);
 if (curl_errno($ch)) {
 echo 'Error:' . curl_error($ch);
 }
 curl_close ($ch);

 echo $result;
?>
```

3. The following line will need to be changed on each DMZ web server and customized with the individual host name for the web services server assigned to the specific DMZ web server. Each DMZ web server should have its own individual web services server:

```
curl_setopt($ch, CURLOPT_URL,'https://CHANGE TO YOUR MACHINE NAME');
```

4. Save this file with a .php extension into the root base directory of the Drupal site created for this demonstration.

### Web Services Server Configuration

The web services server must be configured to check its own time and send the results back to the requesting DMZ web server via secure communication. Use the following guidance to set up the web services server.

1. Log in to the command line for each web services server, and navigate to the Apache document root configured in the *httpd.conf* file for Apache. In most cases it is */var/www/html*.
2. Open a VIM/VI editor and paste the following into that file:

```
<?php

$sourceip = $_SERVER['HTTP_ORIGIN'];

if (isset($_SERVER["HTTP_ORIGIN"]) === true) {
 $origin = $_SERVER["HTTP_ORIGIN"];
}
```

```

$allowed_origins = array(

 // ANY
 $_SERVER['HTTP_ORIGIN']

 // SPECIFIC
 "https://192.168.4.2",
 "https://apache1.ext-nccoe.org",
 "https://tls.nccoe.org",
 "https://apache2.ext-nccoe.org",
 "https://192.168.4.3",
 "https://iis1.ext-nccoe.org",
 "https://192.168.4.4"
);
if (in_array($origin, $allowed_origins, true) === true) {
 header('Access-Control-Allow-Origin: ' . $origin);
 header('Access-Control-Allow-Credentials: true');
 header('Access-Control-Allow-Methods: POST');
 header('Access-Control-Allow-Headers: Content-Type');
}
if ($_SERVER["REQUEST_METHOD"] === "OPTIONS") {
 exit; // OPTIONS request wants only the policy, we can stop
here
}
}

$timetime = exec('date');

echo "WEB SERVICES SERVER2's TIME AN DATE IS: ". $timetime;

?>

```

3. Remember to save the file in the document root directory under the same name used in the previous section with the .php extension.
4. Ensure the Apache service is running: `service httpd restart`

### Web Services Testing Process

1. Navigate to the public IP of the Drupal web servers (should be the F5 virtual ip or if behind a firewall, the IP address of the firewall used to NAT to the web server cluster behind the F5).
2. There should be at least three Basic Pages listed on the main site landing page. These should be the pages created in this section to point to the web services server.
3. Choose one by clicking on its title or **Read more** link beside the title.
4. The time should be automatically updating each second to indicate the web server is using its designated web services server to check time via TLS connection (indicated by the https).

5. If the time updates are not being seen, there could be an issue with the browser application accepting the valid certificate. If self-signed untrusted certificates instead of a trusted certificate are being used on the DMZ web servers, then the web client used (Chrome, Internet Explorer, or Edge) may not trust the individual server being accessed. To discover the issue, press the F12 key on the keyboard, then select the **Console** tab. If there is an error stating:

Net::ERR\_CERT\_AUTHORITY\_INVALID or any other certificate validation error with an associated IP address, open a new tab and navigate directly to the IP address listed by using 192.168.3.85. If there is the standard certificate error for an untrusted site, then accept the risk if this is a laboratory environment. The time should pop up afterward, and the other tabs with the Drupal time connection will also work now. If this is production system, then a valid certificate will need to be placed on the machine with the IP listed. The client that browses that machine should trust the certificate.

### 1.5.5.5 Mail Services

The TLS lab utilizes a Simple Mail Transfer Protocol (SMTP) service to accept alerts from all the configured components on the network. The SMTP service was created on a Linux server running Fedora. The mail system was composed of a Dovecot Mail Transfer Agent (MTA) and a Postfix Mail User Agent (MUA). The following section provides guidance on download, installation, and configuration of each service.

#### 1.5.5.5.1 Mail Services Prerequisites

Before installing Dovecot and Postfix, set up the mail1.int-nccoe.org server by using the guidance in Section [1.5.2.2.2](#).

#### 1.5.5.5.2 Installation and Configuration of Mail Services Postfix Mail Transfer Agent

Postfix is a free and open-source mail transfer agent that routes and delivers electronic mail. To download and install the Postfix MTA, follow the instructions in the following link:

- [https://docs.fedoraproject.org/en-US/Fedora/12/html/Deployment\\_Guide/s3-email-mta-postfix-conf.html](https://docs.fedoraproject.org/en-US/Fedora/12/html/Deployment_Guide/s3-email-mta-postfix-conf.html)

Note: The actual *main.cf* file used in the TLS lab build is in [Appendix F](#).

#### 1.5.5.5.3 Installation and Configuration of Mail Services Dovecot Mail Transfer Agent

Dovecot is an open-source Internet Message Access Protocol (IMAP) and Post Office Protocol 3 Mail User Agent server for Linux systems. It allows TLS administrators to manage and view email received by the Postfix server. To download and install the Dovecot MUA, please refer to the instructions in the following link:

- <https://wiki.dovecot.org/BasicConfiguration>

Note: The actual *dovecot.conf* file used in the TLS lab build is in Appendix F.

### 1.5.5.6 Log Aggregation and Correlation Services

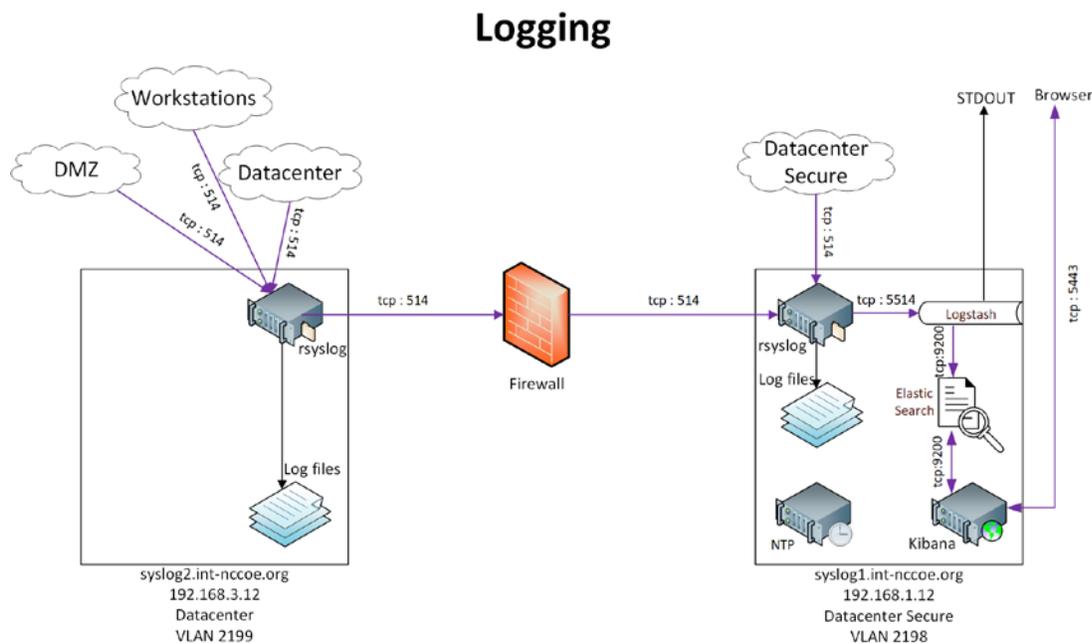
“ELK” stands for three open-source projects:

- Elasticsearch—a search and analytics engine
- Logstash—a server-side data processing pipeline that ingests data from multiple sources simultaneously, transforms it, and then sends it to a “stash” like Elasticsearch
- Kibana—lets users visualize data with charts and graphs in Elasticsearch

The TLS lab utilized the ELK stack log aggregation and correlation services to manage and visualize the remote logging services for all capable supplemental and collaborator products.

The following diagram depicts a view of the TLS lab logging infrastructure.

Figure 1-3 TLS Lab Logging Infrastructure



#### 1.5.5.6.1 Prerequisites for Log Aggregation and Correlation Services

In accordance with the logging architecture above, the TLS lab utilized the hosts below. Both hosts must be configured with Fedora, based on the OS configuration guidance in Section 1.5.2.2.2. Configure both servers with rsyslog.

- `syslog1.int-nccoe.org`
- `syslog2.int-nccoe.org`
- Logstash requires Java 8 or Java 11.

#### 1.5.5.6.2 Remote System Logging Services

Rsyslog is an open-source software utility used on UNIX and UNIX-like computer systems for forwarding log messages in an IP network.

- To install rsyslog use the command `dnf install rsyslog`

For more information on configuring rsyslog, refer to the following link:

- [https://docs.fedoraproject.org/en-US/fedora/rawhide/system-administrators-guide/monitoring-and-automation/Viewing\\_and\\_Managing\\_Log\\_Files/#](https://docs.fedoraproject.org/en-US/fedora/rawhide/system-administrators-guide/monitoring-and-automation/Viewing_and_Managing_Log_Files/#)

#### 1.5.5.6.3 Elasticsearch Installation and Configuration

Elasticsearch is a search engine based on the Lucene library. It provides a distributed, multitenant-capable full-text search engine with an http web interface and schema-free JavaScript Object Notation documents. Elasticsearch is developed in Java.

To install and configure Elasticsearch, please refer to the following link:

- <https://www.elastic.co/guide/en/elasticsearch/reference/current/rpm.html>

#### 1.5.5.6.4 Kibana Installation and Configuration

Kibana is an open-source data visualization plug-in for Elasticsearch and provides visualization capabilities on top of the content indexed on an Elasticsearch cluster. Users can create bar, line, and scatter plots (or pie charts) and maps on top of large volumes of data.

To install and configure Kibana, please refer to the following link:

- <https://www.elastic.co/guide/en/kibana/current/rpm.html>

#### 1.5.5.6.5 Logstash Installation and Configuration

Logstash is an open-source, server-side data processing pipeline that ingests data from a multitude of sources simultaneously, transforms it, and then sends it to the user's favorite stash.

To install and configure Logstash, please refer to the following link:

- <https://www.elastic.co/guide/en/logstash/current/installing-logstash.html#package-repositories>

## 1.5.6 DevOps Services

The NCCoE undertook a limited DevOps demonstration using a Kubernetes cluster. This limited demonstration included basic DevOps functionality for automated system and application deployment. We showed automated management of TLS server certificates in a container-based environment by using Kubernetes with Docker, NGINX, and Jetstack Cert-Manager

### 1.5.6.1.1 Kubernetes Installation and Configuration

Instructions for installing Kubernetes are available at the following link:

- <https://kubernetes.io/docs/setup/>

We installed Kubernetes on three CentOS Linux systems (cluster1, cluster2, cluster3.int-nccoe.org).

### 1.5.6.1.2 Weave

We used Weave as the virtual network to facilitate communications between the Kubernetes primary and nodes. Instructions for installing Weave can be found at the following link:

- <https://www.weave.works/docs/net/latest/install/>

### 1.5.6.1.3 Docker Installation and Configuration

We used the community edition of Docker with Kubernetes. Instructions for installing Docker on CentOS are found at the following link:

- <https://docs.docker.com/install/linux/docker-ce/centos/>

### 1.5.6.1.4 Jetstack Cert-Manager Installation and Configuration

We installed Jetstack Cert-Manager on Kubernetes with the necessary components to request certificates from Venafi TPP by using the following command:

```
kubectl apply -f https://raw.githubusercontent.com/jetstack \
/cert-manager/venafi/contrib/manifests/cert-manager/with-rbac.yaml
```

This automatically created a namespace named “cert-manager,” which we used for the rest of our configuration.

### 1.5.6.1.5 NGINX Installation and Configuration

NGINX was used as the web server and ingress on Kubernetes. Certificates were associated with the NGINX ingress. Instructions for installing and configuring NGINX on Kubernetes are found at the following link:

- <https://www.nginx.com/>

In our implementation, we installed NGINX on Kubernetes with the following command into the cert-manager namespace.

```
kubectl create deployment nginx --image=nginx -n cert-manager
```

We then created a service for NGINX by using the following command:

```
kubectl create service nodeport nginx --tcp=80:80 -n cert-manager
```

## 2 Product Installation and Configuration Guides

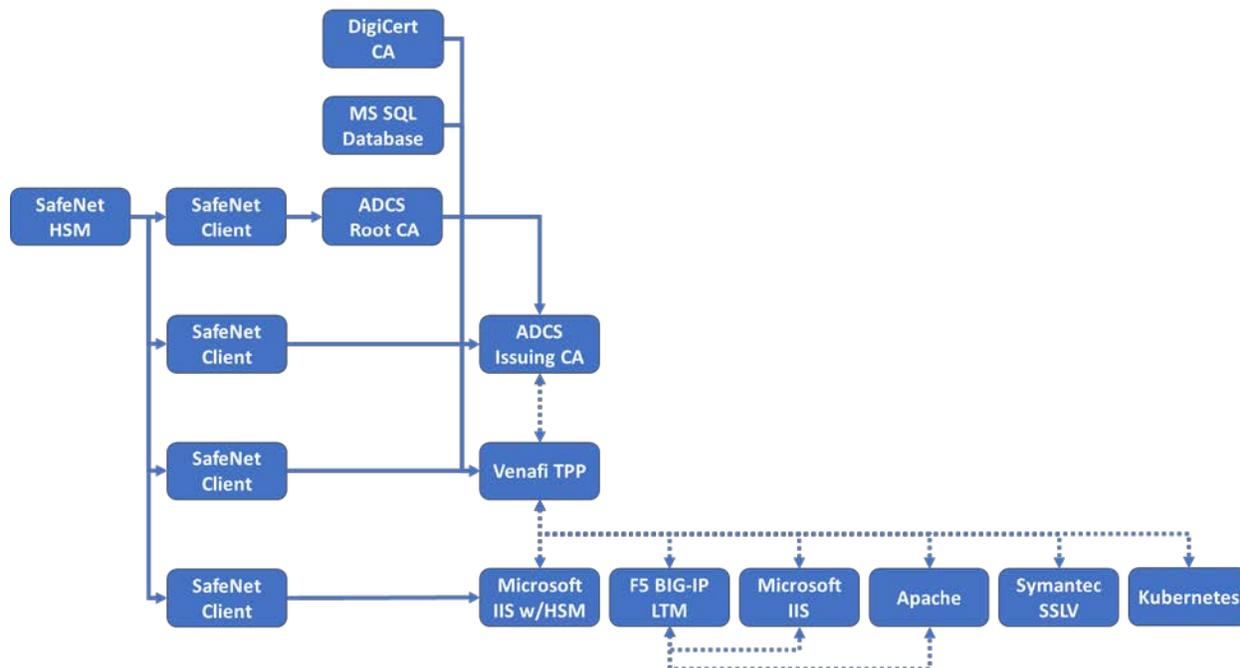
This section of the practice guide contains detailed instructions for installing and configuring all of the TLS collaborator products used to build an instance of the example solution. Each major subsection (2.1, 2.2, 2.x) is dedicated to a collaborator's product capability. Within each product capability section, descriptions of each product capability align with a Day 0, Day 1, and Day N concept. It is important to note that each day builds on the previous day(s) for prerequisites, and each collaborator capability does the same. So, if the implementer's intent is to fully replicate the TLS lab environment, then following the order of days and component installations will help make that endeavor more successful.

- **Day 0** provides how-to guidance from a first-day installation perspective. It is assumed the implementer is getting acclimated with the collaborator product. The implementer should complete all prerequisites, which include complete installations of other collaborator products in some instances or the Supporting Architecture described in [Section 1.3](#). The expectation is for only basic crucial configuration functions to get the system up and running. Otherwise, other configurations should be executed on Day 1, or there may be issues with prerequisites that have not been executed.
- **Day 1** assumes all Day 0 activities have been completed, including all prerequisites. Expected activities include how-to guidance on more advanced security configuration of functioning in the TLS environment. Day 1 also assists the implementer with configuration guidance for integration with any other collaborator product capabilities.
- **Day N** assists the implementer with all necessary configurations and integrations of systems that help facilitate ongoing security management and maintenance. In most cases, the minimum Day N configuration and integration include security event audit and event logging for TLS systems. In all cases, there are variations of services and offerings, which each collaborator describes in their respective sections.

### 2.1 Product Installation Sequence (Example Build)

Figure 2-1 shows the dependencies among components deployed for the example build. A solid line with a single arrow signifies hard dependencies. The component from which the arrow points should be installed before the component to which the arrow points. This facilitates phased and secure deployment. A dashed line with a double arrow indicates that integration between the components is not dependent on the installation sequence (i.e., either component can be installed first).

Figure 2-1 Overview of Dependencies Among Components Deployed for the Example Build



## 2.2 Thales TCT Luna SA 1700 Hardware Security Module

HSMs are specialized hardware devices dedicated to maintaining the security of sensitive data throughout its life cycle. HSMs provide tamper-evident and intrusion-resistant protection of critical keys and other secrets, and off-loading of processing-intensive cryptographic operations. By performing cryptographic operations within the HSM, sensitive data never leaves the secure confines of the hardened device.

The Thales TCT Luna SA for Government is a network-attached HSM with multiple partitions to effectively provide a many-in-one solution to multiple tenants—each with its own security officer management credentials. Depending on security needs, the Luna SA can be used with or without a secure personal identification number entry device (PED) for controlling management access to the HSM partitions. Utilizing the PED takes the HSM from a Federal Information Processing Standards (FIPS) 140-2 [1] Level 2 certified device to Level 3. The Luna SA also comes in two performance models: the lower performance 1700, and the high-performance 7000 for transaction-intensive use cases.

## 2.2.1 Day 0: Product Installation and Standard Configuration

### 2.2.1.1 Prerequisites

#### 2.2.1.1.1 Rack Space

Installation of the HSM requires rack space with the following characteristics:

- standard 1u 1 gin rack mount chassis
- dimensions: 19" x 21" x 1.725" (482.6 millimeters [mm] x 533.4 mm x 43.815 mm)
- weight capacity: 28 pounds (lb) (12.7 kilograms [kg])
- input voltage: 100-240 V.50-60 hertz
- power consumption: 180 watts (W) maximum, 155 W typical
- temperature: operating 0 degrees Celsius (C)–35 degrees C, storage 20 degrees C–60 degrees C
- relative humidity: 5% to 95% (38 degrees C) noncondensing

#### 2.2.1.1.2 Networking

One of two approaches to networking may be used. The steps for the commands in this document assume the NCCoE's laboratory networking environment will be replicated. An organization may also opt to use its own network settings. In either case, the following Luna SA HSM appliance parameters information will be needed:

- IP address that will be assigned to this device (Static IP is recommended)
- Host name for the HSM appliance (registered with network DNS)
- a domain name where the device will reside
- default gateway IP address
- DNS Name Server IP address(es)
- Search Domain name(s)
- device subnet mask
- Ethernet device (use eth0, which is the uppermost network jack on the HSM appliance back panel, closest to the power supply, and labeled 1 )

The network must be configured for optimal use of Luna appliances. The following bandwidth and latency recommendations are optimal for performance settings:

- bandwidth
  - minimum supported: 10 megabit (Mb) half-duplex

- recommended: at least 100 Mb full duplex—full gigabit Ethernet is supported

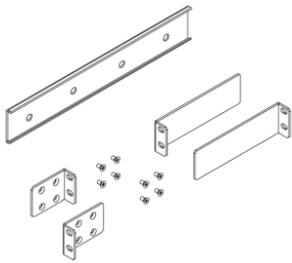
Note: Ensure the network switch is set to AUTO negotiation, as the Luna appliance negotiates at AUTO. If the network switch is set to use other than automatic negotiation, there is a risk that the switch and the Luna appliance will settle on a much slower speed than is actually possible in the organization’s network conditions.

- network latency
  - maximum supported: 500 milliseconds (ms)
  - recommended: 0.5 ms

### 2.2.1.1.3 Unpacking the Appliance

Follow this checklist to verify that all of items required for the installation are in hand.

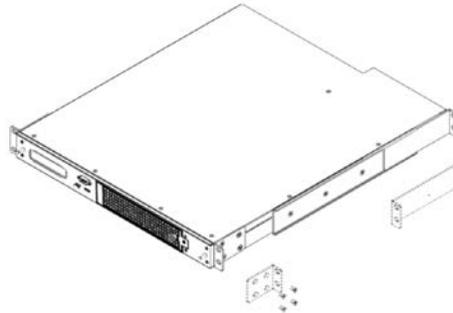
| Qty | Item                                                                                                                                                                                                                  |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1   |  <p data-bbox="691 1121 997 1157">Luna SA HSM appliance</p>                                                                        |
| 2   |  <p data-bbox="315 1562 1377 1633">power supply cord (one for each power supply; style to suit country for which was ordered)</p> |

| Qty | Item                                                                                                                                                                                                                                                                                                   |
|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1   |  <p data-bbox="690 682 998 714">null modem serial cable</p>                                                                                                                                                          |
| 1   |  <p data-bbox="544 1008 1161 1039">Universal Serial Bus 2.0 to RS232 serial adapter</p>                                                                                                                               |
| 1   |  <p data-bbox="267 1386 1347 1543">Set of:<br/>- 2 front mounting brackets with screws<br/>- 2 side bracket guides<br/>- 2 sliding rear brackets (Fit into the guides for rear support adjustable positioning.)</p> |
| 1   |                                                                                                                                                                                                                     |

| Qty | Item                                           |
|-----|------------------------------------------------|
|     | client/software development kit (SDK) software |

### 2.2.1.2 Rack-Mount the Appliance

1. Install and adjust rails and brackets to suit the equipment rack.

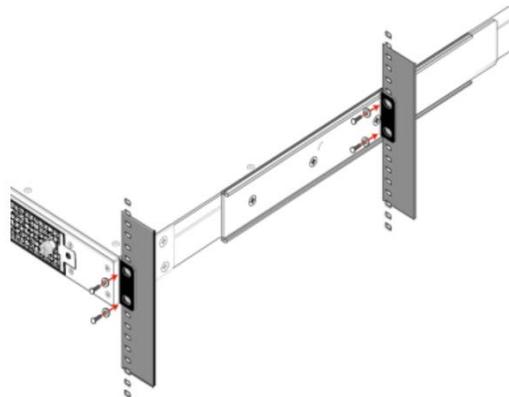


2. Mount the appliance in the equipment rack. Alternatively, ignore the rails and mounting tabs, and rest the Luna SA appliance on a mounting tray or shelf suitable for the organization's specific style and brand of equipment rack.

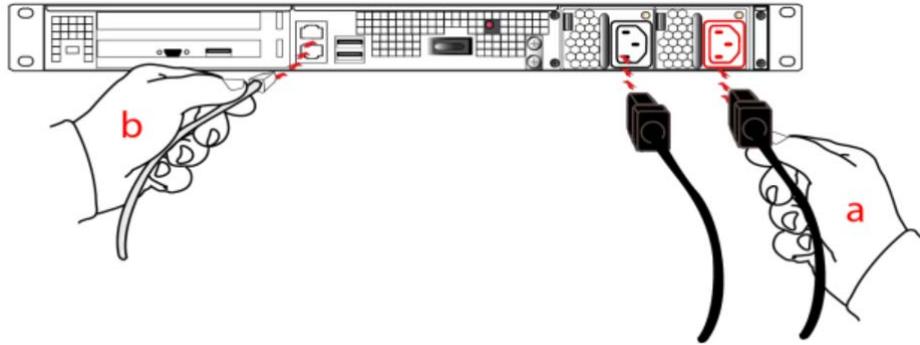
---

**CAUTION:** Support the weight of the appliance until all four brackets are secured.

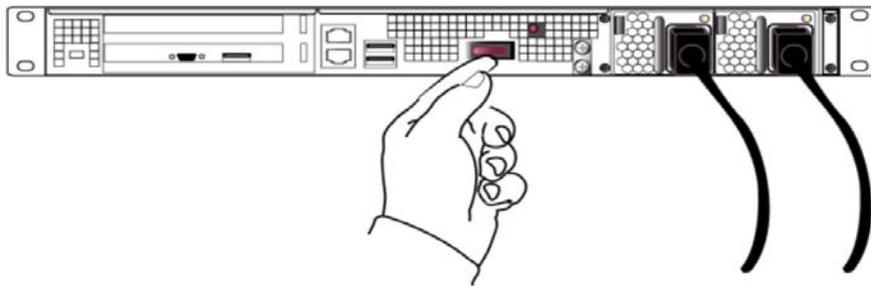
---



3. Insert the power (a) and network (b) cables at the rear panel. For proper redundancy and best reliability, the power cables should connect to two completely independent power sources.



4. Press and release the Start/Stop switch, on the rear panel.



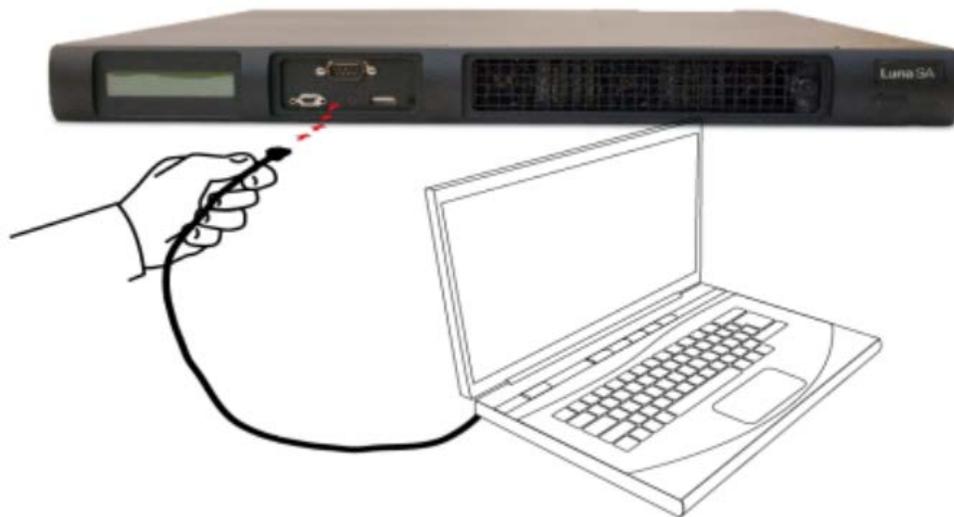
### 2.2.1.3 Initial Appliance Configuration

This section describes the process to prepare the new HSM Server and one client system for operation with the application. It includes the following steps:

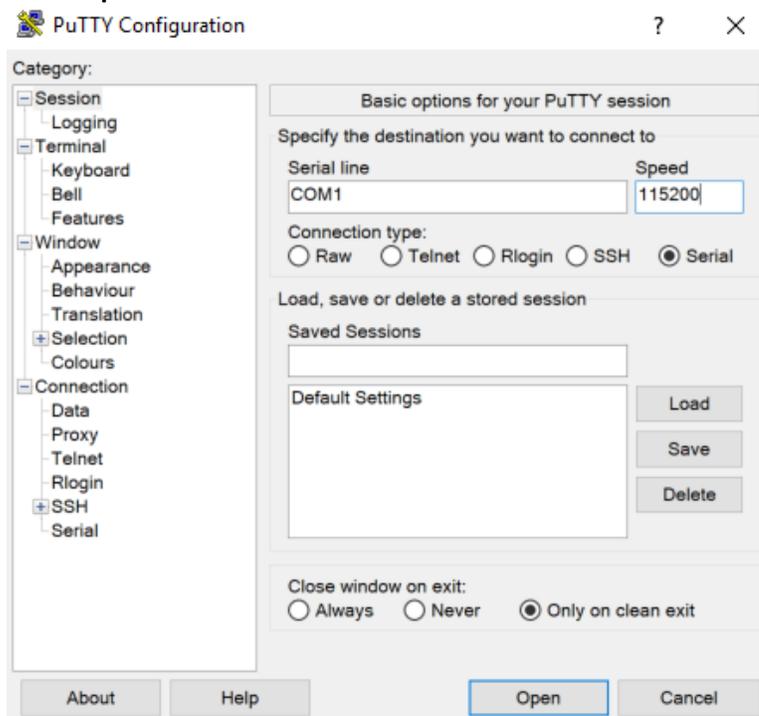
- process for first-time login and changing passwords
- verify and set the date and time
- configure HSM appliance's IP and network parameters (using static or Dynamic Host Configuration Protocol [DHCP]. In general, we strongly recommend against using DHCP for HSM appliances.)
- make network connections (To make a network connection, refer to Section [2.2.1.1.2.](#))
- HSM initialization process
- restart services so configuration changes can take effect

#### 2.2.1.3.1 Process for First-Time Login and Changing Passwords

1. To perform initial login to the HSM appliance, connect a serial cable to serial port on the front of the appliance.

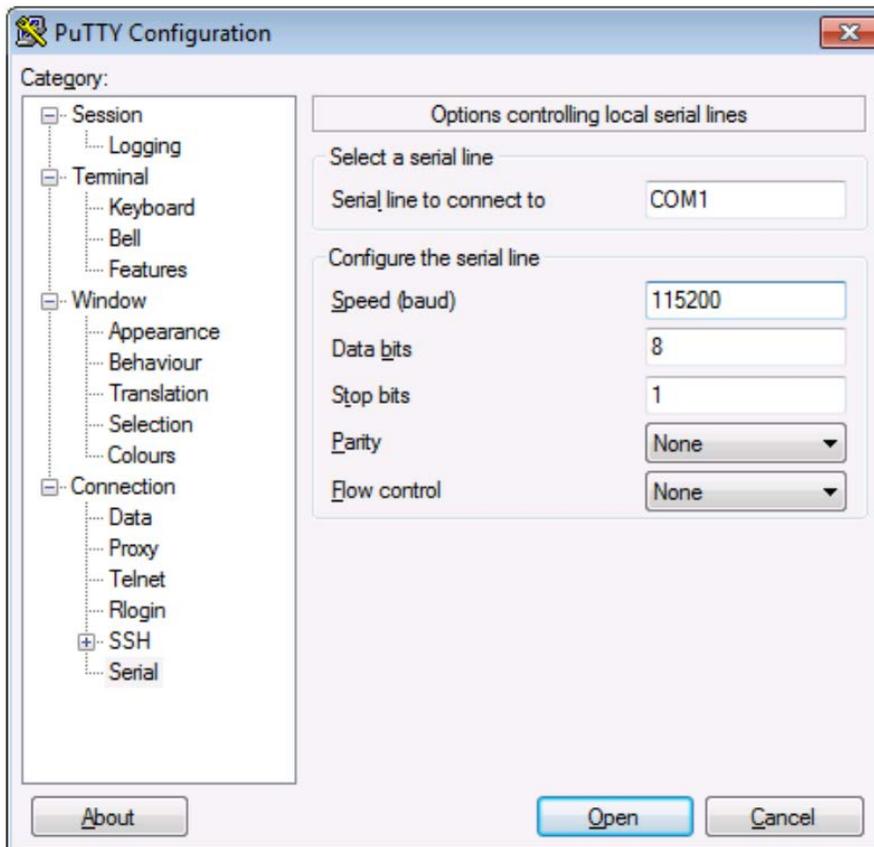


2. On the management laptop, open the PuTTY application and select a **Connection type** of **Serial** with a **Speed** of **115200**.



3. Navigate to the **Serial** Category on the bottom left side of the window.
4. Configure the serial connection to support the SSL Visibility Appliance's console speeds by selecting the following options:

- **Speed (baud):** 115200
- **Data bits:** 8
- **Stop bits:** 1
- **Parity:** None
- **Flow control:** None



5. Log in to the appliance by using the default credentials of:
  - **username:** bootstrap
  - **password:** bootstrap
6. For security purposes, the user is immediately prompted to change the factory-default password for the admin account.

[localhost] ttyS0 login: admin

Password:

You are required to change your password immediately (root enforced)

## Changing password for admin

(current) UNIX password:

```
A valid password should be a mix of upper and lower case letters, digits, and
other characters. You can use an 8 character long
password with characters from at least 3 of these 4 classes.
An upper case letter that begins the password and a digit that
ends it do not count towards the number of character classes used.
```

```
Enter new password:
```

```
Re-type new password:
```

```
Luna SA 5.4.0-14 Command Line Shell - Copyright (c) 2001-2013 SafeNet, Inc. All
rights reserved.
```

```
Command Result: 0 (Success)
```

```
lunash:>
```

The above represents a local serial connection; text will differ slightly for a Secure Shell (SSH) connection.

**Note:** The username and passwords are case-sensitive.

**Note:** To protect the HSM appliance and its HSM from vulnerabilities due to weak passwords, new passwords must be at least eight characters in length and must include characters from at least three of the following four groups:

- lowercase alphabetic (abcd...xyz)
- uppercase alphabetic (ABCD...XYZ)
- numeric (0123456789)
- special (nonalphanumeric, #\*@\$%&...)

**Note:** Login must occur within two minutes of opening an administration session, or the connection will time out.

### 2.2.1.3.2 Date and Time

To configure the HSM's date and time, perform the following steps:

1. Verify the current date and time on the HSM Server.
2. At the lunash prompt, type the command:  

```
lunash:> status date
```
3. If the date, time, or time zone is incorrect for the location, change them by using the `lunash sysconf` command. For example: 

```
lunash:> sysconf timezone set Canada/Eastern
Timezone set to Canada/Eastern
```

4. Use `sysconf time` to set the system time and date <HH:MM YYYYMMDD> in the format shown. Note that the time is set on a 24-hour clock (00:00 to 23:59).

```
lunash:> sysconf time 12:55 20190410 Sun April 10 12:55:00 EDT 2019
```

5. Optionally to configure Network Time Protocol (NTP), use the following command:

```
lunash:> sysconf ntp addserver 192.168.1.12
```

6. Activate the NTP service with the following command:

```
sysconf ntp enable
```

### 2.2.1.3.3 Network Configuration

1. Use the `network show` command to display the current settings and to see how they need to be modified for the network.

```
lunash:>net show
```

```
Hostname: HSM
Domain: int-nccoe.org

IP Address (eth0): 192.168.1.13
HW Address (eth0): 00:15:B2:AB:D6:D6
Mask (eth0): 255.255.255.0
Gateway (eth0): 192.168.1.1
```

```
Name Servers: 192.168.1.6
Search Domain(s): <not set>
```

```
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
Link status
eth0: Configured
 Link detected: yes
eth1: Configured
 Link detected: no
```

```
Command Result : 0 (Success)
lunash:>
```

2. Use `network hostname` to set the host name of the HSM appliance (use lowercase characters).  
lunash:> **network hostname HSM**
3. Use `network domain` to set the name of the network domain in which the HSM Server (appliance) is to operate.  
lunash:> **net domain int-nccoe.org**
4. Use `network dns add nameserver` to set the Nameserver IP Address (address for the local name server).  
lunash:> **net dns add nameserver 192.168.1.6**
5. Use `net dns add searchdomain` to set the DNS Search Domain (the search list to be used for host name lookups).  
lunash:> **net dns add searchdomain int-nccoe.org**

6. Use `network interface` to change network configuration settings.

All of the `network interface` parameters are required for the IP setup of the Ethernet device and must be set at the same time for the HSM appliance to connect with the network.

```
[HSM] lunash:>net interface -device eth0 -ip 192.168.1.13 -netmask 255.255.255.0 -
gateway 192.168.1.1
```

7. View the new network settings with `network show`.  
lunash:> `network show`

#### 2.2.1.3.4 Generate a New HSM Server Certificate

Although the HSM appliance came with a server certificate, good security practice dictates that a new one be generated.

1. Use `sysconf regenCert` to generate a new server certificate:

```
lunash:> sysconf regenCert 192.168.1.13
WARNING !! This command will overwrite the current server certificate and private
key.
All clients will have to add this server again with this new certificate.
If you are sure that you wish to proceed, then type 'proceed', otherwise type
'quit'
> proceed
Proceeding...
'sysconf regenCert' successful. NTLS must be (re)started before clients can
connect.
Please use the 'ntls show' command to ensure that NTLS is bound to an appropriate
network device or IP address/hostname for the network device(s) NTLS should be
active on. Use 'ntls bind' to change this binding if necessary.

Command Result: 0 (Success)
lunash:>
```

#### 2.2.1.3.5 Bind the Network Trust Link Service

From the factory, the network trust link service (NTLS) is bound to the loop-back device by default. To use the appliance on the network, bind the NTLS to one of the two Ethernet ports— ETH0 or ETH1—or to a host name or IP address. Use the `ntls show` command to see current status.

1. Use `ntls bind` to bind the service:

```
lunash:>ntls bind eth0 -bind 192.168.1.13
Success: NTLS binding hostname or IP Address 192.168.1.13 set.
NOTICE: The NTLS service must be restarted for new settings to take effect.
If you are sure that you wish to restart NTLS, then type 'proceed', otherwise
type 'quit'
> proceed
Proceeding...
Restarting NTLS service...
Stopping ntlsl: [OK]
Starting ntlsl: [OK]
```

```
Command Result : 0 (Success)
[myluna] lunash:>ntls show
NTLS bound to network device: eth0 IP Address: "192.168.1.13" (eth0)
Command Result : 0 (Success)
```

---

**NOTE:** The “Stopping ntl” operation might fail in the above example, because NTLS is not yet running on a new HSM appliance—ignore this message. The service restarts regardless if the stop was needed.

---

### 2.2.1.3.6 Enabling Federal Information Processing Standards 140-2 Mode

In many areas of the information security industry, validations against independent or government standards are considered a desirable or essential attribute of a product. NIST’s FIPS 140 is the pre-eminent standard in the field of cryptography. Enabling FIPS 140-2 [1] ensures the HSM uses strong cryptographic modules in its operations.

1. Log in to the APPLIANCE management console (LunaSH) as admin.
  - a. SSH into the APPLIANCE
  - b. Use these credentials: Username: admin Password: \*\*\*\*YOUR admin PASSWORD\*\*\*\*
2. Check if FIPS 140 mode is enabled.
  - a. Command: `hsm show`
  - b. In the results, look for “The HSM is in FIPS 140-2 approved operation mode.” If this is seen, then stop: FIPS 140-2 mode is already enabled on the HSM. Otherwise, continue.
3. Log in to the admin role.
  - a. Command: `hsm login`
  - b. Password: \*\*\*\*YOUR admin PASSWORD\*\*\*\*
4. View HSM Capabilities and Policies.
  - a. Command: `hsm showPolicies`
  - b. In the results, look for “Allow non-FIPS algorithms” and record its value and code.
5. Edit HSM Capabilities and Policies.
  - a. Command: `hsm changePolicy -policy <code> -value <desired_value>`
    - i. `hsm changePolicy -policy 12 -value 1`
    - ii. When prompted type: `proceed`
6. Confirm FIPS 140 mode is enabled.
  - a. Command: `hsm show`
  - b. In the results, look for “The HSM is in FIPS 140-2 approved operation mode.” If this is seen, then stop: FIPS 140-2 mode is already enabled on the HSM. Otherwise, further investigation is required.

### 2.2.1.4 HSM Initialization

In this section, initialize the HSM portion of the Luna appliance and set any required policies. In normal operations, these actions are performed when first commissioning the Luna appliance.

#### 2.2.1.4.1 Initialize a Password-Authenticated HSM

1. To initialize the HSM, type the following command:

```
hsm -init -label HSM
```

```
[HSM] lunash:> hsm -init -label HSM
> Please enter a password for the security officer
> *****
Please re-enter password to confirm:
> *****
Please enter the cloning domain to use for initializing this
HSM (press <enter> to use the default domain):
> *****
Please re-enter domain to confirm:
> *****
CAUTION: Are you sure you wish to re-initialize this HSM?
All partitions and data will be erased.
Type 'proceed' to initialize the HSM, or 'quit'
to quit now.
>proceed
'hsm - init' successful.
```

2. When activity is complete, lunash displays a “success” message.

## 2.2.2 Day 1: Product Integration Configuration

### 2.2.2.1 Prerequisites

- NTL—This step will need to be completed for each system; refer to Section 2.2.2.2.
- ADCS—Windows server needs to be running; refer to guide.
- IIS—Windows server needs to be running; refer to guide.
- Venafi—must be installed and configured; refer to Section 2.2.2.2.

### 2.2.2.2 Network Trust Link

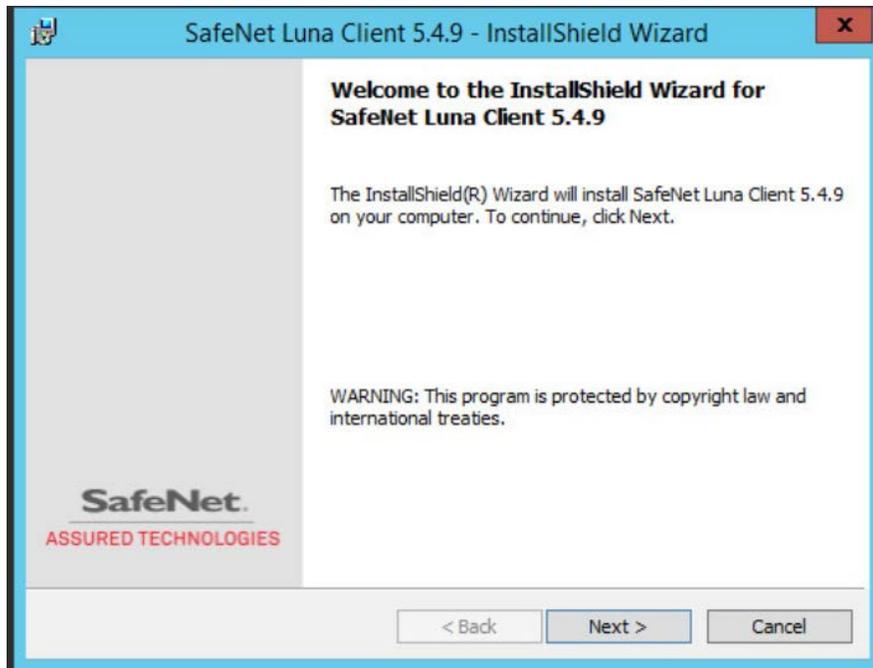
This section provides directions to configure a Luna Client to communicate with the network-attached Luna SA HSM. A client may have multiple Luna SA HSMs connected—using a slot designation when referencing an assigned Luna SA. The client also assumes the Luna SA is installed and operational but without a partition created for the new client.

The Luna Client is available in Windows and Linux. For Linux systems, refer to Thales TCT’s Configuring a Network Trust Link documentation. In this document, the necessary commands and screenshots are listed for Windows-based systems.

#### 2.2.2.2.1 Install the Luna Client Software

To install the Luna Client software, perform the following steps:

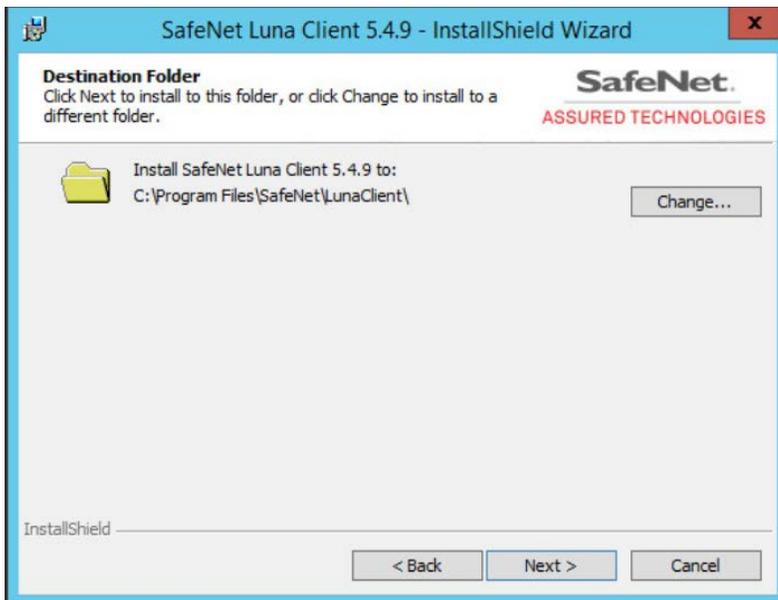
1. Log in to Windows as Administrator or as a user with administrator privileges.
2. Insert the Luna Client Software DVD into the optical drive.
3. Open a file explorer and navigate to **D:\windows\64\**.
4. Double-click **Luna Client.msi**.
5. Click **Next** at the welcome screen.



6. Accept the software license agreement by clicking **"I accept the terms in the license agreement"** and clicking **Next**.



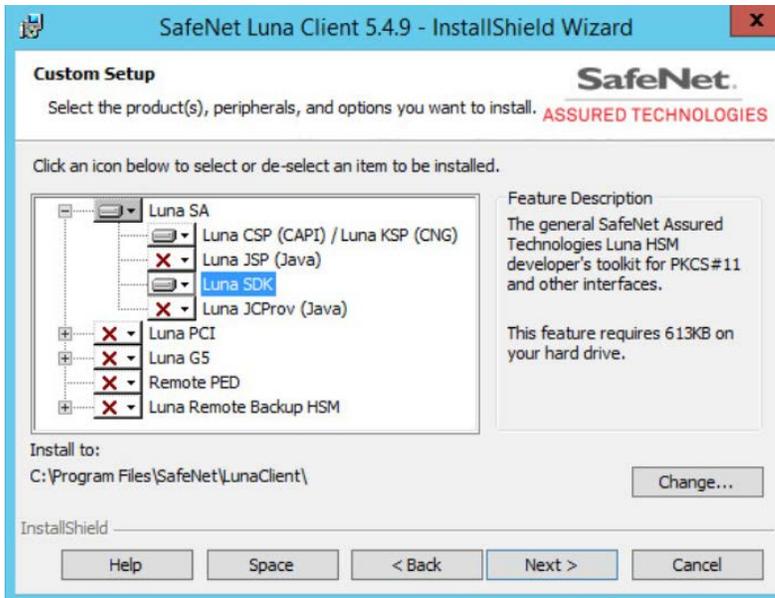
7. In the Choose Destination Location dialogue, accept the default offered and click **Next**.



8. Ensure the following options are selected and click **Next**:

- **Luna CSP (CAPI)/Luna KSP (CNG)**

- **Luna SDK**



9. On the **Ready to Install** page, click **Install**.
10. If Windows presents a security notice asking if the user wishes to install the device driver from Thales TCT, click **Install** to accept.



11. When the installation completes, click **Finish**.

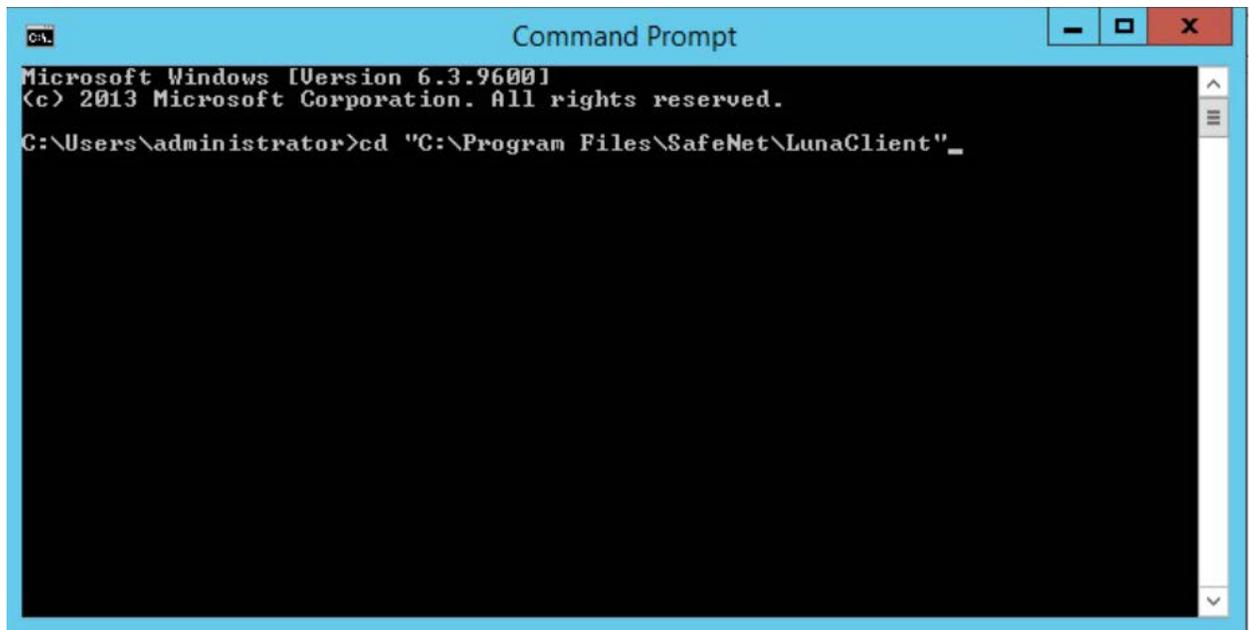
#### 2.2.2.2.2 Configure the Luna Client

To establish the NTL, first create a client certificate, and then the client and server certificates are exchanged. The Luna SA appliance is then added as a trusted server in the client.

### 2.2.2.2.3 Create the Client Certificate

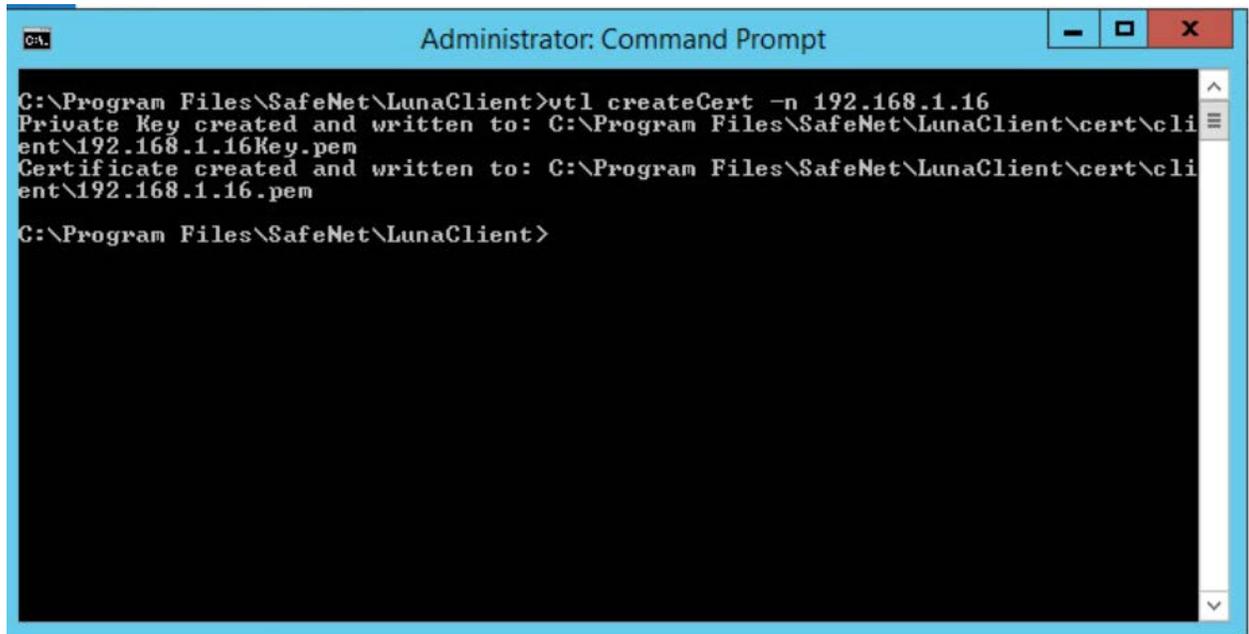
First, create the client certificate by using the Thales TCT VTL command line. This results in a *.pem* certificate file being created in a `\cert\client` subfolder.

1. On the client system, from the Windows command environment, run as administrator and navigate to the folder `C:\Program Files\Safenet\LunaClient`.



2. Enter the following command:

```
vtl createcert -n <client IP address>
```



The screenshot shows a Windows Command Prompt window titled "Administrator: Command Prompt". The window has a blue title bar and standard Windows window controls (minimize, maximize, close) in the top right corner. The command prompt shows the following text:

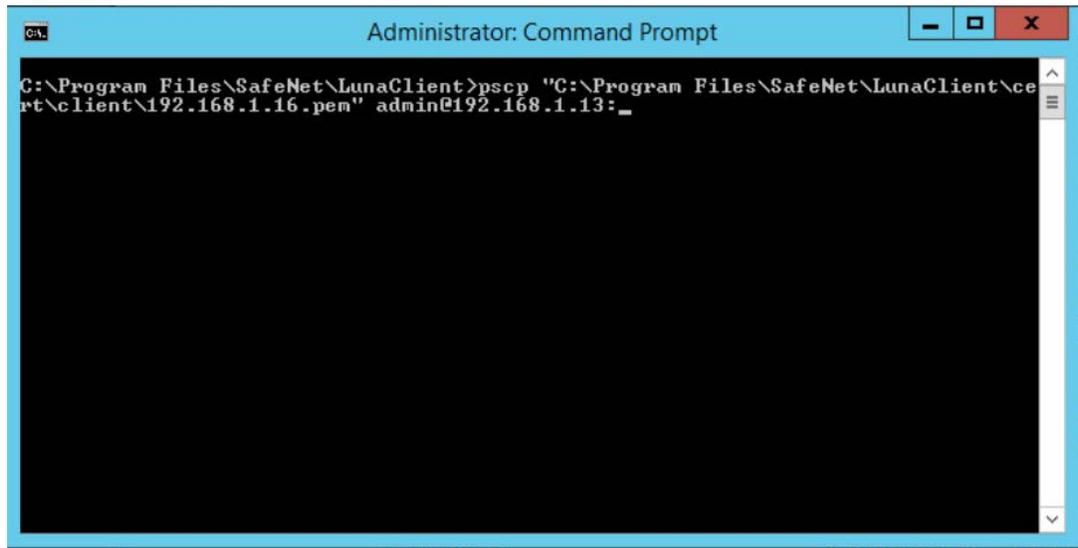
```
C:\Program Files\SafeNet\LunaClient>vtl createCert -n 192.168.1.16
Private Key created and written to: C:\Program Files\SafeNet\LunaClient\cert\cli
ent\192.168.1.16Key.pem
Certificate created and written to: C:\Program Files\SafeNet\LunaClient\cert\cli
ent\192.168.1.16.pem
C:\Program Files\SafeNet\LunaClient>
```

#### 2.2.2.2.4 Transfer the Client Certificate to the Luna SA

Now, transfer the newly created client certificate to the Luna SA by using the PuTTY Secure Copy Protocol (PSCP) or Secure Copy Protocol (SCP) tool.

1. On the client system using Windows, enter the following command:

```
pscp "C:\Program Files\SafeNet\LunaClient\cert\client\192.168.1.16.pem"
admin@192.168.1.13:
```



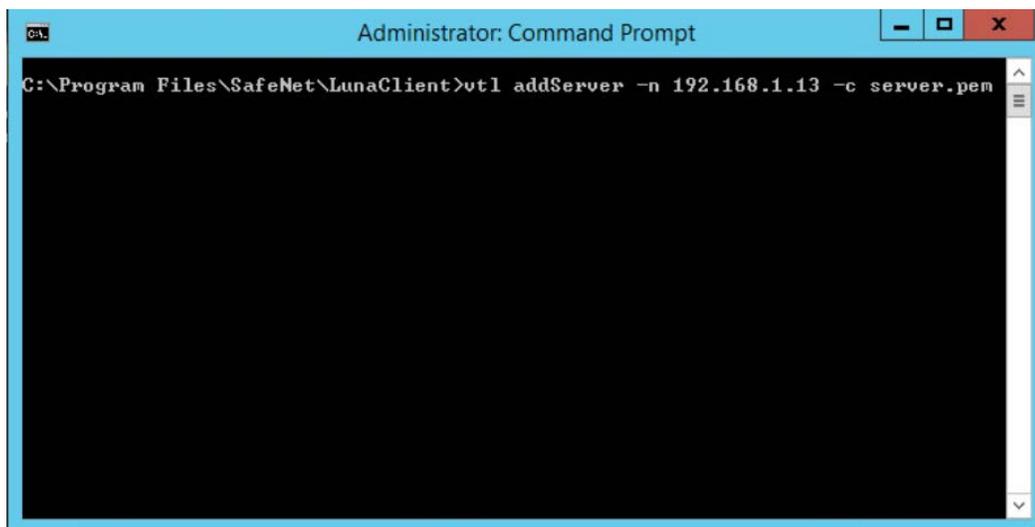
2. When prompted, enter the appliance administrative password for the Luna SA. The transfer automatically takes place.

#### 2.2.2.2.5 Transfer the Server Certificate from the Luna SA

Using PSCP or SCP, transfer the Luna SA's server certificate to the client.

1. On a client system using Windows, enter the following command:

```
pscp admin@192.168.1.13:server.pem
```



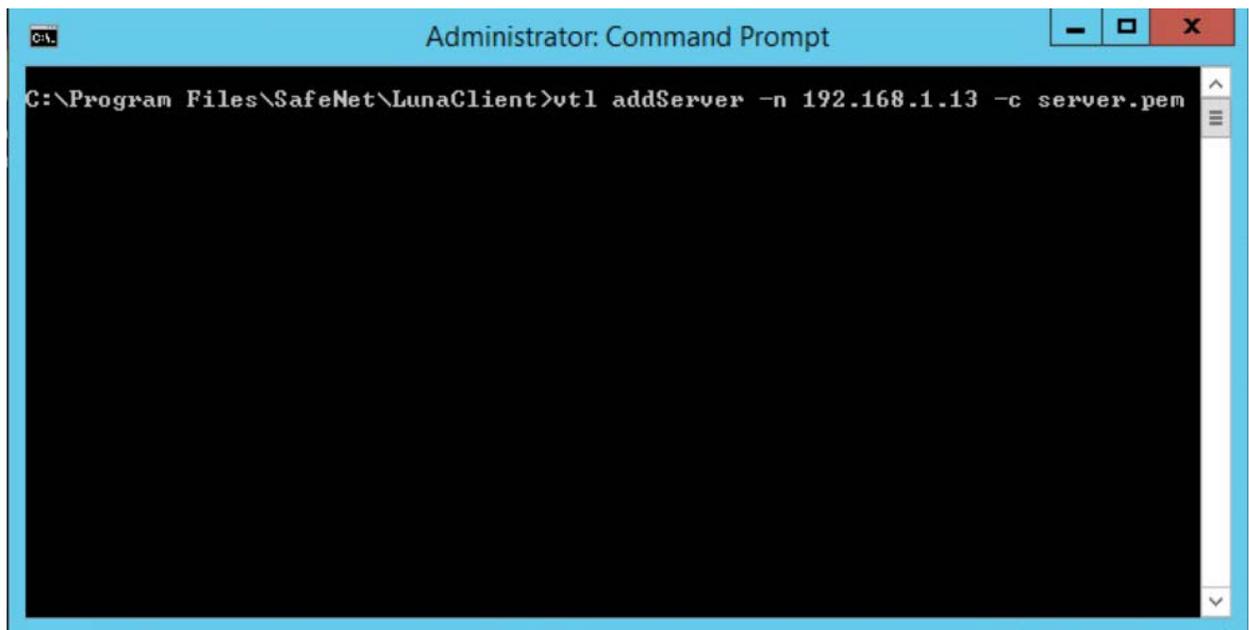
2. When prompted, enter the administrative password for the Luna SA. The transfer will automatically take place.

#### 2.2.2.2.6 Register the HSM on the Client

The final step in configuring the client is to register the Luna SA's certificate with the client.

1. On a client system, enter the following command:

```
vtl addServer -n <HSM IP Address> -c server.pem
```



At this point, the client is fully configured and ready to establish a secure link with the HSM.

#### 2.2.2.2.7 Create a Partition (Password Authentication)

1. Connect into the HSM via SSH or Serial.
2. At the `lunash:>` prompt on the Luna SA, enter the following command:

```
partition create -partition <partition name> -domain <domain name>
```

```

[HSM] lunash:>partition create -partition HRhsmiis

Please ensure that you have purchased licenses for at least this number of partitions: 5

Please enter a password for the partition:
> *****

Please re-enter password to confirm:
> *****

Please enter a cloning domain to use when creating this partition:
> *****

Please re-enter cloning domain to confirm:
> *****

If you are sure to continue then type 'proceed', otherwise type 'quit'
> proceed
Proceeding...

'partition create' successful.

```

3. When prompted, enter and re-enter to confirm the partition password.
4. Enter `proceed` when prompted.

#### 2.2.2.2.8 Register the Client on the HSM and Assign It to a Partition

Register the client on the HSM and assign it to a partition. Because the HSM was previously created and the client certificate was transferred to it, the HSM can find the certificate file based on the IP address. Assign a name for the client for easy recognition.

1. On the Luna SA, enter the following command to register the client:

```
client register -client HRhsmiis -ip 192.168.1.16
```

```
[HSM] lunash:>client register -client HRhsmiis -ip 192.168.1.16
```

2. On the Luna SA, enter the following command to assign the client to the previously created partition.

```
client assignPartition -client <client name> -partition <partition name>
```

```
[HSM] lunash:>client assignPartition -client HRhsmiis -partition HRhsmiis_
```

3. On the Luna SA, enter the following command to verify the client is assigned to the proper partition.

```
client show -client <client name>
```

```
[HSM] lunash:>client show -client HRhsmiis

ClientID: HRhsmiis
IPAddress: 192.168.1.16
HTL Required: no
OTT Expiry: n/a
Partitions: "HRhsmiis"

Command Result : 0 (Success)
```

At this point, the HSM is configured, and in the next section, the user will return to the client to verify connectivity and the ability to request cryptographic operations from the client.

#### 2.2.2.2.9 Verify the Network Trust Link

Return to the client and verify it can view the Luna SA and its associated slot and partition. Run the Multitoken2 utility to verify the client can request cryptographic operations from the HSM.

#### 2.2.2.2.10 Verify the Luna SA in Client Server Lists

Verify the Luna SA is in the client's server lists.

1. On the client system, from the Windows command environment run as administrator, navigate to the folder *C:\Program Files\Safenet\LunaClient*.
2. On the client system, enter the following command and verify the Luna SA is in the list of servers:

```
vtl listservers
```

```
C:\Program Files\Safenet\LunaClient>vtl listservers
Server: 192.168.1.13 HTL required: no
```

#### 2.2.2.2.11 Verify the Slot and Partition

Verify the slot and the assigned HSM partition can be seen.

1. On the client system using either Windows and Linux, enter the following command to verify the Luna SA slot and partition are known to the client:

```
vtl verify
```

```
C:\Program Files\SafeNet\LunaClient>vtl verify
The following Luna SA Slots/Partitions were found:
Slot Serial # Label
==== ===== =====
1 575342049 HRhsmiis
C:\Program Files\SafeNet\LunaClient>_
```

Should this verification fail, check the times on the client and HSM to ensure they are set properly.

#### 2.2.2.2.12 Request Cryptographic Operations on the HSM

Request an actual crypto operation on the HSM to verify full functionality. The Multitoken utility to use is described in the Luna SA product documentation.

1. On the client system, enter the following command:

```
multitoken2 -mode rsasigver -key 1024 -slots 1,1,1,1,1
```

2. When prompted, if continuing, enter **y**.
3. Enter the partition password when prompted. The test will begin.
4. Press the **Enter** key to terminate the test after verifying that RSA signatures were successfully performed in the statistics table.

```

Command Prompt - multitoken2 -mode rsasigver -key 1024 -slots 1,1,1,1,1
C:\Program Files\SafeNet\LunaClient>multitoken2 -mode rsasigver -key 1024 -slots 1,1,1,1,1
Initializing library...Finished Initializing
...done.

Do you wish to continue?
Enter 'y' or 'n': y

Constructing thread objects.
Logging in to tokens...
slot 1... Enter password: NCC0e123456!
Serial Number 575342049

Please wait, creating test threads.
Test threads created successfully. Press ENTER to terminate testing.

RSA sign/verify 1024-bit : <packet size = 16 bytes>

1, 0 1, 4 | operations/second | elapsed
 | total average | time (secs)
-----|-----|-----
136.9 136.7 | 679.0 672.187* | 10_

```

### 2.2.2.3 ADCS Integration Configuration

This section provides the necessary steps for configuring an ADCS CA to use the Thales TCT Luna SA 1700 HSM for Government, to secure the CA's private key. This section assumes the Luna HSM client has been installed and configured, as detailed in Section [2.2.1](#).

Perform the following steps:

- Verify the Network Trust Link (NTL) between the Windows Server and the HSM.
- Register the Key Storage Provider (KSP) on the Windows Server.
- Add the CA role.
- Verify the private key for the CA was created on the HSM.

#### 2.2.2.3.1 Prerequisites

To configure Microsoft CA to use the Luna HSM, the following prerequisites must be met:

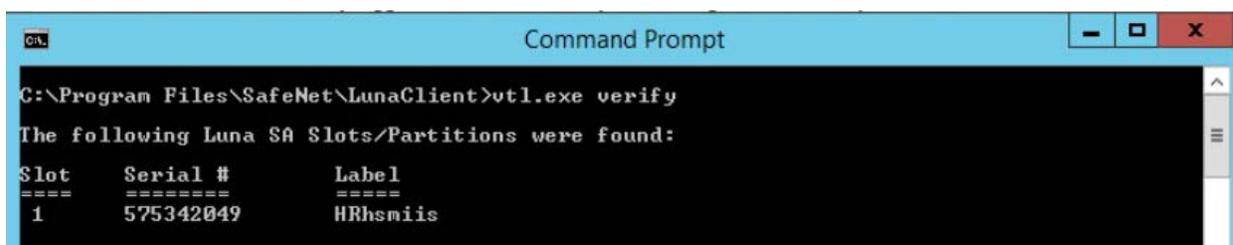
- The Thales TCT Luna HSM is installed and operational.
- The Thales TCT Luna Client is installed on the Windows Server where the CA is being added.

- The NTL is established between the Luna Client and the Luna HSM. If not, see Section [2.2.2.2](#).

### 2.2.2.3.2 Verify the HSM Configuration

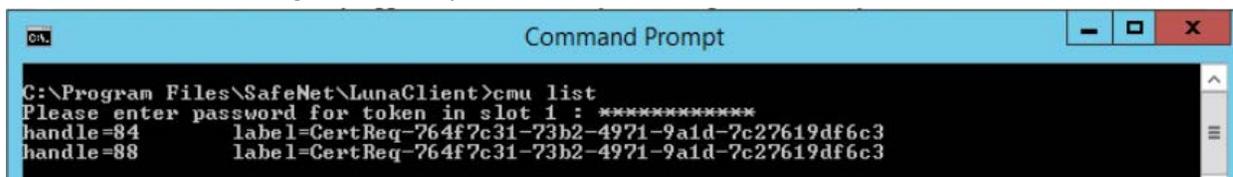
Verify the HSM client configuration prior to proceeding by following the steps below:

1. Open a Command Prompt as Administrator, and change into the Luna Client directory, typically `C:\Program Files\SafeNet\LunaClient\`.
2. Execute the command `VTL.exe verify` to check that the client is configured correctly and the partition is visible. Slot/Partition information should be displayed in response.



```
CA. Command Prompt
C:\Program Files\SafeNet\LunaClient>vtl.exe verify
The following Luna SA Slots/Partitions were found:
Slot Serial # Label
==== ===== =====
1 575342049 HRhsmiis
```

3. Execute the command `cmu list` to see the list of current objects on the HSM, and enter the password when prompted. If nothing has been created on the partition, this list will be blank. Once the CA is configured, the keys created on the HSM are listed.

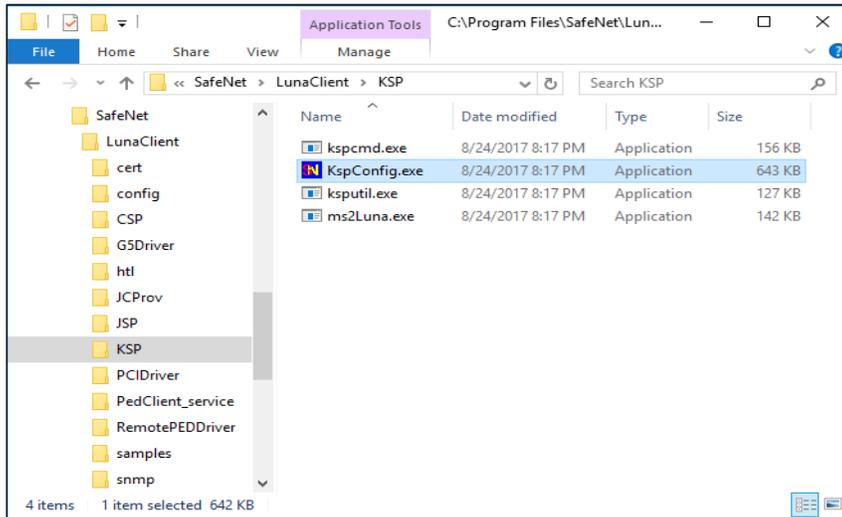


```
CA. Command Prompt
C:\Program Files\SafeNet\LunaClient>cmu list
Please enter password for token in slot 1 : *****
handle=84 label=CertReq-764f7c31-73b2-4971-9a1d-7c27619df6c3
handle=88 label=CertReq-764f7c31-73b2-4971-9a1d-7c27619df6c3
```

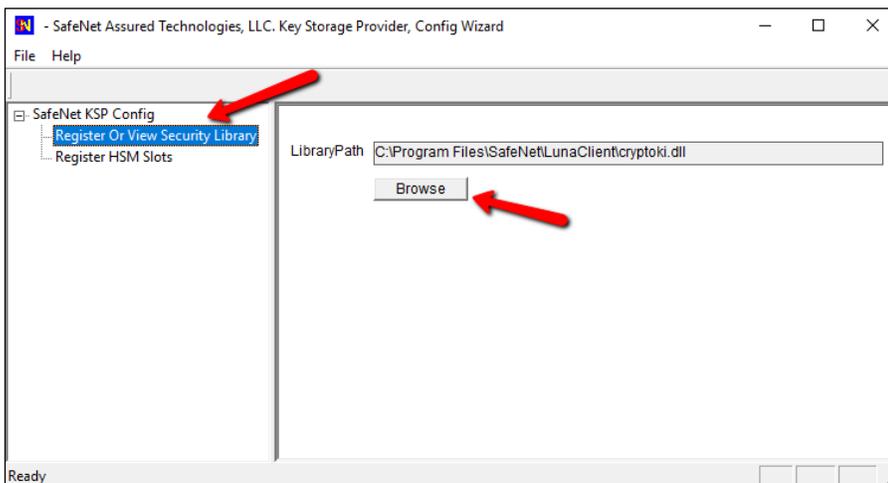
### 2.2.2.3.3 Register the Key Storage Provider

Beginning with Windows Server 2008, the older CryptoAPI CSP has been superseded by the newer CNGKSP. The Luna Client installation includes a utility to register the Thales TCT HSM for Government as a KSP for use in Windows applications. To register, follow these instructions:

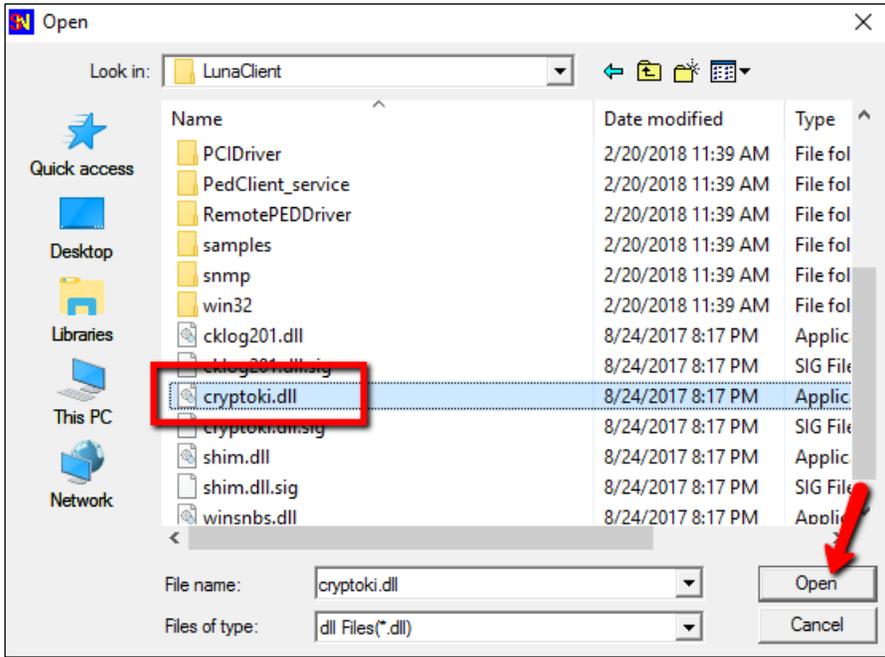
1. Open Windows Explorer, browse to the KSP folder in the Luna Client installation folder, and double-click on the **KSPConfig.exe** utility.



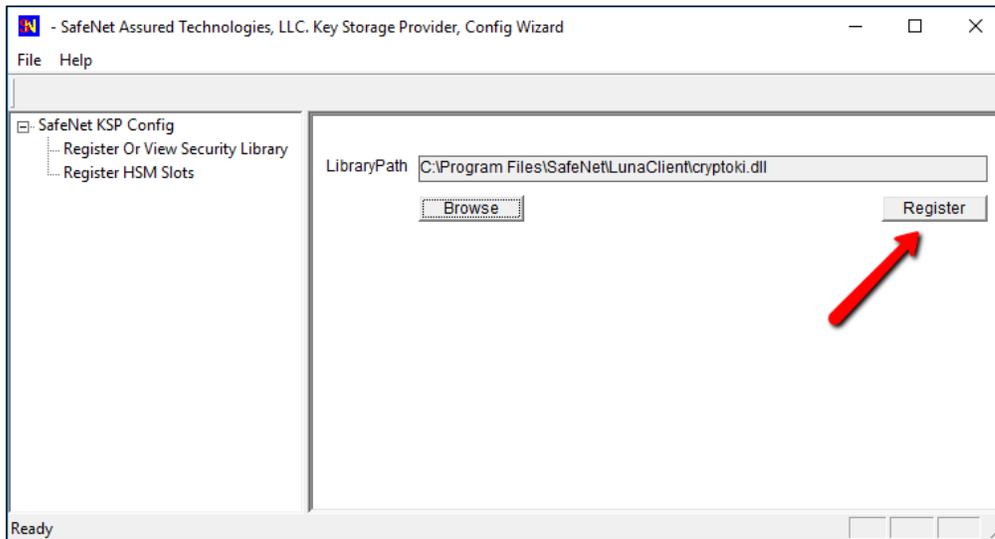
2. Double-click on **Register Or View Security Library**, then click **Browse**.



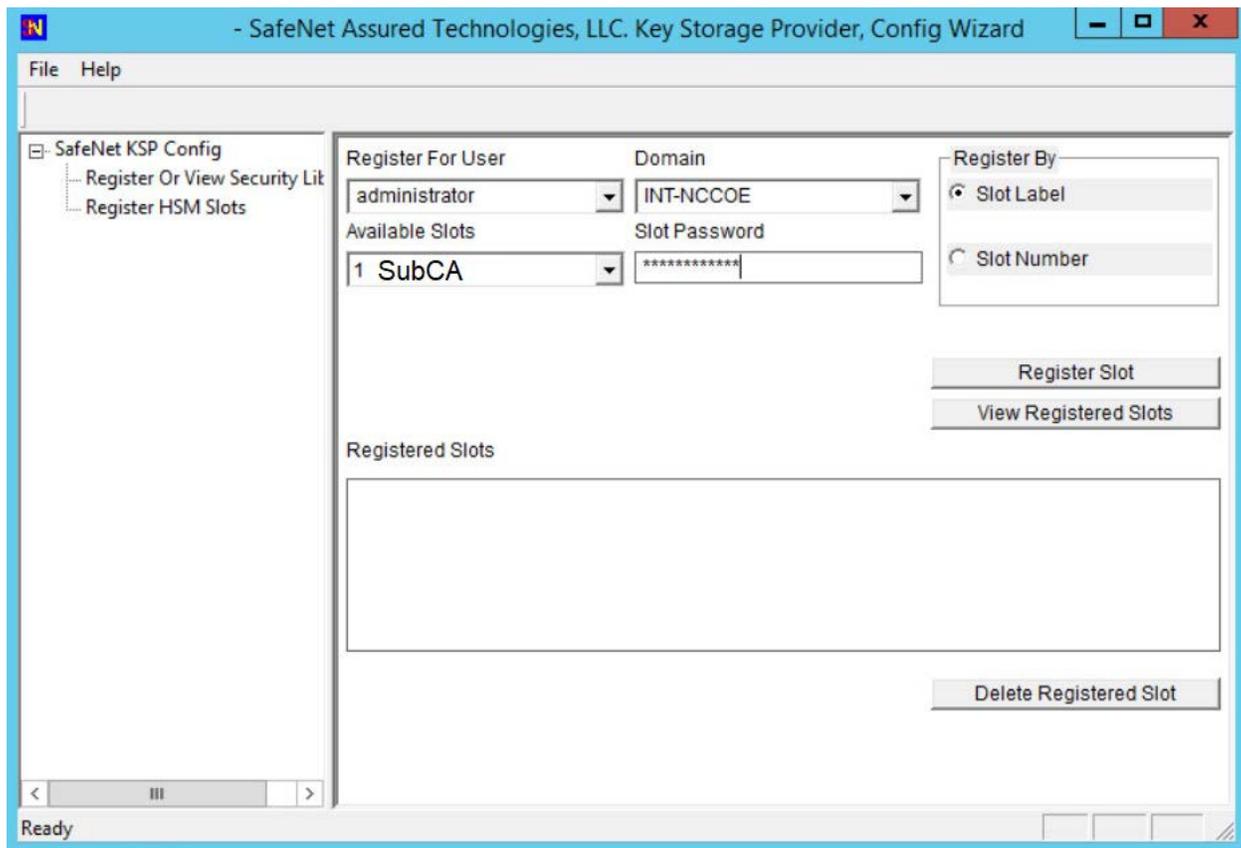
3. Browse to the Luna Client folder, select **cryptoki.dll**, and click **Open**.



4. Click on **Register** to complete the library registration.



5. Double-click **Register HSM Slots** on the left to open the slot registration page. Select the **Administrator** account and the Domain for the user that will be configuring the CA role. For a server joined to a domain, this should be a Domain or Enterprise Admin account rather than the local machine Administrator. Select the slot for the HSM, enter the **Slot Password**, and click **Register Slot**.



- Repeat the slot registration for the user **SYSTEM** with Domain **NT AUTHORITY**, and click **Register**. This is the account used for the CA service—it must also have access to the HSM. Verify the registration by selecting user and domain and clicking **View Registered Slots**.

#### 2.2.2.3.4 Add CA Role

For instructions on CA installation and configuration, refer to Section [1.5.3.3.2](#) on root CAs.

#### 2.2.2.3.5 Verify the Successful Integration on the HSM

As a final step, verify the private key and the public key are stored on the HSM.

- Open a command prompt and change to the Luna Client directory, typically C:\Program Files\SafeNet\LunaClient\.
- Run **cmu list** to verify the private and public keys for the CA are present on the HSM. They are represented by two “handles.”

The screenshot below shows running the `cmu list` command before configuring the CA and then after the configuration has been completed.

This completes integration of the Thales TCT Luna SA 1700 HSM for Government with Microsoft Active Directory Certificate Services.

#### 2.2.2.4 IIS Integration Configuration

This section provides the steps necessary to integrate the Microsoft IIS web server and the Thales TCT Luna SA 1700 HSM. The benefit of the integration is that the root private key for IIS is stored in a hardened, FIPS 140-2-certified device.

The following steps explain how to register the Thales TCT Luna SA 1700 HSM as a KSP to store the root certificate's private key in the HSM.

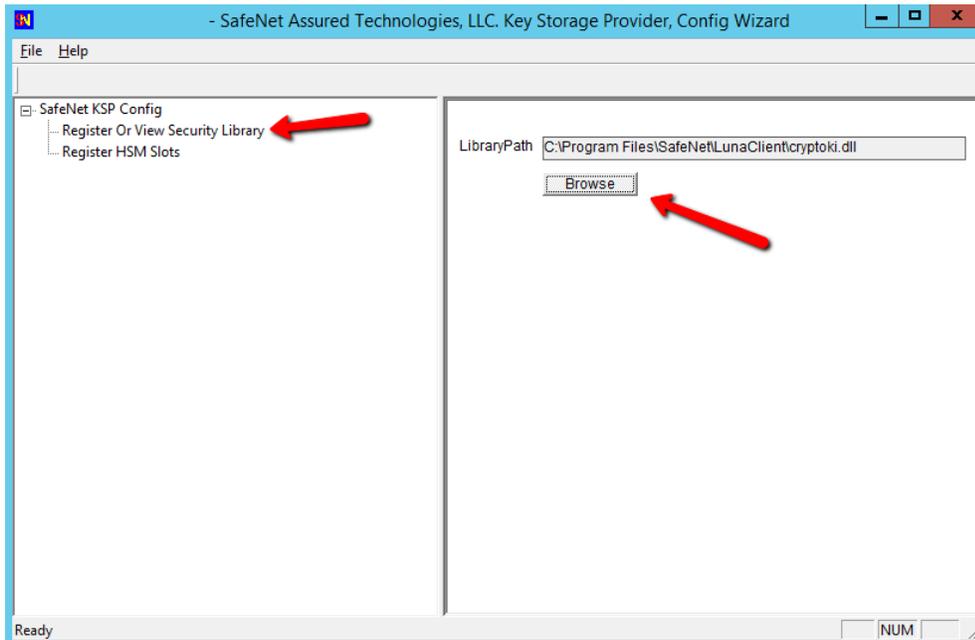
##### 2.2.2.4.1 Prerequisites

- IIS is installed or ready to be installed. The firewall rules may need to be edited to allow https access (typically port 443) and optionally block http (port 80).
- If mutual authentication is being performed, the trusted CA's certificate has been installed.

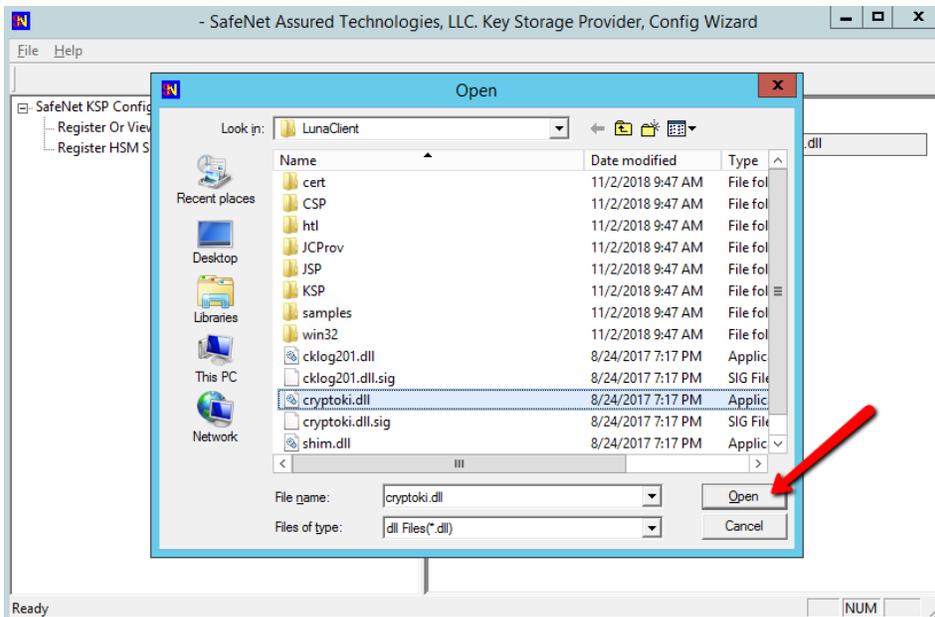
##### 2.2.2.4.2 Register the Luna KSP

For IIS integration, two accounts need access to the HSM. First, the `DOMAIN\Administrator` account is used for setting up the server—creating the certificate request and installing the certificate. Second, the `NT Authority\System` account is used by the server to start the IIS service. The **KSPConfig** utility is used to register the HSM as a KSP for these accounts.

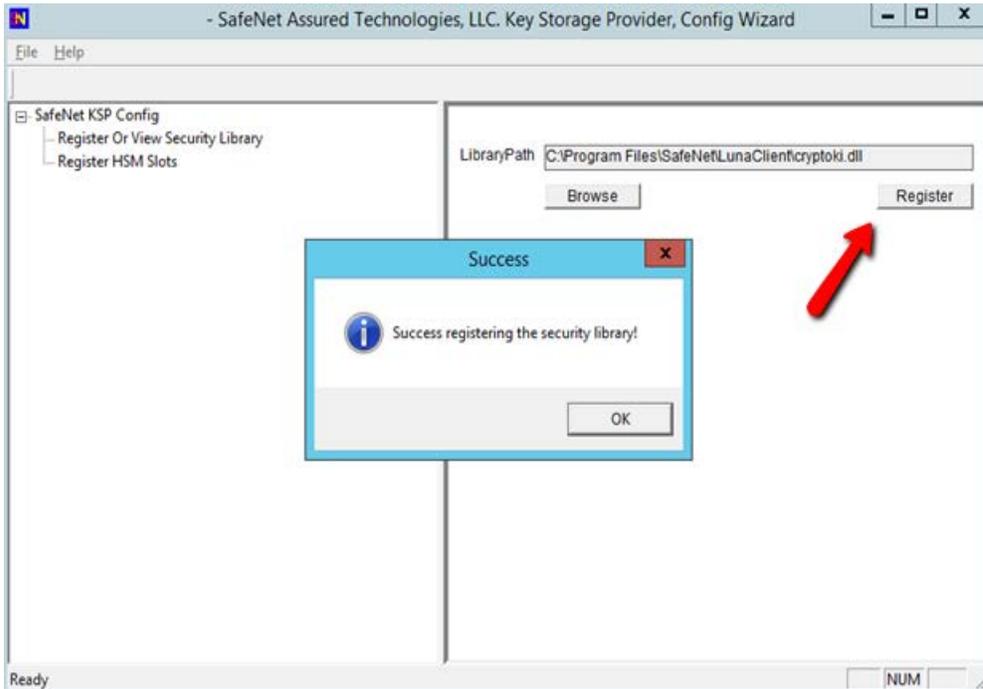
1. Navigate to the **KSP** directory under the Luna installation directory, which is typically `C:\ProgramFiles\SafeNet\LunaClient`.
2. Run **KspConfig.exe** to launch the wizard.
3. When the wizard launches, double-click **Register Or View Security Library** on the left side of the pane, and then click the **Browse** button on the right.



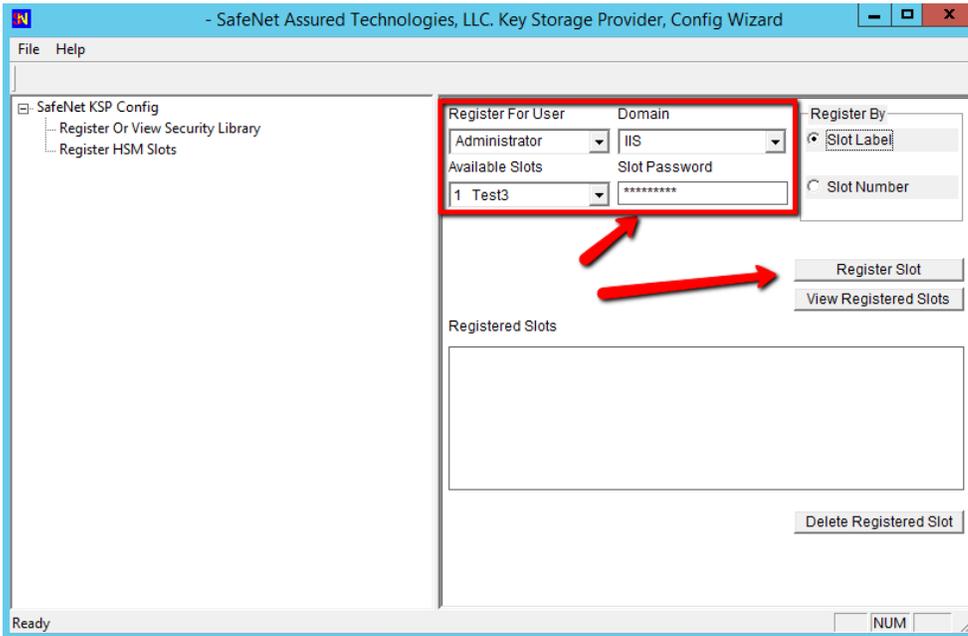
4. Browse to and select the **cryptoki.dll** library in the Luna Client directory.



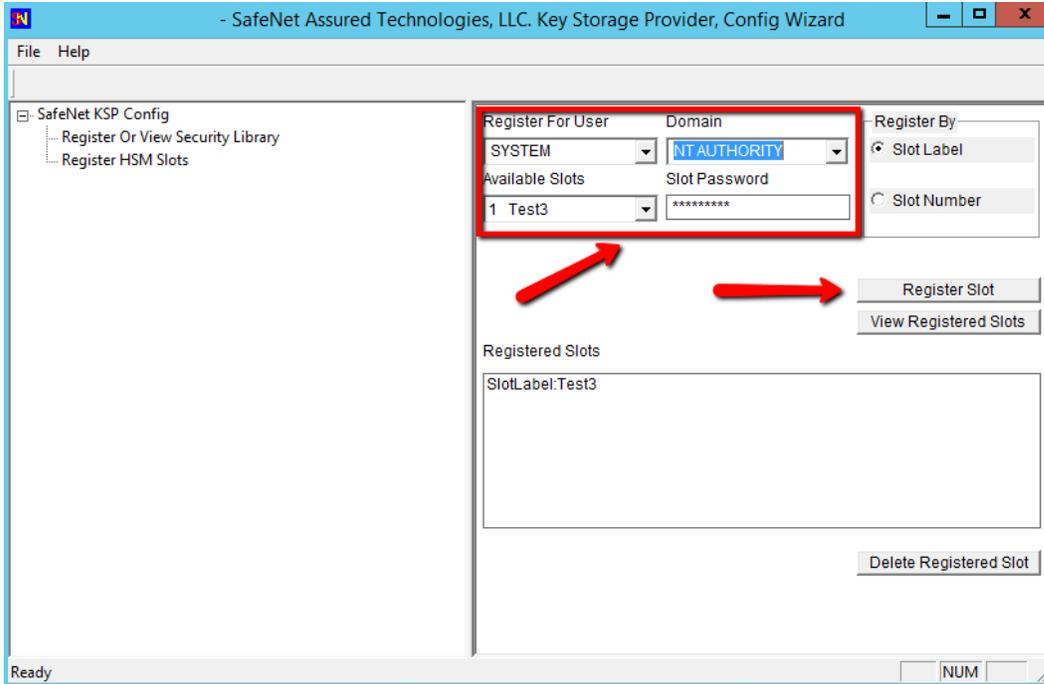
5. Having selected the dll, click the **Register** button. The message “**Success registering the security library!**” displays.



6. Double-click **Register HSM Slots** on the left side of the pane.
7. Verify the correct **User** and **Domain** are selected (the Administrator account on the server) and slot is selected (can be registered by slot label or slot number), and enter the **Slot Password** (HSM partition password).
8. Click **Register Slot** to register the slot for that User/Domain. Upon successful registration, a message **“The slot was successfully and securely registered”** displays.



9. Repeat the steps above to register the slot for the **User SYSTEM** and **Domain NT AUTHORITY**.



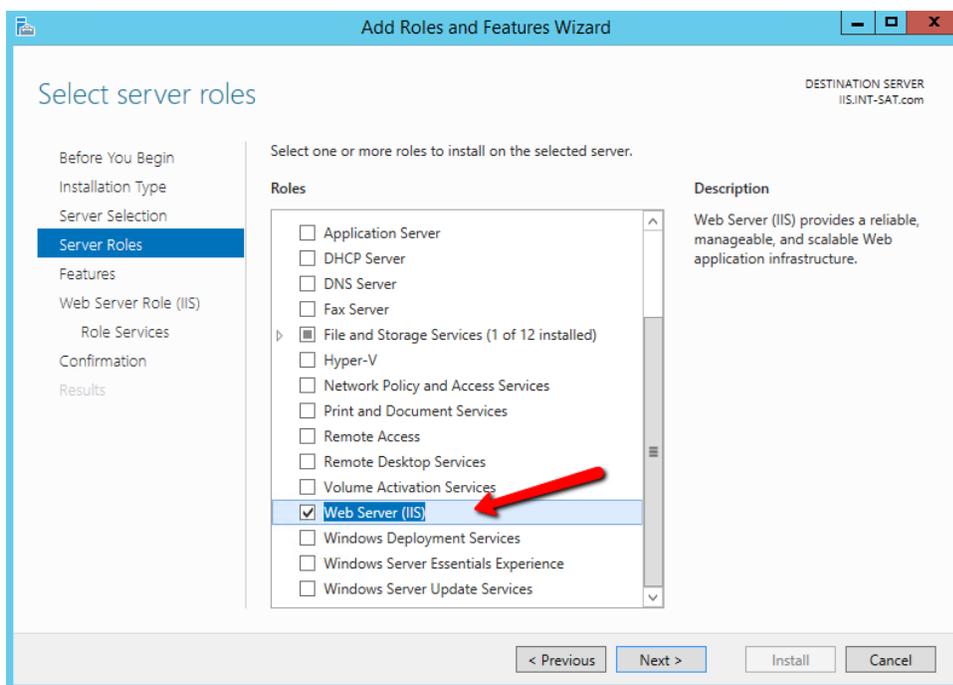
To verify the registered slot, select a **User/Domain**, and click the **View Registered Slots** button.

#### 2.2.2.4.3 Setup Synopsis

- Verify the NTL between the server and the HSM.
- Register the HSM as a KSP.
- Install IIS and configure it to use an HSM.
- Create a certificate request for IIS, and get it signed.
- Install the signed certificate.
- Bind the certificate to the web server.

#### 2.2.2.4.4 Install Microsoft IIS

The next step is to install the **Web Server (IIS)** role by using **Server Manager**. There are no special considerations surrounding the IIS integration with an HSM. Please follow the installation and configuration steps in Section [1.5.5.2](#).



#### 2.2.2.4.5 Create and Install a Certificate for IIS

IIS will need a certificate installed that has been signed by a trusted CA. This involves creating a certification signing request (CSR), then the CA signs it and installs it back in the server. **IIS Manager**

provides an easy way for creating a CSR, but it cannot be used when a key is generated on an external HSM. Instead, use a Microsoft command line utility.

Clients attempting to securely connect to the web server will see an alert if the fully qualified domain name (FQDN) in the Common Name (CN) field (or on more recent browsers, the FQDN in the Subject Alternate Name field) does not match the uniform resource locator (URL) they are accessing. An alert also occurs if the certificate was not issued by a trusted root CA. For this integration, use the FQDN in the CN and Subject Alternative Name (SAN) fields.

#### 2.2.2.4.6 Create a Certificate Signing Request and Private Key

Instructions follow for using the **certreq.exe** utility to create the CSR and private key in the HSM.

1. Create a file called ***request.inf*** that will contain the necessary information for the utility to create the CSR. The contents of the file are as follows—only those items in blue italics will vary per the organization’s environment and requirements. The **CN** in the subject and the **dns** name in the **SAN** extension must match the full host name that clients enter as the URL in a web browser.

Copying and pasting the text may insert line breaks or change quotation marks to smart (curly) quotation marks. Ensure that each entry is on a single line and that all quotation marks are standard, straight, and double.

In this document, some entries may appear with line breaks such as the **Subject=...** and **%szOID\_ENHANCED\_KEY\_USAGE...** lines, but they must be on a single line. In addition, if using Notepad, change the file type to “all files” so it does not create the file with an extension of .txt. The “hide extensions for known file types” option may need to be disabled in Windows Explorer to verify the file is an *.inf* file rather than a *.txt* file. The text of the *.inf* file follows, as well as an image of the how the file should look.

```
[Version]
 Signature= "$Windows NT$"

 [NewRequest]
 Subject = "C=US,CN=HRhsm.int-
nccoe.org,O=SafeNetAT,OU=TLSLAB,L=Gaithersburg,S=Maryland"
 HashAlgorithm = SHA256
 KeyAlgorithm = RSA
 KeyLength = 2048
 ProviderName = "Safenet Key Storage Provider"
 KeyUsage = 0xf0
 MachineKeySet = True
 [EnhancedKeyUsageExtension]
 OID=1.3.6.1.5.5.7.3.1

[Strings]

szOID_SUBJECT_ALT_NAME2 = "2.5.29.17"
 szOID_ENHANCED_KEY_USAGE = "2.5.29.37"
```

```
szOID_PKIX_KP_SERVER_AUTH = "1.3.6.1.5.5.7.3.1" szOID_PKIX_KP_CLIENT_AUTH =
"1.3.6.1.5.5.7.3.2"
```

```
[Extensions]
```

```
%szOID_SUBJECT_ALT_NAME2% = "{text}dns=HRhsm.int-nccoe.org"
%szOID_ENHANCED_KEY_USAGE% =
"{text}%szOID_PKIX_KP_SERVER_AUTH%, %szOID_PKIX_KP_CLIENT_AUTH%"
```

Example image of file with correct line breaks:

```
request - Notepad
File Edit Format View Help
[Version]
Signature= "$Windows NT$"

[NewRequest]
Subject = "C=US,CN=HRhsm.int-nccoe.org,O=SafeNetAT,OU=TLSLAB,L=Gaithersburg,S=Maryland"
HashAlgorithm = SHA256
KeyAlgorithm = RSA
KeyLength = 2048
ProviderName = "Safenet Key Storage Provider"
KeyUsage = 0xf0
MachineKeySet = True
[EnhancedKeyUsageExtension]
OID=1.3.6.1.5.5.7.3.1

[Strings]

szOID_SUBJECT_ALT_NAME2 = "2.5.29.17"
szOID_ENHANCED_KEY_USAGE = "2.5.29.37"
szOID_PKIX_KP_SERVER_AUTH = "1.3.6.1.5.5.7.3.1"
szOID_PKIX_KP_CLIENT_AUTH = "1.3.6.1.5.5.7.3.2"

[Extensions]
%szOID_SUBJECT_ALT_NAME2% = "{text}dns=HRhsm.int-nccoe.org"
%szOID_ENHANCED_KEY_USAGE% = "{text}%szOID_PKIX_KP_SERVER_AUTH%, %szOID_PKIX_KP_CLIENT_AUTH%"
```

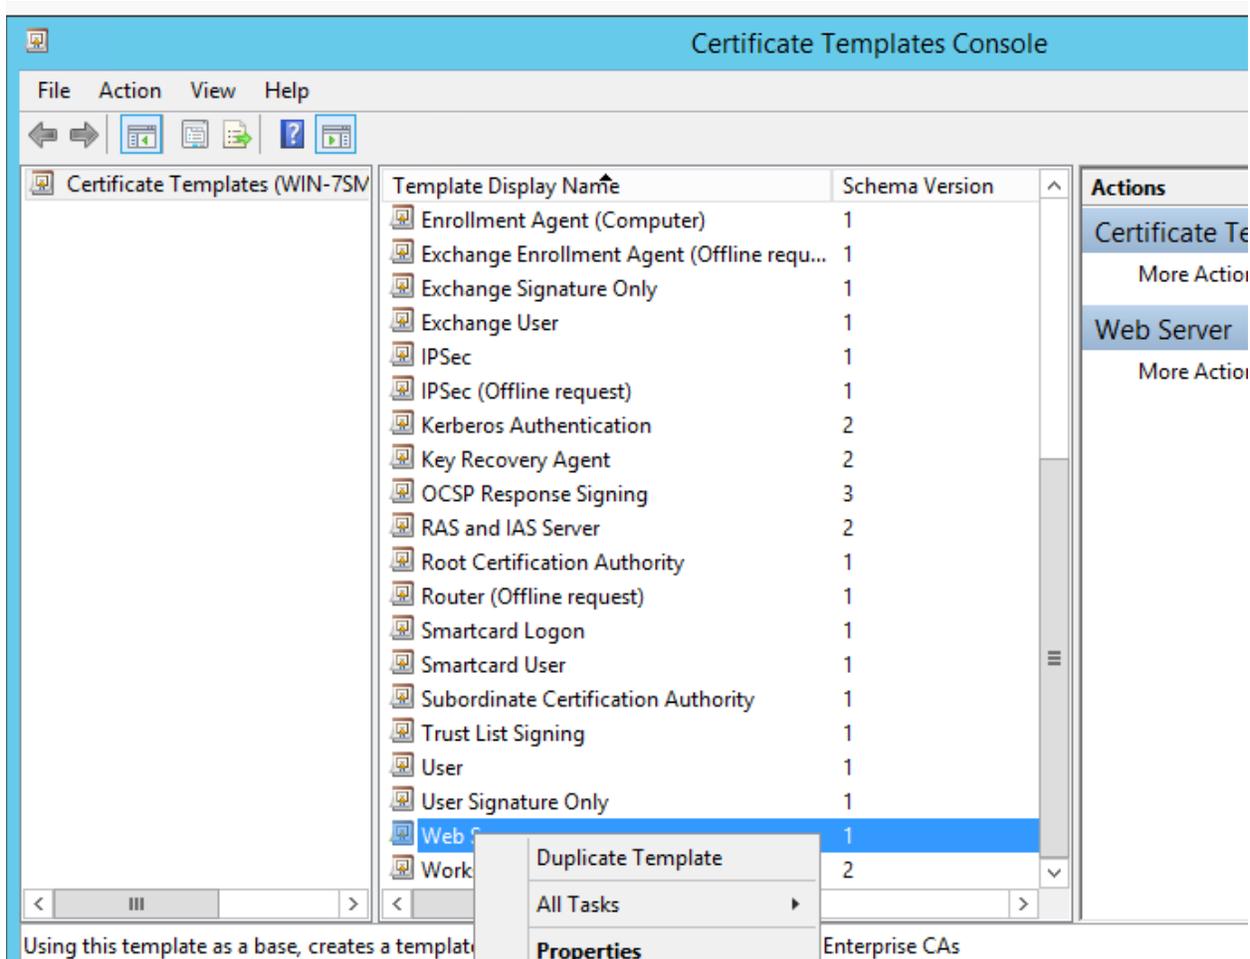
2. With the information file created, execute the **certreq** utility to generate a key on the HSM, and the certificate request. The CSR will be output to the file name that the user provides.

```
certreq.exe -new request.inf <CSR_filename>
```



The CSR was signed by using an Enterprise CA. Follow the steps below to create a new template and to sign the certificate request:

1. Search for and run **certsrv.msc**, or from Server Manager select **Tools > Certification Authority** to view the CA. Expand the CA > right-click **Certificate Templates** > select **Manage**.
2. In the **Certificate Templates Console**, scroll down to find the **Web Server** template and right-click > select **Duplicate Template**.



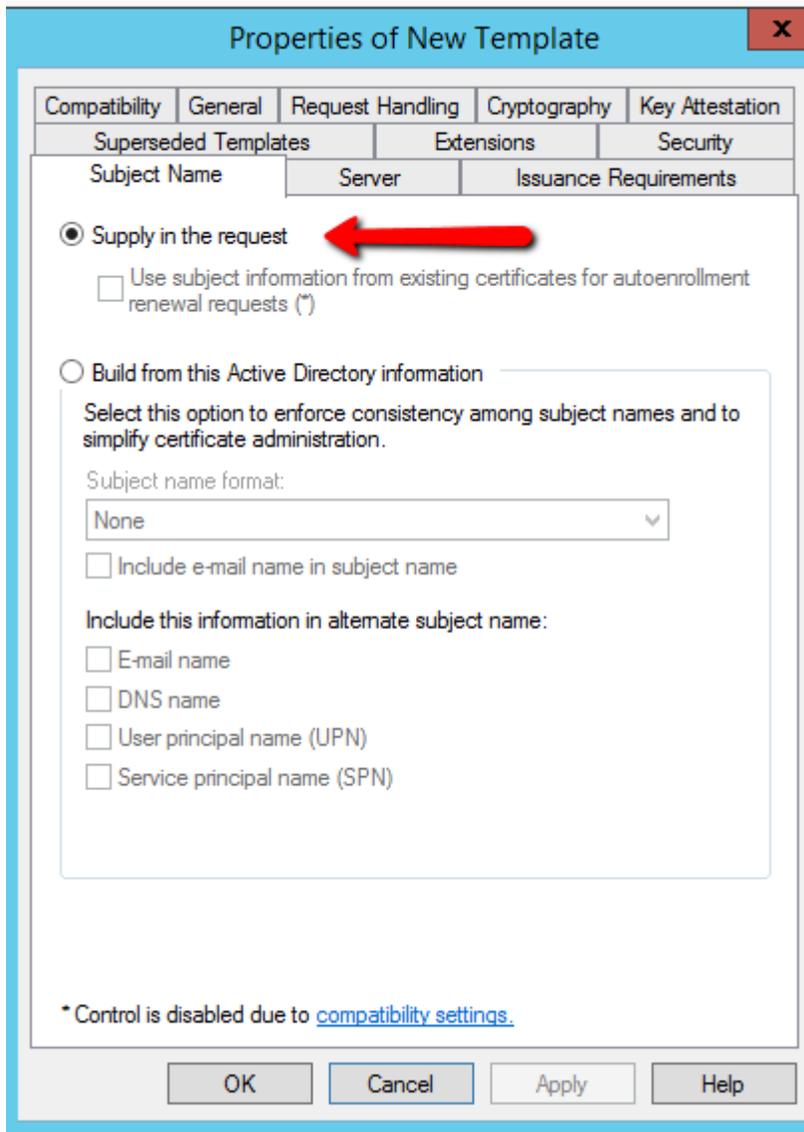
- Fill out the various sections of the properties with settings that adhere to the company's security policies. For this guide, the only thing altered is the **Template name** in the **General** tab. This will be the name used when signing the request on the command line.

The screenshot shows a dialog box titled "Properties of New Template" with a close button (X) in the top right corner. The dialog has several tabs: "Subject Name", "Server", "Issuance Requirements", "Superseded Templates", "Extensions", "Security", "Compatibility", "General", "Request Handling", "Cryptography", and "Key Attestation". The "General" tab is currently selected. Inside the dialog, there are the following fields and options:

- Template display name:** A text box containing "Copy of Web Server".
- Template name:** A text box containing "WebServer2".
- Validity period:** A dropdown menu showing "2" and "years".
- Renewal period:** A dropdown menu showing "6" and "weeks".
- Publish certificate in Active Directory
  - Do not automatically reenroll if a duplicate certificate exists in Active Directory

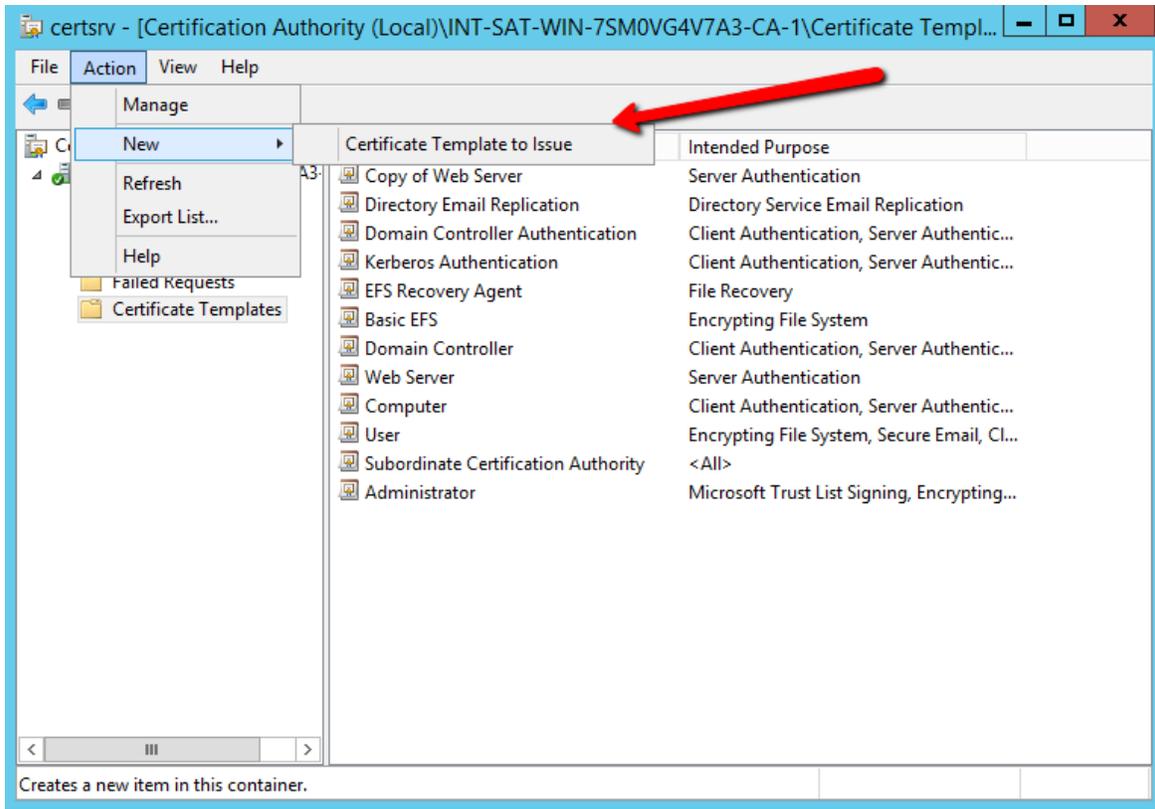
At the bottom of the dialog, there are four buttons: "OK", "Cancel", "Apply", and "Help".

4. Select the **Subject Name** tab, and verify that **Supply in the request** is selected. The FQDN is specified in both the CN and SAN fields in the request file created, and the certificate will use these values.

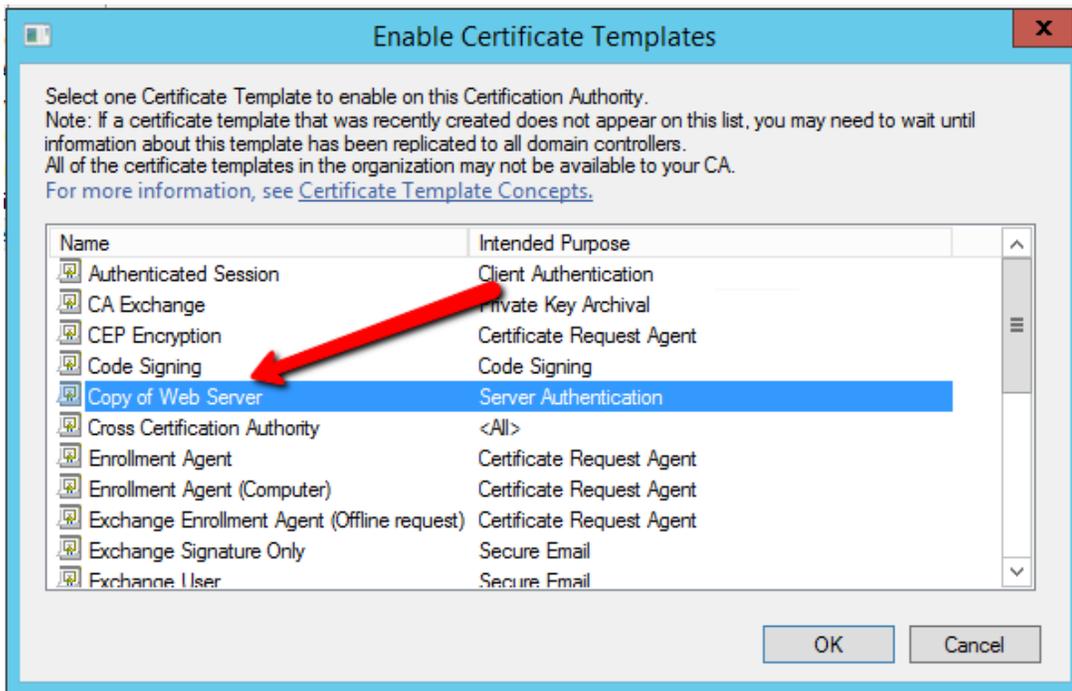


5. Click **OK** to finish creating the new template.
6. Close the **Certificate Templates Console** > return to the **Certificate Authority window**.

7. Click on **Action > New > Certificate Template to Issue**

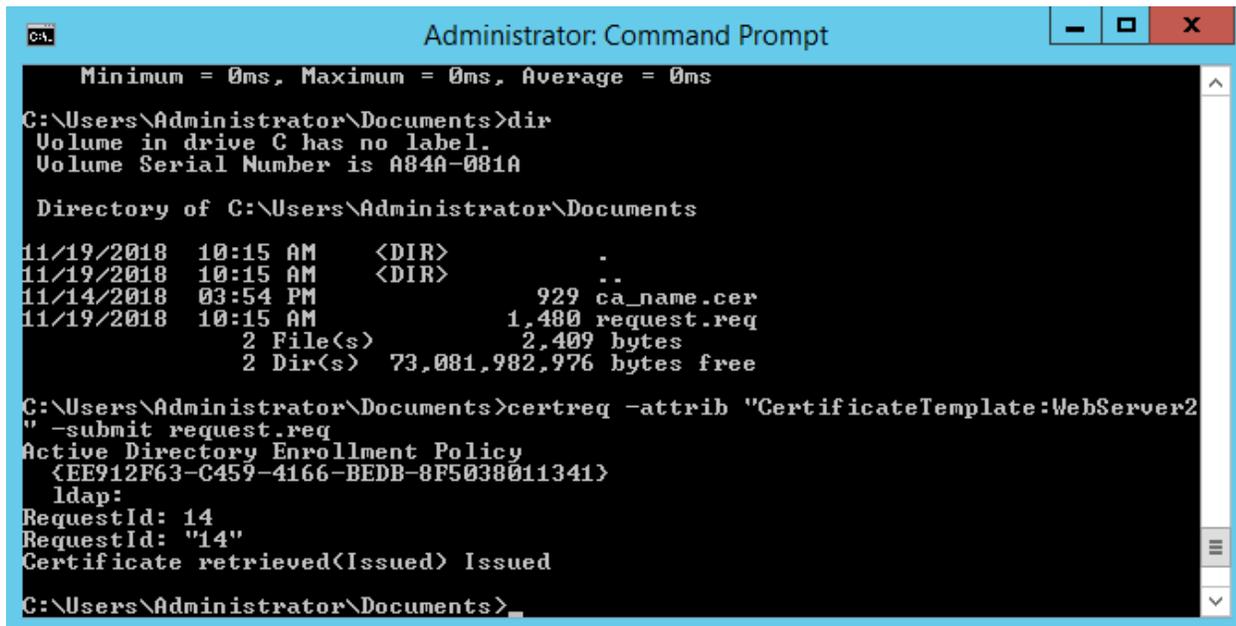


8. Select the certificate template created > click **OK**.



9. Generate a certificate from the certificate request:

```
certreq -attrib "CertificateTemplate:<TemplateName>" -submit <certificate request filename>
```

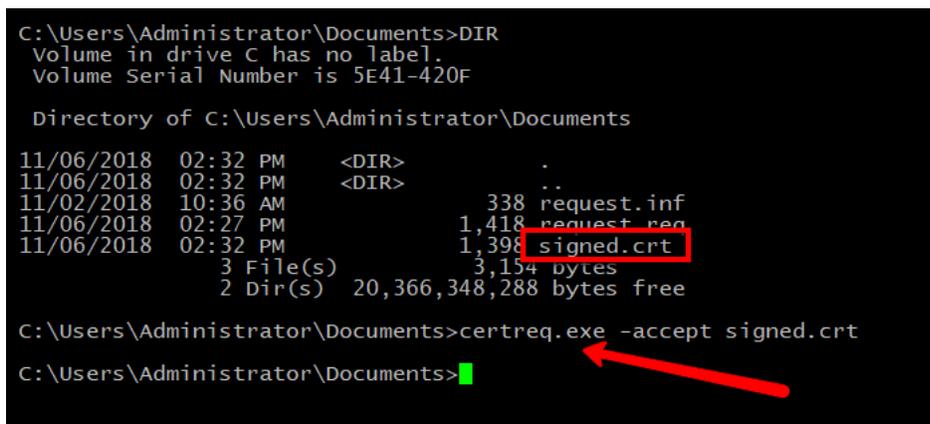


The user will be prompted to select the CA to use for signing, and a location and file name to save the signed certificate. Once the signed certificate file is created, it can be copied to the IIS server to continue with the integration.

#### 2.2.2.4.8 Install the Signed Certificate

Once the CSR is signed and the signed certificate file is received back, accept and install it by using the `certreq` utility.

```
certreq.exe -accept <newcert.crt>
```



```
C:\Users\Administrator\Documents>DIR
Volume in drive C has no label.
Volume Serial Number is 5E41-420F

Directory of C:\Users\Administrator\Documents

11/06/2018 02:32 PM <DIR> .
11/06/2018 02:32 PM <DIR> ..
11/02/2018 10:36 AM 338 request.inf
11/06/2018 02:27 PM 1,418 request.req
11/06/2018 02:32 PM 1,398 signed.crt
 3 File(s) 3,154 bytes
 2 Dir(s) 20,366,348,288 bytes free

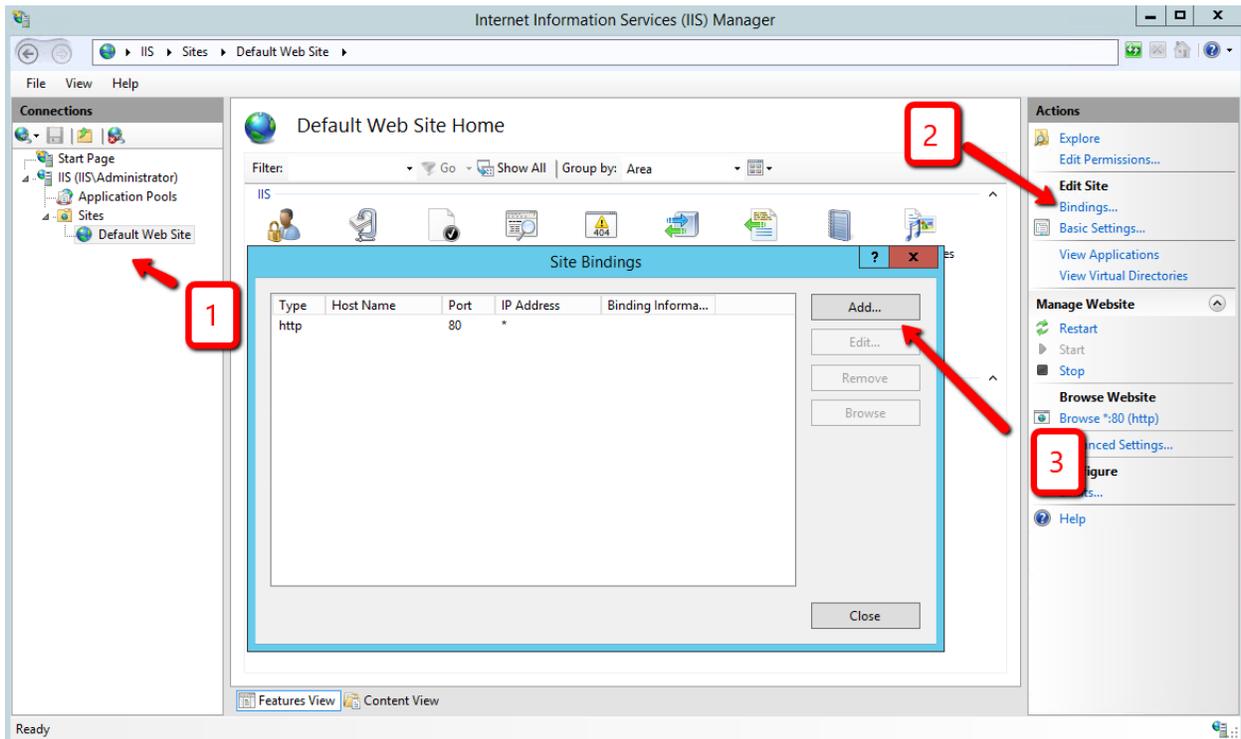
C:\Users\Administrator\Documents>certreq.exe -accept signed.crt
C:\Users\Administrator\Documents>
```

If this step fails, the most common cause is that the issuing CA root certificate is not installed in the server's certificate store. Verify the issuing CA is trusted, or install the CA certificate into the Local Machine—Trusted Root CA certificate store.

#### 2.2.2.4.9 Bind the Certificate to the IIS Web Server

The final step is to bind the certificate to the IIS web server:

1. Open the **IIS Manager** from **Start > Administrative Tools > Internet Information Services (IIS) Manager**.
2. Under **Sites** on the left side of the IIS Manager window, select the desired website.
3. On the right side of the IIS Manager, click **Bindings**.
4. In the **Site Bindings** window, click **Add**.



5. Select the protocol as **https**.
6. Select the IP address of the machine running IIS from the **IP Address** drop-down list, or leave blank to use all available network interfaces.
7. Enter port **443**.

**Add Site Binding**

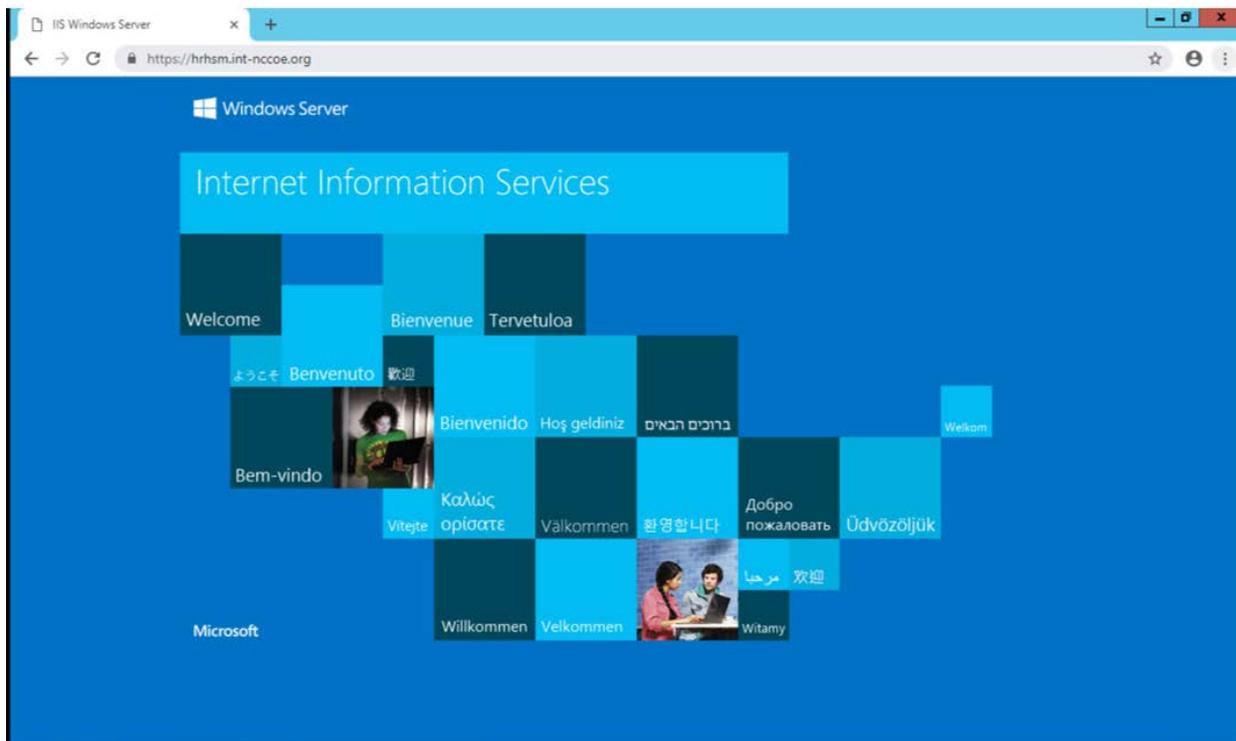
Type:  IP address:  Port:

Host name:

Require Server Name Indication

SSL certificate:

8. In the **SSL certificate:** drop-down, select the certificate that was just installed.
9. Complete the certificate binding in support of SSL/TLS, then click **OK**.
10. Verify the connection is working, open a browser, and enter your URL (e.g., *https://hrhsm.int-nccoe.org:443*). There may be a prompt to accept the certificate for the site. The host name must match the name used in the certificate request and must be registered with the DNS server to resolve the host name to the IP address of the IIS server.



### 2.2.2.5 Venafi Integration Configuration

This section covers the necessary information to integrate Venafi with the Thales TCT Luna SA 1700 for Government HSM. When integrated with the Luna, Venafi can create and store the primary encryption key used to encrypt and decrypt the Venafi database. In this configuration, the Venafi TPP services will not start unless the key stored in the HSM is accessible. This provides an additional hardened layer of security to protect data in the database.

#### 2.2.2.5.1 Prerequisites

To integrate Venafi with the Luna SA HSM, the following prerequisites must be met:

- The Thales TCT Luna HSM is installed and operational.
- The Thales TCT Luna Client is installed on the Venafi server.
- The NTL is established between the Luna Client and the Luna HSM as described in Section [2.2.2.2.9](#).
- The NTL between the Venafi server and the HSM has been verified.
- Venafi has been configured to use the Luna SA HSM.
- The primary encryption key was created on the Luna SA HSM and has been verified.

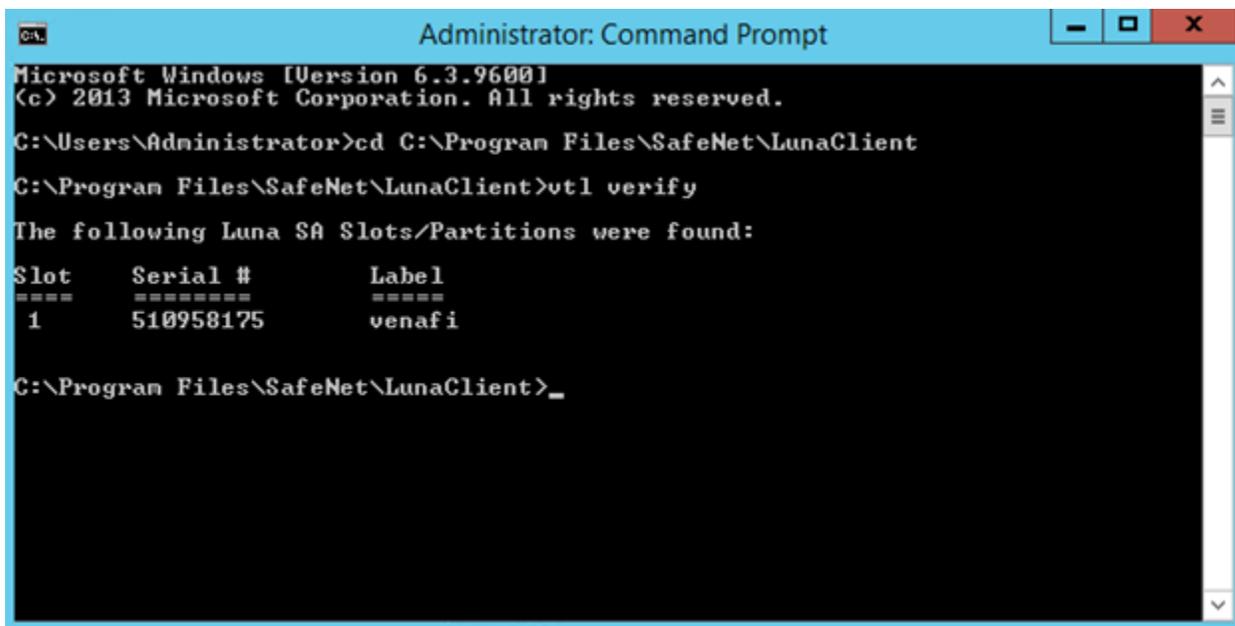
### 2.2.2.5.2 Verify the Network Trust Link Between Venafi and the HSM

The Luna Client installed on the server enables communication between Venafi and the HSM via a secure connection or an NTL. If the NTL has not been set up during HSM/client installation, reference Section [2.2.2.2](#) of this guide.

Use the `vtl verify` command in the installed client directory (typically `C:\Program Files\SafeNet\LunaClient`) to determine if the connection was established and that a partition exists on the HSM that the client can access. If no slot and partition are found, the NTL is not established.

The slot number and partition password will be needed when configuring Venafi to use the HSM.

```
vtl verify
```



```
Administrator: Command Prompt
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.
C:\Users\Administrator>cd C:\Program Files\SafeNet\LunaClient
C:\Program Files\SafeNet\LunaClient>vtl verify
The following Luna SA Slots/Partitions were found:
Slot Serial # Label
==== ===== =====
1 510958175 venafi
C:\Program Files\SafeNet\LunaClient>_
```

For further configuration between the HSM and Venafi TPP, please reference Section [2.6.13.3](#).

## 2.2.3 Day N: Ongoing Security Management and Maintenance

### 2.2.3.1 Prerequisites

- remote system logging server

### 2.2.3.2 Remote System Logging

Refer to the Luna SA syslog commands to use the remote system logging on any UNIX/Linux system that supports the standard syslog service. Refer to the Luna SA syslog commands under “syslog remotehost”

(subcommands “add,” “delete,” and “list”) for more information. The remote host must have User Datagram Protocol (UDP) port 514 open to receive the logging. Refer to the host’s OS and firewall documentation for more information.

1. Type the command below on the Luna SA appliance:

```
lunash:>syslog remotehost add 192.168.1.12
```

2. Start syslog with the “-r” option on the receiving or target system to allow it to receive the logs from the Luna SA appliance(s).

### 2.2.3.3 Audit Logging

With Luna SA, the audit logs can be sent to one or more remote logging servers. Either UDP or Transmission Control Protocol (TCP) protocol can be specified. The default is UDP and port 514.

#### 2.2.3.3.1 UDP Logging

If using UDP protocol for logging:

- The following is required in /etc/rsyslog.conf

```
$ModLoad imudp
```

```
 $InputUDPServerRun (PORT)
```

- Possible approaches include:

1. With templates:

```
$template AuditFile,"/var/log/luna/audit_remote.log"
```

```
 $syslogfacility-text == 'local3' then ?AuditFile;AuditFormat
```

2. Without templates:

```
local3.* /var/log/audit.log;AuditFormat
```

3. Dynamic file name:

```
$template DynFile,"/var/log/luna/%HOSTNAME%.log"
```

```
 if $syslogfacility-text == 'local3' then ?DynFile;AuditFormat
```

- The important thing to remember is that the incoming logs go to local3, and the Port/Protocol that is set on the Luna appliance must be the same that is set on the server running rsyslog.

### 2.2.3.3.2 TCP Logging

Here is an example to set up a remote Linux system to receive the audit logs by using TCP.

- Register the remote Linux system IP address or host name with the Luna SA:

```
lunash:> audit remotehost add -host 172.20.9.160 -protocol tcp -port 1660
```

## 2.3 DigiCert Certificate Authority

### 2.3.1 Day 0: Installation and Standard Configuration

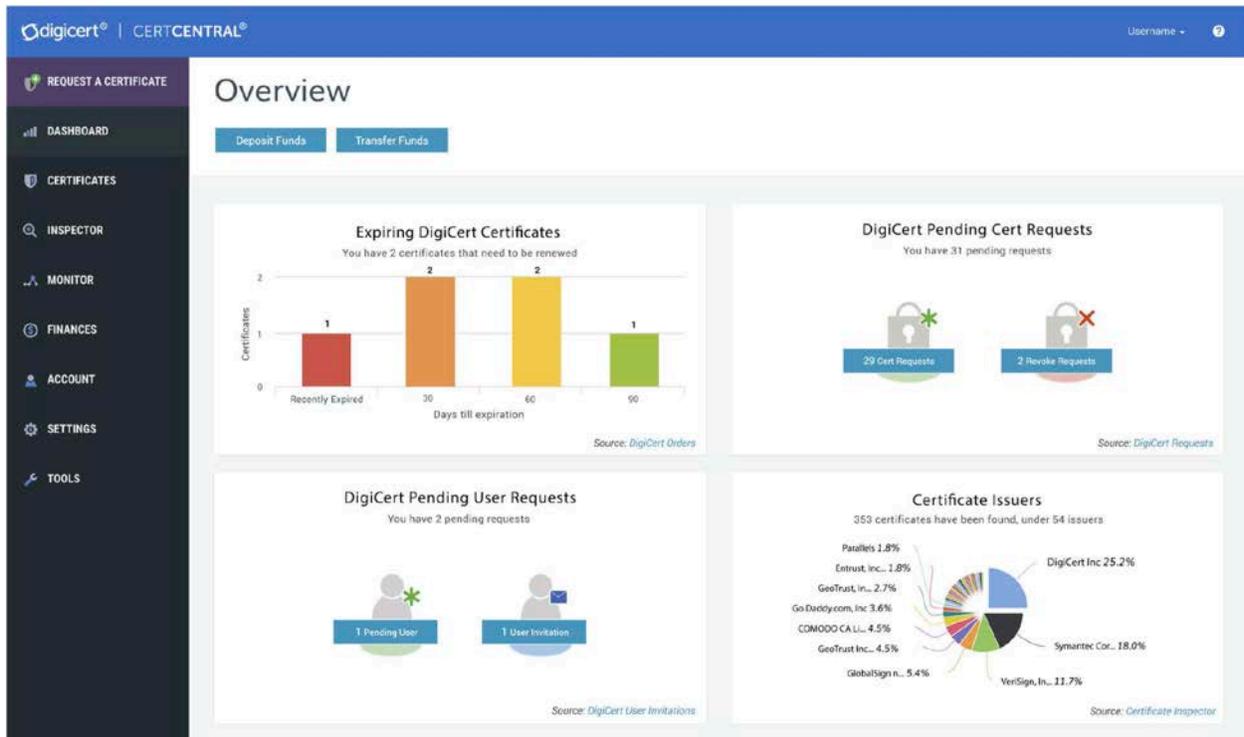
#### 2.3.1.1 *Certificate Prerequisites for Domain Validation and Organization Validation*

- organization validation—can be an individual or group/team
- domain validation process—DNS text (TXT) record validation
- must have resolvable FQDN entered in zone file (*tls.nccoe.org*, *app1.tls.nccoe.org*)
- access to DigiCert’s web-based registration system
- account sign-up

#### 2.3.1.2 *Standard Configuration*

##### 2.3.1.2.1 Account Sign-Up

1. Start the account sign-up process at <https://www.digicert.com/account/signup/>.
2. Complete the **Your information**, **Organization information**, and **Account information** sections.
3. Read and accept the terms of the Certificate Services Agreement. Check the box to acknowledge acceptance of the terms.
4. Click the **Sign Up** button to create a CertCentral account.



### 2.3.1.2.2 Language Preferences

Currently, CertCentral supports the following languages:

- Deutsch
- English
- Español
- Français
- Italiano
- Português
- 한국어
- 日本語
- 简体中文
- 繁體中文

1. To change the language in the CertCentral account, click the account name at the upper-right side of the screen and select **My Profile** from the drop-down list.

2. On the Profile Settings page in the **Language** drop-down list, select the language preference for the account.
3. Click **Save Changes**. The language in CertCentral should now be the same as the one selected.

#### 2.3.1.2.3 Billing Contact

To edit the assigned Billing Contact in the CertCentral account:

1. In the sidebar menu, click **Finances > Settings**.
2. On the Finance Settings page, click **Edit** under **Billing Contact** in the right column.
3. In the **Edit Billing Contact** window, set or change the contact information.
4. Click **Update Billing Contact** to save the change.

#### 2.3.1.2.4 Authentication Settings

Authentication settings allow control over the user login options for the CertCentral account and to set security standards for password requirements and alternative authentication methods.

To access the CertCentral authentication options:

1. In the CertCentral account in the sidebar menu, click **Settings > Authentication Settings**.  
On this page, the following settings can be changed:
  - **Minimum Length:** Change the minimum allowed password character length.
  - **Minimum Categories:** Change the variety of characters allowed (uppercase, lowercase, numbers, and symbols).
  - **Expires After:** Change the password expiration policy.
  - **Two-Factor Authentication:** Enable or disable onetime password two-factor authentication for CertCentral users.
2. Configure the authentication settings as desired, then click **Save Settings**.

#### 2.3.1.2.5 Security Assertion Markup Language (SAML) Single Sign-On Prerequisites

SAML is a highly recommended DigiCert feature for secure user authentication. However, it is not required to duplicate the TLS lab setup. For more information on SAML, please refer to guidance at:

- <https://pages.nist.gov/800-63-3/sp800-63-3.html>

Before beginning, make sure the following prerequisites are met:

- Have a CertCentral account.
- Have SAML enabled on the CertCentral account. (To get the SAML features turned on for the CertCentral account, contact the DigiCert account representative or the DigiCert support team.)

Once activated, in the sidebar menu, under Settings, see the Single Sign-On and SAML Certificate Request menu options.)

- Have an identity provider (IdP).
- Have the IdP metadata (dynamic or static).
- Have admin privileges on the CertCentral account (or have manager privileges on the CertCentral account with the Allow access to SAML settings permission).

#### 2.3.1.2.6 Organization Validation

To validate an organization, DigiCert firsts verifies the organization requesting a certificate is in good standing. This may include confirming good standing and active registration in corporate registries. It may also include verifying the organization is not listed in any fraud, phishing, or government-restricted entities and anti-terrorism databases. Additionally, DigiCert verifies the organization requesting a certificate is, in fact, the organization to which the certificate will be issued. DigiCert also verifies the organization contact.

1. In the CertCentral account, using the sidebar menu, click **Certificates > Organizations**.
2. On the **Organizations** page, click **New Organization**.
3. On the **New Organization** page, under **Organization Details**, enter the specified organization information:

|                                                         |                                                                                                                                                        |
|---------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Legal Name</b>                                       | Enter the organization's legally registered name.                                                                                                      |
| <b>Assumed Name</b>                                     | If the organization has a doing-business-as name and the name should appear on the certificates, enter the name here.<br>If not, leave this box blank. |
| <b>Organization Phone Number</b>                        | Enter a phone number at which the organization can be contacted.                                                                                       |
| <b>Country</b>                                          | In the drop-down list, select the country where the organization is legally located.                                                                   |
| <b>Address 1</b>                                        | Enter the address where the organization is legally located.                                                                                           |
| <b>Address 2</b>                                        | Enter a second address, if applicable.                                                                                                                 |
| <b>City</b>                                             | Enter the city where the organization is legally located.                                                                                              |
| <b>State/Province/<br/>Territory/Region/<br/>County</b> | Enter the state, province, territory, region, or county where the organization is legally located.                                                     |

|                             |                                                               |
|-----------------------------|---------------------------------------------------------------|
| <b>Zip Code/Postal Code</b> | Enter the zip or postal code for the organization's location. |
|-----------------------------|---------------------------------------------------------------|

4. Under **Validation Contact**, provide the contact's information:

|                        |                                                             |
|------------------------|-------------------------------------------------------------|
| <b>First Name</b>      | Enter the contact's first name.                             |
| <b>Last Name</b>       | Enter the contact's last name.                              |
| <b>Job Title</b>       | Enter the contact's job title.                              |
| <b>Email</b>           | Enter an email address at which the contact can be reached. |
| <b>Phone Number</b>    | Enter a phone number at which the contact can be reached.   |
| <b>Phone Extension</b> | Enter the contact's extension, if applicable.               |

5. When finished, click **Save Organization**.  
Submit an organization for validation.
6. In the CertCentral account, using the sidebar menu, click **Certificates > Organizations**.
7. On the **Organizations** page, use the drop-down list, search box, and column headers to filter the list of organizations.
8. Click the link for the organization being submitted for validation and authorization for certificates.
9. On the organization's information page in the **Submit Organization for Validation** section, select the validation types (certificates) needed for DigiCert to validate the organization's information below:
- OV—Normal Organization Validation (Recommended)
  - EV—Extended Organization Validation (EV)
  - Private SSL—DigiCert Private SSL Certificate
  - CS—Code Signing Organization Validation
  - EV CS—Code Signing Organization Extended Validation (EV CS)
  - DS—Document Signing Validation
  - Add verified contact (EV/EV CS, and CS).

If the organization validation chosen is not OV, refer to <https://docs.digicert.com/manage-certificates/organization-domain-management/managing-domains-cc-guide/> for additional details.

10. When finished, click **Submit for Validation**.

### 2.3.1.2.7 Domain Validation

DigiCert's domain validation process ensures the organization requesting a certificate is authorized to request a certificate for the domain in question. Domain validation can include emails or phone calls to the contacts listed in a domain's WHOIS record as well as emails to default administrative addresses at the domain. For example, DigiCert may send an authorization email to the administrator@domain.com or webmaster@domain.com but would not send an authorization email to [tech@domain.com](mailto:tech@domain.com).

Note: To validate a domain by using DNS TXT, see the steps below. To use an alternative method, refer to <https://docs.digicert.com/manage-certificates/organization-domain-management/managing-domains-cc-guide/>.

#### Step I: Add and Authorize a Domain for TLS/SSL Certificates

1. In the CertCentral account in the sidebar menu, click **Certificates > Domains**.
2. On the **Domains** page, click **New Domain**.
3. On the **New Domain** page, under **Domain Details**, enter the following domain information:
  - a. **Domain Name**  
In the box, enter the domain name that the certificates will secure (for example, *yourdomain.com*).
  - b. **Organization**  
In the drop-down list, select the organization to assign to the domain.
4. Under **Validate This Domain For**, check the validation types needed for the domain to be validated:
  - a. **OV—Normal Organization Validation (Recommended)**  
Use this option to order Standard SSL, Secure Site SSL, Wildcard SSL, Secure Site Wildcard SSL, Multi-Domain SSL, and Secure Site Multi-Domain SSL certificates for this domain.
5. Under **Domain Control Validation (DCV) Method**, select **DNS TXT Record**.  
Note: The default DCV method is by verification email.
6. When finished, click **Submit for Validation**.

#### Step II: Use DNS TXT Record to Demonstrate Control Over the Domain

1. **Create the DNS TXT record:**
  - a. Under **User Actions** in the **Your unique verification token** box, copy the verification token.

To copy the value to the clipboard, click in the text field.

Note: The unique verification token expires after 30 days. To generate a new token, click the **Generate New Token** link.

- b. Go to the organization's DNS provider's site and create a new TXT record.
- c. In the **TXT Value** field, paste the verification code copied from the CertCentral account.
- d. Host field
  - i. **Base Domain**  
If validating the base domain, leave the **Host** field blank, or use the @ symbol (dependent on the DNS provider requirements).
  - ii. **Subdomain**  
In the **Host** field, enter the subdomain being validated.
- e. In the record type field (or equivalent), select **TXT**.
- f. Select a Time-to-Live value, or use the organization's DNS provider's default value.
- g. Save the record.

## 2. Verify the DNS TXT record:

- a. In the CertCentral account, using the sidebar menu, click **Certificates > Domains**.
- b. On the **Domains** page in the **Domain Name** column, click the link for the domain.
- c. On the domain information page (e.g., *example.com*) at the bottom of the page, click **Check TXT**.

## 2.3.2 Day 1: Integration Configuration

### 2.3.2.1 Generate API Key

DigiCert Services API provides the foundation for the CertCentral web portal. Because DigiCert developed CertCentral as an API-first web application, the DigiCert Services API allows one to automate CertCentral web application workflows and typical certificate processes and to streamline certificate management. To access DigiCert Services API documentation, see the [DigiCert Developers Portal](#). The services API uses RESTful conventions. The DigiCert Services API requires a DigiCert Developer API key, which is included in the header as part of each request.

#### Generate API Key

1. In the CertCentral account, using the side bar menu, click **Account > Account Access**.
2. On the **Account Access** page in the **API Key** section, click **Add API Key**.
3. In the **Add API Key** window, in the **Description** box, enter a description/name for the API key.

4. In the **User** drop-down, select the user to whom they key should be assigned/linked.  
Note: When linking a key to a user, link that user's permissions to the key. The API key has the same permissions as the user and can perform any action that the user can.
5. Click **Add API Key**.
6. In the **New API Key** window, click on the generated key to copy it.
7. Save the key in a secure location.  
Note: The API keys will be displayed only one time. If the window is closed without recording the new API key, the key cannot be recorded again.
8. When done, click **I understand I will not see this again**.

### 2.3.2.2 *Venafi Integration (Automated)*

Venafi integrates with the DigiCert Services API. The integrated solution leverages DigiCert's Online Certificate Status Protocol (OCSP) infrastructure and API integration with Venafi's machine identity protection platform. Customers can customize specific features, from fully automating certificate provisioning to enforcing internal policies, allowing them to address industry regulations such as Payment Card Industry Data Security Standard, Health Insurance Portability and Accountability Act of 1996, and General Data Protection Regulation. The integrated solution also simplifies integration of machine identity protection across a wide variety of systems and allows customers to fulfill certificate requests.

### 2.3.2.3 *Order Certificate Directly Through CertCentral (Manual Process)*

The TLS certificate life cycle begins when a TLS certificate is ordered. The process for requesting any of the available certificates is the same:

- Create a CSR.
- Fill out the order form by clicking the **Request a Certificate** button from the left navigation bar.
- Complete domain control validation for the domains on the order (in other words, demonstrate control over the domains).
- Complete organization validation for the organization on the certificate order.

### 2.3.2.4 *Order an OV Single- or Multi-Domain TLS Certificate*

When ordering Multi-Domain SSL certificates, add **Other Hostnames (SANs)** to the certificate order. This option is not available for the single-domain certificates.

1. **Create the CSR.**
2. **Select the OV Single- or Multi-Domain SSL/TLS certificate.**
  - a. In the CertCentral account in the sidebar menu, click **Request a Certificate**, and then under All Products, click **Product Summary**.

- b. On the Request a Certificate page, look over the certificate options and select the certificate.

### 3. **Add the CSR.**

On the Request page, under Certificate Settings, upload the CSR to or paste it in the **Add Your CSR** box.

When copying the text from the CSR file, make sure to include the -----**BEGIN NEW CERTIFICATE REQUEST**----- and -----**END NEW CERTIFICATE REQUEST**----- tags.

### 4. **Common Name**

Type the common name in the box, or under Common Name, expand **Show Recently Created Domains**, and select the domain from the list.

### 5. **Other Hostnames (SANs)**

In the **Other Hostnames (SANs)** field, enter the additional host names needed for the certificate to be secure.

For Multi-Domain certificates, four SANs are included in the base price of each certificate. Additional SANs (over those included in the base price) increase the cost of the certificate.

### 6. **Validity Period**

Select a validity period for the certificate: one year, two years, custom expiration date, or custom length.

#### **Custom Validity Periods**

- a. Certificate pricing is prorated to match the custom certificate length.
- b. Certificate validity cannot exceed the industry-allowed maximum life-cycle period for the certificate.  
For example, a 900-day validity period cannot be set for a certificate.

### 7. **Additional Certificate Options**

The information requested in this section is optional.

Expand **Additional Certificate Options** and provide information as needed.

a. **Signature Hash**

Unless there is a specific reason for choosing a different signature hash, DigiCert recommends using the default signature hash: Secure Hash Algorithm 256.

b. **Server Platform**

Select the server or system generated on the CSR.

c. **Organization Unit(s)**

Adding organization units is optional. This field can be left blank. If the CSR includes an organization unit, we use it to populate the Organization Unit(s) box.

Note: If an organization's units are included in the order, DigiCert will need to validate them before issuing a certificate.

d. **Auto-Renew**

To set up automatic renewal for this certificate, check **Auto-renew order 30 days before expiration**.

With auto-renew enabled, a new certificate order will be automatically submitted when this certificate nears its expiration date. If the certificate still has time remaining before it expires, DigiCert adds the remaining time from the current certificate to the new certificate (as long as 825 days or approximately 27 months).

Note: Auto-renew cannot be used with credit card payments. To automatically renew a certificate, the order must be charged to an account balance.

8. To add an organization, click **Add Organization**. Add a new organization or an existing organization in the account.

Note: When adding a new organization, DigiCert will need to validate the organization before issuing a certificate.

9. **Add Contacts**

Two different contacts can be added to the order: Organization and Technical.

**Organization Contact (required)**

The **Organization Contact** is someone who works for the organization included in the certificate order. DigiCert will contact the **Organization Contact** to validate the organization and verify the request for OV TLS/SSL certificates. DigiCert also sends this person an order confirmation and renewal emails.

### Technical Contact (optional)

In addition to the **Organization Contact**, the **Technical Contact** will receive order emails, including the one with the certificate attached, as well as renewal notifications.

### 10. Additional Order Options

The information asked for in this section is optional.

Expand **Additional Order Options** and add information as needed.

a. **Comments to Administrator**

Enter any information the administrator might need for approving the request, such as the purpose of the certificate.

b. **Order Specific Renewal Message**

To create a renewal message for this certificate right now, type a renewal message with information possibly relevant to the certificate's renewal.

Note: Comments and renewal messages are not included in the certificate.

### 11. Additional Emails

Enter the email addresses (comma separated) for the people who want to receive the certificate notification emails, such as certificate issuance, duplicate certificate, and certificate renewals.

Note: These recipients cannot manage the order; however, they will receive all the certificate-related emails.

### 12. Select Payment Method

Under **Payment Information**, select a payment method to pay for the certificate.

### 13. Certificate Services Agreement

Read the agreement and check **I agree to the Certificate Services Agreement**.

### 14. Click **Submit Certificate Request**.

### 2.3.2.5 *Manage Order Within CertCentral (Manual)*

After submitting the TLS certificate order, DCV and organization validation must be completed before DigiCert can issue the certificate.

If the certificate does not immediately issue, please ensure all Day 0 activities have been completed (Organization Validation and Domain Validation).

### 2.3.2.6 *Download a Certificate from the CertCentral Account*

After DigiCert issues the certificate, access it from inside the CertCentral account.

1. In the CertCentral account, go to the **Orders** page.  
In the sidebar menu, click **Certificates > Orders**.
2. On the **Orders** page, use the filters and advanced search features to locate the certificate to be downloaded.
3. In the **Order #** column of the certificate to be downloaded, click the **Quick View** link.
4. In the **Order #** details pane (on the right), using the **Download Certificate As** drop-down, select the certificate format to be used.
  - a. **.crt (best for Apache/Linux)**  
Download the certificate in a .crt format, best for Apache/Linux platforms.
  - b. **.pb7 (best for Microsoft and Java)**  
Download the certificate in a .pb7 format, best for Microsoft and Java platforms.
5. (OPTIONAL) In the **Download Certificate As** drop-down, click **More Options** to see more **Server Platform** options and **File Type** options or to download only the **Certificate**, the **Intermediate Certificate**, or the **Root Certificate**.
6. **Download a Combined Certificate File**

In the **Download Certificate** window, under **Combined Certificate Files**, use any of these options to download the combined SSL certificate file.

- a. **Platform specific**

In the **Server Platform** drop-down, select the server where the SSL/TLS certificate will be installed, and then click **Download**.

b. **File type specific**

In the **File Type** drop-down, select the SSL/TLS file format to be downloaded, and then click **Download**.

7. In the **Download Certificate** window, under **Individual Certificate Files**, use one of these options to download an individual certificate file.
  - a. **Server certificate file**  
Under **Certificate**, click the **Download** link. Save the server certificate file to the server or workstation, making sure to note the location.
  - b. **Intermediate certificate file**  
Under **Intermediate Certificate**, click the **Download** link. Save the intermediate certificate file to the server or workstation, making sure to note the location.
  - c. **Root certificate file**  
Under **Root Certificate**, click the **Download** link. Save the root certificate file to the server or workstation, making sure to note the location.

## 2.3.3 Day N: Ongoing Security Management and Maintenance

### 2.3.3.1 *Ongoing Auditing*

Once the users, divisions, domains, and organizations have been added, an account audit may need to be executed to highlight areas where training is required, reconstruct events, detect intrusions, and discover problem areas.

### 2.3.3.2 *Run an Audit*

1. In the CertCentral account, using the sidebar menu, click **Account > Audit Logs**.
2. On the **Audit Logs** page, use the filters to filter the results of the audit.
  - a. Choose a filter (for example, User).
  - b. In the filter drop-down, select an option (for example, select a user).
  - c. Wait for the filter to modify the audit log before using another filter.

### 2.3.3.3 *Set Up Audit Log Notifications*

To be of help to the organization, log data must be reviewed. The audit log notifications feature can be used to keep aware of certain activities as well as make log review more meaningful.

1. In the CertCentral account, using the sidebar menu, click **Account > Audit Logs**.
2. On the **Audit Logs** page, click **Audit Log Notifications**.

3. On the **Audit Log Notifications** page, under **Create a New Notification**, take the following steps:

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Email Address</b>   | Enter the email address of the person to whom the audit log notifications are to be sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Division</b>        | In the drop-down, select the divisions whose account activity needs to be monitored.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Notify me about</b> | Check any of the following options: <ul style="list-style-type: none"><li>• <b>Order Changes</b><br/>Alerts if any changes are made to certificate orders.</li><li>• <b>User Changes</b><br/>Alerts if any edits are made to any user accounts.</li><li>• <b>User Logins</b><br/>Alerts of all account logins.</li><li>• <b>Logins from Invalid IP Addresses</b><br/>Alerts if any account logins are made from invalid IP addresses.</li><li>• <b>Certificate Revocations</b><br/>Alerts to all certificates are revocations.</li></ul> |

4. When finished, click **Save Changes**.

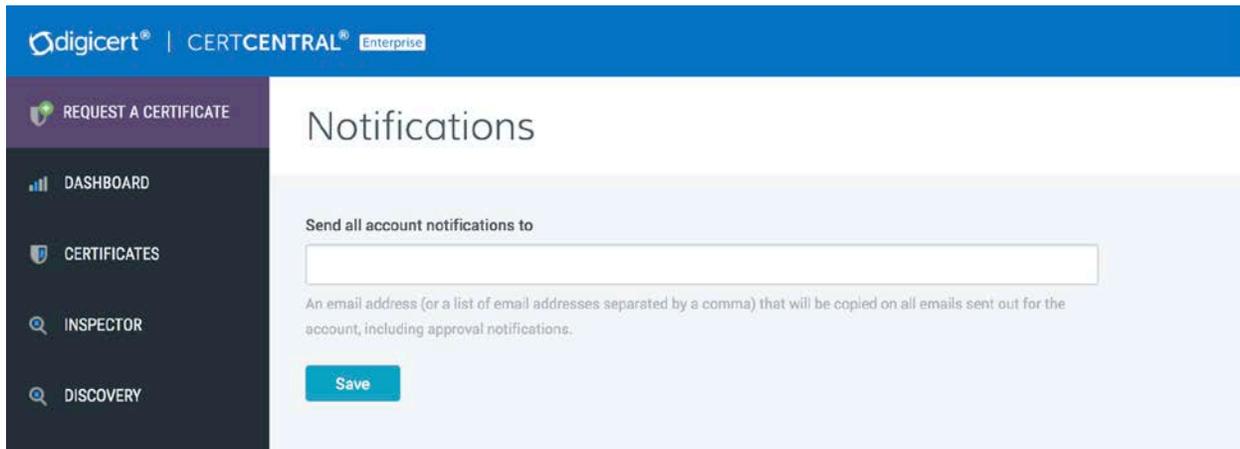
The designated individual should start receiving the selected audit log notifications.

#### 2.3.3.4 *Notification Management*

Typically, notifications are not strictly required when utilizing Venafi to manage certificates, as expiring certificates are renewed automatically (or not) based on configured policy within Venafi. However, it is beneficial to configure renewal notifications within CertCentral.

##### 2.3.3.4.1 *Account Notifications*

Before sending email from an account, assign an email address to receive a copy of any message sent (e.g., approval notifications). Configure renewal notifications and add default renewal messages that include renewal notifications.



#### 2.3.3.4.2 Set Up Email Notification Accounts

1. In the CertCentral account's sidebar menu, click **Settings > Notifications**.
2. On the **Notifications** page in the **Send all account notifications to** box, add the email addresses that should be copied on all emails sent from the account.

**Note:** When setting up multiple notification accounts, use commas to separate the email addresses.

3. When finished, click **Save**.

#### 2.3.3.4.3 Certificate Renewal Notifications

After DigiCert has issued the first certificate, configure the **Certificate Renewal Settings** (such as when renewal notifications are sent and to whom notifications are sent) to help prevent unexpected certificate expirations.

When configuring the certificate renewal settings, there are two options:

1. **Nonescalation Certificate Renewals**  
This option sends renewal notifications to the same email addresses at every stage as certificates get closer to expiration or after they have expired.
2. **Escalation Certificate Renewals**  
This option configures email escalation settings in which additional email addresses can receive renewal notifications at critical stages as certificates get closer to expiring or after they have expired. This allows additional oversight of certificate expiration.

#### 2.3.3.4.4 Configure Nonescalation Renewal Notifications

Use the steps below to send all renewal notifications to the same email addresses at every stage as certificates get closer to expiring or after they have expired.

1. In the CertCentral account's sidebar menu, click **Settings > Preferences**.
2. On the **Division Preferences** page, scroll down to the **Certificate Renewal Settings**, and uncheck **Enable Escalation**.
3. In the **Send request renewal notifications to** box, enter the email addresses for the people who should receive the renewal notifications (comma separated).
4. Under **When certificates are scheduled to expire in**, check the boxes to indicate when to send renewal notices.

Note: These options determine when email notifications are sent. For example, if only **30 days**, **7 days**, and **3 days** are checked, no email notifications will be sent **90 days** or **60 days** before certificates expire.

5. In the **Default Renewal Message** box, type an optional renewal message for inclusion in all the renewal notification emails.
6. Click **Save Settings** when finished.

#### 2.3.3.4.5 Configure Escalation Renewal Notifications

Email escalation settings allow control over what email addresses will receive renewal notifications at each stage as certificates approach or reach expiration.

1. In the CertCentral account's sidebar menu, click **Settings > Preferences**.
2. On the **Division Preferences** page, scroll down to **Certificate Renewal Settings**, and check **Enable Escalation**.
3. Under **Days before expiration**, check the boxes for when renewal notices should be sent.
4. Under **Additional email addresses or distribution lists**, enter the email addresses for the people who should receive each renewal notification (comma separated).
5. In the **Default Renewal Message** box, type an optional renewal message for inclusion in all renewal notification emails.
6. Click **Save Settings** when finished.

### 2.3.3.5 *Managing Custom Order Fields*

CertCentral allows users to add custom fields to certificate order forms. Use the custom field metadata to search or sort a set of certificate orders that match the metadata search criteria.

Note: The **Custom Fields** feature is off by default. To enable this feature for a CertCentral account, please contact a DigiCert account representative.

Once enabled for a CertCentral account, the **Custom Order Fields** menu option is added to the sidebar menu under **Settings (Settings > Custom Order Fields)**.

#### 2.3.3.5.1 Custom order form field features

- Apply to Future and Present Requests—When a custom order form field is added, the field is also added to pending requests. If the field is required, the pending requests cannot be approved until the field is completed.
- Apply to Entire Account—When custom order form fields are added, the fields are applied to the order forms for the entire account. Custom order form fields cannot be set per division.
- Apply to All Certificate Types—When custom order form fields are created, the fields are added to the order forms for all certificate types (SSL, Client, Code Signing, etc.). A custom order form field cannot be added to the order forms for only SSL certificate types.
- Apply to Guest URLs—When custom order form fields are added, these fields are added to the certificates ordered from directly inside the CertCentral account as well as from any guest URLs that have been sent.
- Different Types to Choose From—When custom order form fields are created, different types of fields can be added such as single-line and multiple-line text boxes and email address and email address list boxes.
- Required or Optional—When custom order form fields are added, they can be required or optional. Required fields must be completed before the order can be approved. Optional fields can be left blank.
- Deactivated or Activated—After a custom order form field has been added, the field can be deactivated (removed) and activated (added back) as needed. Deactivated fields are removed from pending requests but not from issued orders. Activated fields are added to pending requests. If the field is required, it must be completed before the request can be approved.

#### 2.3.3.5.2 Add a Custom Field to Request Forms

1. In the CertCentral account in the sidebar menu, click **Settings > Custom Order Fields**.
2. On the **Custom Order Form Fields** page, click the **Add Custom Order Form Field** link.
3. In the **Add Custom Order Form Field** window, configure the custom field:

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Label</b>                                              | In the box, type a name/label for the field (e.g., Direct Report’s Email Address).                                                                                                                                                                                                                                                                                                                                                           |
| <b>Input Type</b>                                         | In the drop-down list, select an input type for the field (i.e., email address).<br>Input Types: <ul style="list-style-type: none"> <li>▪ <b>Anything:</b> Single-line text box</li> <li>▪ <b>Text:</b> Multiline text box</li> <li>▪ <b>Integer:</b> Number box (limited to nondecimal whole numbers)</li> <li>▪ <b>Email Address:</b> Single email address box</li> <li>▪ <b>Email Address List:</b> Multiple email address box</li> </ul> |
| <b>This field should be required for all new requests</b> | If the field needs to be completed before the request can be submitted (or approved for pending requests), check this box.<br>Note: If this box is not checked, the field appears on the order form with the word “optional” in the box. The requester does not need to complete the box for the request to be submitted (or approved for pending requests).                                                                                 |

4. When finished, click **Add Custom Form Field**.

### 2.3.3.6 *User Management*

Add a user to the CertCentral account.

1. In the CertCentral account in the sidebar menu, click **Account > Users**.
2. On the **Users** page, click **Add User**.
3. On the **Add User** page in the **User Details** section, enter the new user’s information.
4. In the **User Access** section, assign the user a role, and configure their division access if applicable:

|                                                      |                                                                                                                                                                      |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Username</b>                                      | We recommend using the user’s email address.                                                                                                                         |
| <b>Restrict this user to specific divisions</b>      | Check this box if the role should be restricted to specific divisions.<br>Note: This option appears only if divisions within the CertCentral account are being used. |
| <b>User is restricted to the following divisions</b> | Select the divisions to which the role is restricted.<br>Note: This drop-down appears only if “Restrict this user to specific divisions” is checked.                 |

|                                                                       |                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Allow this user to log in only through SAML Single Sign-On SSO</b> | Check this box if this user should be restricted from being able to log in with username and password. Note: SAML SSO must be configured in the account and the IdP must be configured with this user's information. |
| <b>Role</b>                                                           | Select a role for the new user: Administrator, Standard User, Finance Manager, or Manager.                                                                                                                           |
| <b>Limit to placing and managing their own orders</b>                 | To create a Limited User role, select Standard User, and check this box.                                                                                                                                             |

- When finished, click **Add User**.

### What's next

The newly added user will receive an email with instructions for setting up their account credentials and can use them to sign in to their CertCentral account.

#### 2.3.3.7 Revalidation Processes

Organization and domain validation typically expire in two years. When the validation status nears expiration, CertCentral sends a notification and automatically initiates a revalidation process. The user should complete the steps outlined in Day 0 Organization Validation and Domain Validation. The standards governing the requirements surrounding (re)validation processes are encapsulated in the CA/Browser Forum's Baseline Requirements (<https://cabforum.org/baseline-requirements-documents/>). The specific allowed methods of validation will change over time.

Note: This revalidation process is outside the Venafi certificate management processes.

- OV validation and revalidation: two years
- DV validation and revalidation: two years
- EV validation and revalidation: one year

Note: Extended Validation provides additional levels of vetting surrounding the legal entity represented in a certificate. Vetting ensures that a complete picture of the identity, which has proven control over the domain in the certificate, is available to user agents verifying the certificate.

## 2.4 F5 BIG-IP Local Traffic Manager (LTM)

BIG-IP Virtual Edition (VE) is a version of the BIG-IP system that runs as a virtual machine in specifically supported hypervisors. BIG-IP VE emulates a hardware-based BIG-IP system running a VE-compatible version of BIG-IP software.

## 2.4.1 Day 0: Installation and Standard Configuration

### 2.4.1.1 Prerequisites

- VMware ESX 6.5
- 2 virtual Central Processing Units (CPUs)
- 4 GB RAM
- 1 x VMXNET3 virtual network adapter or Flexible virtual network adapter (for management)
- 1 x virtual VMXNET3 virtual network adapter
- 1 x 100 GB Small Computer System Interface disk, by default
- connection to a common NTP source
- SMTP for BIG-IP to send email alerts
- a computer with internet (browser) access to activate license
- license key for F5 BIG-IP
- F5 Support ID account

### 2.4.1.2 Download the Virtual Appliance

To deploy BIG-IP VE, download the open virtualization appliance (OVA) file to your local system.

1. Open the F5 Downloads page at <https://downloads.f5.com>.
2. Log in with an F5 Support ID.
3. In the Downloads Overview page, click **Find a Download** button.
4. In the Select a Product Line page, click the **BIG-IP v13.x / Virtual Edition...** link.
5. In the Select a Product Version... page, click the **13.1.1.4\_Virtual-Edition** link.
6. In the Software Terms... page, review, then click **I Accept** button to agree to terms and conditions.
7. In the Select a Download page, click the **BIGIP-13.1.1.4-0.0.4.ALL-scsi.ova** link.
8. In the Download Locations page, click the link nearest to the correct region.
9. Save the OVA file to the local computer.

### 2.4.1.3 *Deploying the BIG-IP OVA*

Use the Deploy Open Virtualization Format (OVF) Template wizard from within the VMware vSphere client. Follow the steps in this procedure to create an instance of the BIG-IP system that runs as a virtual machine on the host system.

1. Start the vSphere Client and log in.
2. Launch the **Deploy OVF Template** wizard.
3. Select an OVF template from Local file. Select the previously downloaded OVA file.
4. In the Virtual machine name field, type in `F51b1.ext-nccoe.org`. Then select the location for this virtual machine. Click **Next**.
5. Select the compute resource and click **Next**.
6. Verify that the OVF template details are correct, then click **Next**.
7. Review the template details, then click **Next**.
8. Review License agreements. Select “I accept...” and click **Next**.
9. Read and accept the license agreement, and click **Next**.
10. Accept the default value **2 CPUs** and click **Next**.
11. Accept the default value **Thick Provision Lazy Zeroed** and click **Next**.
12. Assign the networks to the network interface cards (NICs) and click **Next**.
  - NIC 1: VLAN 2199 (Datacenter Secure)
  - NIC 2: VLAN 2201
  - NIC 3: VLAN 2197 (DMZ)
13. Review information and click **Finish**.

### 2.4.1.4 *Assigning a Management IP Address to a BIG-IP VE Virtual Machine*

The BIG-IP VE virtual machine needs an IP address assigned to its virtual management port.

1. In the main vSphere client window, **Power On** the BIG-IP.
2. Launch a Console session for the BIG-IP.
3. At the login prompt, log in as `root / default`.
4. At the config # prompt, type `config`.

The Configure Utility panel appears.

5. Press **Enter** for **OK**.

The Configure IP Address panel appears.

6. For “Automatic configuration...”, choose **No**.
7. For IP Address, type 192.168.3.85 Choose **OK**.
8. For Netmask, type 255.255.255.0. Choose **OK**.
9. For Management Route, choose **Yes**.
10. For Management Route, type 192.168.3.1 Choose **OK**. The Confirm Configuration panel appears. (This Gateway address is used for management traffic.)
11. Review the IP information, and choose **Yes**. Return to the config # prompt.

#### 2.4.1.5 *Log in to BIG-IP for the First Time*

After the initial login to the BIG-IP, the Setup Utility will guide through the initial setup process.

1. Open the browser and navigate to the BIG-IP address <https://192.168.3.85>.
2. Log in as the default admin/admin.



3. The Setup Utility panel appears, then click **Next**.
4. For License, click **Activate**.
5. As a prerequisite, the user should already have a BIG-IP VE license key. Copy the key and paste in the Base Registration Key field.
6. This step is dependent on internet access for the BIG-IP.

- a. If the management route configured in the previous section has a path to internet, select **Automatic**. Click **Next**. Review the End User License Agreement (EULA) and click **Agree**. Then go to step 7.
  - b. Otherwise, select **Manual**. Click **Next**.
  - c. **Left-click** in the Dossier field, and select all the encrypted text with **Ctrl-A**. Copy the selected text with **Ctrl-C**.
  - d. Assuming the administration computer has internet access, click the “Click here to access F5...” link. A new browser tab appears.
  - e. In the Enter Your Dossier field, paste in the copied text. Click **Next**.
  - f. Review the EULA, and select “I have read and agree... ” Click **Next**.
  - g. Left-click the license text field, and select all text with **Ctrl-A**. Copy selected text with **Ctrl-C**.
  - h. Return to the BIG-IP Setup Utility. In the License field, paste in the copied text. Click **Next**.
7. Some BIG-IP services will restart and log the user off the BIG-IP. It will automatically resume. Click **Continue**.
  8. Review the License page. Click **Next**.
  9. On the Resource Provisioning page, verify that the only default value, **Local Traffic (LTM)**, is selected and set to **Nominal**. Click **Next**.
  10. On the Device Certificates page, leave the default as self-sign device Certificate. Click **Next**.
  11. On the Platform page, fill these values. Then click **Next**.

| Field                         | Value               | Comments |
|-------------------------------|---------------------|----------|
| Management Port Configuration | 443                 |          |
| IP Address                    | 192.168.3.85        |          |
| Network Mask                  | 255.255.255.0       |          |
| Management Route              | 192.168.3.1         |          |
| Host Name                     | f51b1.ext-nccoe.org |          |

|               |                 |                                                 |
|---------------|-----------------|-------------------------------------------------|
| Time Zone     | EST             |                                                 |
| Root Account  | <your password> | Refer to NIST SP 800-63B for password guidance. |
| Admin Account | <your password> | Refer to NIST SP 800-63B for password guidance. |

**General Properties**

Management Port Configuration:  Automatic (DHCP)  Manual

Management Port: IP Address[/prefix]: 192.168.3.85  
Network Mask: 255.255.255.0 (255.255.255.0 ▼)  
Management Route:

Host Name: f5lb1.ext-nccoe.org  
Host IP Address: Use Management Port IP Address ▼  
Time Zone: America/New York ▼

**Redundant Device Properties**

Root Folder Device Group: None  
Root Folder Traffic Group: traffic-group-1 ▼

**User Administration**

Root Account:  Disable login  
 Disable default admin, use alternate

Admin Account: Password:   
Confirm:

SSH Access:  Enabled  
SSH IP Allow: \*All Addresses ▼

12. System logs off the user with password change. Log back in with the new admin password.
13. In the Standard Network Configuration page, click **Next**.
14. In the Redundant Device Wizard Options page, **Un-Select** Display configuration synchronization options.
15. In the Internal Network Configuration page, fill in these values.

|                 |               |
|-----------------|---------------|
| Address         | 192.168.4.85  |
| Netmask         | 255.255.255.0 |
| VLAN Interfaces | internal      |
| Tagging         | untagged      |

16. Click **Add**, then click **Next**.

17. In the External Network Configuration page, fill in these values.

|                 |                      |
|-----------------|----------------------|
| Address         | <i>192.168.5.86</i>  |
| Netmask         | <i>255.255.255.0</i> |
| VLAN Interfaces | <i>external</i>      |
| Tagging         | <i>untagged</i>      |

18. Click **Add**, then click **Finished**.

### 2.4.1.6 BIG-IP Configuration Utility

There are at least two ways to administer the BIG-IP.

- Use SSH to connect to the BIG-IP to access the command line interface, referred to as traffic management shell (TMSH).
- With a web browser, navigate to the management URL—referred to as Configuration utility and mainly used in this guide.
  1. Open browser and navigate to the BIG-IP address [https://192.168.3.85\\_](https://192.168.3.85_)
  2. Log in as admin, and use the password modified from the default during Setup wizard.



## BIG-IP Configuration Utility

F5 Networks, Inc.

**Hostname**

f5lb1.ext-nccoe.org

**IP Address**

192.168.3.85

**Username**

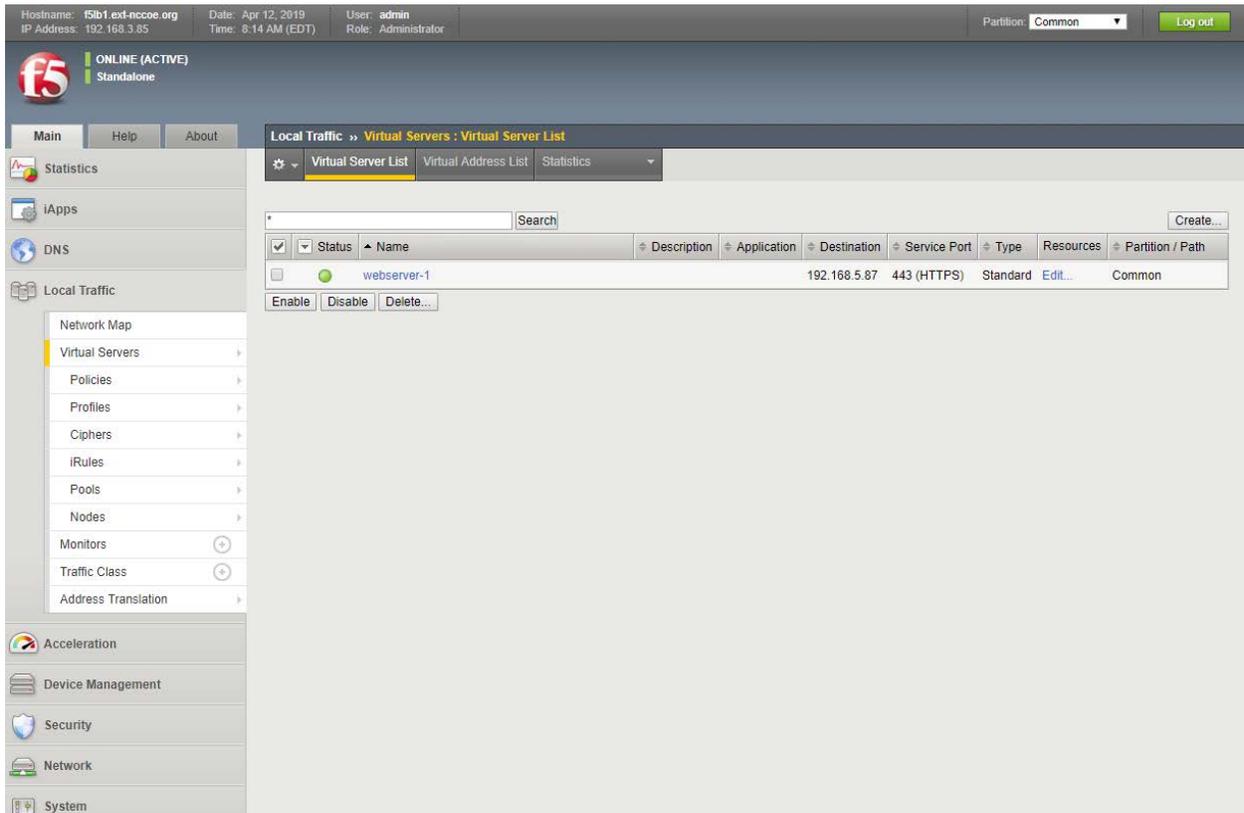
**Password**

Log in

Welcome to the BIG-IP Configuration Utility.

Log in with your username and password using the fields on the left.

(c) Copyright 1996-2017, F5 Networks, Inc., Seattle, Washington. All rights reserved.  
[F5 Networks, Inc. Legal Notices](#)

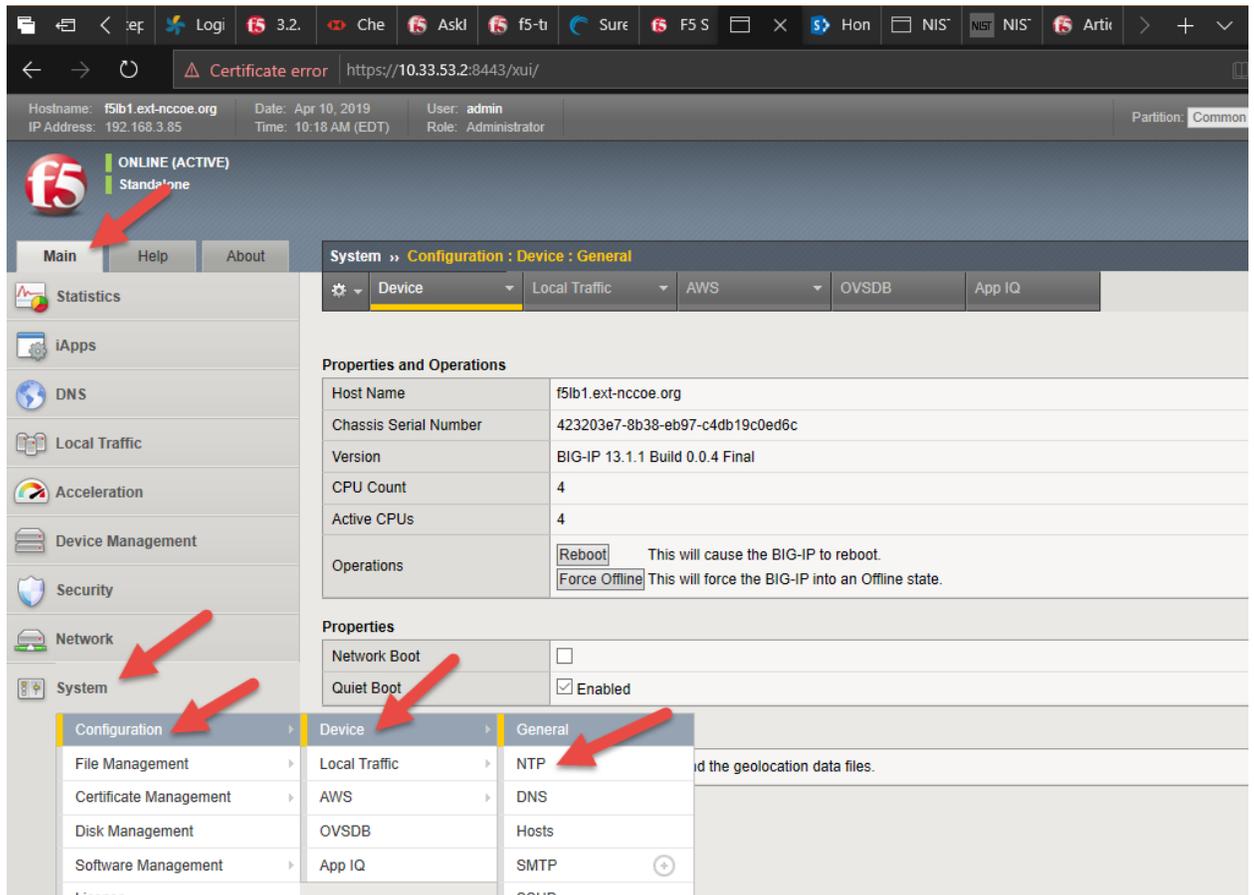


### 2.4.1.7 Configure NTP

Time synchronization is crucial when multiple BIG-IPs are in a cluster (not covered in this guide). It is also necessary for accuracy of logging information.

1. Log on to the Configuration utility.
2. Navigate to **Main > System**. Then click **Configuration > Device > NTP**.

The NTP panel appears.



3. In the Address field, type `time-a-g.nist.gov`. Click **Add**.
4. In the Address field, type `time-b-g.nist.gov`. Click **Add**.
5. Click **Update**.

#### 2.4.1.8 Configure SMTP

BIG-IP can be configured to send email alerts.

1. Navigate to **Main > System**. Then click **Configuration > Device > SMTP**.

The SMTP panel appears.

2. In the upper right corner, click the **Create** button.

The New SMTP Configuration panel appears.

3. Fill in these values.

|                       |                     |
|-----------------------|---------------------|
| Name                  | mail1               |
| SMTP Server Host Name | mail1.int-nccoe.org |
| Local Host Name       | f51b1-ext-nccoe.org |
| From Address          | f5-big-ip@nccoe.org |

4. Click **Finish**.

#### 2.4.1.9 *Configure Syslog*

Log events either locally on the BIG-IP system or remotely by configuring a remote syslog server.

1. Log on to the Configuration utility.
2. Navigate to **System > Logs > Configuration > Remote Logging**.
3. In Remote IP field, type 192.168.3.12.
4. Click **Add**.
5. Click **Update**.

#### 2.4.1.10 *Secure BIG-IP to NIST SP 800-53*

This section provides guidance on using the F5 iApp for NIST SP 800-53 (Revision 5) to configure a BIG-IP device to support security controls according to NIST SP 800-53 (Revision 4): *Security and Privacy Controls for Federal Information Systems and Organizations* (updated January 2, 2015).

Some controls (policies plus supporting technical measures) that organizations adopt by complying with NIST SP 800-53 (Revision 5) relate to the BIG-IP configuration.

This practice guide discusses the security controls in [Appendix F](#) of NIST SP 800-53 (Revision 5) [2] that apply to BIG-IP configuration and shows how to support them. It also focuses on configuring the management features of the BIG-IP system rather than the network-traffic-processing modules of a system such as BIG-IP Local Traffic Manager. This approach helps the user manage the BIG-IP system as an entity responsive to NIST SP 800-53 (Revision 5) controls. Using BIG-IP as a tool to help control other entities, such as network-based applications, is beyond the scope of this project.

#### 2.4.1.10.1 F5 iApp

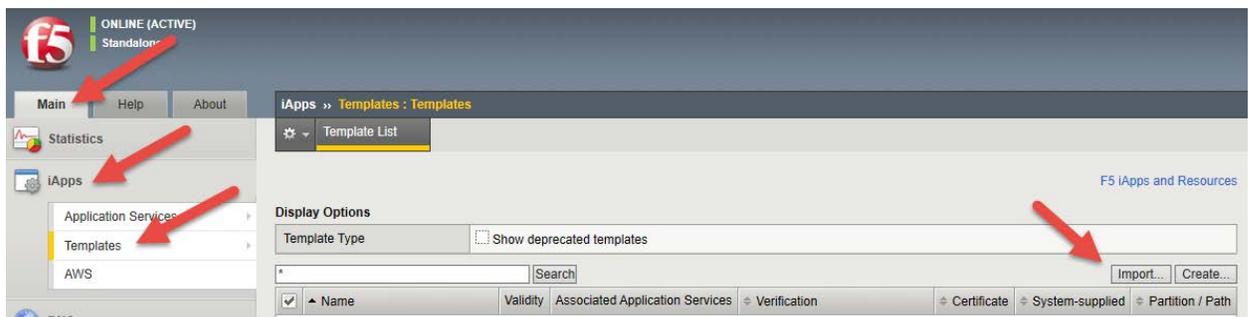
F5 iApp is a feature in the BIG-IP system that provides a way to simplify BIG-IP configurations. An iApp template brings together configuration elements, architectural rules, and a management view to deliver an application reliably and efficiently.

#### 2.4.1.10.2 Download the iApp for NIST SP 800-53 (Revision 5) Compliance

1. In a browser, open the F5 Downloads page at <https://downloads.f5.com>.
2. Log in with an F5 Support ID.
3. In the Downloads Overview page, click **Find a Download** button.
4. In the Select a Product Line page, under Product Line column, click **iApp Templates**.
5. In the Select a Product Version... page, click **iApp-Templates**.
6. Review the EULA, then click **I Accept**.
7. In the Select a Download page, click **iapps-1.0.0.546.0.zip**.
8. In the Download Locations page, click on the link nearest to the user's region.
9. Save the zip file to the local computer.

#### 2.4.1.10.3 Import iApp to BIG-IP

1. Unzip the downloaded file.
2. Open browser and navigate to the BIG-IP address <https://192.168.3.85>.
3. Log in as admin/admin.
4. On the left menu, click **Main > iApps > Templates**. Then on the right side, click **Import** button.



5. Browse to the file unzip location and to the subfolder **\iapps-1.0.0.546.0\Security\NIST\Release\_Candidates**. Select the file **f5.nist\_sp800-53.v1.0.1rc5.tmpl**, then click **Open**.
6. Click **Upload**.
7. On page 2 of the Template List, verify that the **f5.nist\_sp800-53.v1.0.1rc5** template has been uploaded.

#### 2.4.1.10.4 Deploy the NIST iApp

1. On the left menu, click **Main > iApps > Application Services**. Then on the right side, click **Create** button.

The Template Selection panel appears.

2. In the Name field, type `nist-800-53`.
3. In the Template pull-down, select **f5.nist\_sp800-53.v1.0.1rc5**.

The New Application Service panel appears.

iApps » Application Services : Applications » New Application Service...

Template Selection: Basic

Name: nist-800-53

Template: f5.nist\_sp800-53.v1.0.1rc5

Show deprecated templates

**Welcome to the BIG-IP NIST Special Publication 800-53r4 iApp Template f5.nist\_sp800-53.v1.0.1rc5**

**EARLY RELEASE** This template has not yet been fully tested at f5. It has limited support. When testing is complete it will be made available to all users.

**Introduction** This iApp helps you configure BIG-IP to support security controls consonant with NIST Special Publication 800-53r4 on management of the BIG-IP itself rather than control of application traffic through the BIG-IP. For more details on how BIG-IP supports NIST Special Publication 800-53r4, please consult the Deployment Guide or the Help tab (in the left-hand navigation pane).

Do you want to see inline help? No, do not show inline help

Should the iApp show blocks containing only advice? No, do not show advice-only blocks

**User Authentication/Directory Service -- AC-6, IA-2**

Configure authentication/directory service for BIG-IP management.

Which authentication/directory service do you want to use? Local to the BIG-IP system

**Password Strength Policy -- IA-5(1)**

Set local policy for password valid life, strength, and reuse. This policy governs local accounts (such as 'admin') and external user authentication/directory server.

Do you want to enforce custom local password policy? Yes, enforce a custom local password policy

4. Fill in the iApps with parameters in the following table. Leave everything else as default values.

|                                                             |                            |
|-------------------------------------------------------------|----------------------------|
| <b>Password Strength Policy—IA-5(1)</b>                     |                            |
| Do you want to enforce custom local password policy?        | "Yes, enforce a custom..." |
| How many days should pass before the password expires?      | 0                          |
| How many changes before reuse?                              | 0                          |
| How many characters should be the minimum for each setting? | Length = 8                 |
| <b>Maximum Failed Login Attempts—AC-7</b>                   |                            |

|                                                                                 |                               |
|---------------------------------------------------------------------------------|-------------------------------|
| Disable account after several failed login attempts?                            | "Yes, limit fail..."          |
| Allow how many consecutive login failures before disabling the account?         | 9                             |
| <b>NTP Configuration—AU-8(1,2)</b>                                              |                               |
| What is the IP address or FQDN of the primary NTP server?                       | time-a-g.nist.gov             |
| What is the IP address or FQDN of the first alternate NTP server?               | time-b-g.nist.gov             |
| <b>Syslog Configuration—AU-8, AU-9(2), AU-12(2)</b>                             |                               |
| Should log messages use International Standards Organization (ISO) date format? | "Yes, log messages..."        |
| Do you want to add syslog servers?                                              | "Yes, use this iApp..."       |
| Which syslog servers do you want to add?                                        | Server: syslog2.int-nccoe.org |

5. Click **Finished**.

## 2.4.2 Day 1: Product Integration Configuration

### 2.4.2.1 Prerequisites

- Venafi installed
- web servers for load balance

### 2.4.2.2 Venafi Integration

For information on integration with Venafi TPP, see Section [2.6.13.1](#).

### 2.4.2.3 Load Balance Web Servers

#### 2.4.2.3.1 Create a Pool to Manage https Traffic

A pool (a logical set of devices, such as web servers, that are grouped together to receive and process https traffic) can be created to efficiently distribute the load on the server resources.

1. On the Main tab, click **Local Traffic > Pools**.

The Pool List screen opens.

2. Click **Create**.

The New Pool screen opens.

3. In the Name field, type `app1_pool`.
4. For the Health Monitors setting, assign `https` by moving it from the Available list to the Active list.
5. Use the New Members setting to add each resource to include in the pool:
  - a. In the Address field, type `192.168.4.2`.
  - b. In the Service Port field type `443`.
  - c. Click **Add**.
6. Repeat step 5 for these three IP addresses.
  - a. `192.168.4.3`
  - b. `192.168.4.4`
  - c. `192.168.4.7`
7. Click **Finished**.

The `https` load balancing pool appears in the Pool List screen.

#### 2.4.2.3.2 Create Client SSL Profile

Profile for BIG-IP to decrypt traffic from browser

1. On the Main tab, click **Local Traffic > Profiles > SSL > Client**.

The SSL Client List screen opens.

2. Click **Create**.

The New Client SSL Profile screen opens.

3. In the Name field, type `app1_client-ssl`.
4. In the Certificate Key Chain setting, select the checkbox on the right. Then click **Add**.

The Add SSL Certificate to Key Chain screen opens.

5. For **Certificate** pull-down, select `app1.tls.nccoe.org-<value>`.

6. For **Key** pull-down, select app1.tls.nccoe.org-<value>.
7. Click **Add**.
8. Click **Finished**.

#### 2.4.2.3.3 Create Server SSL Profile

Profile for BIG-IP to encrypt traffic to web servers:

1. On the Main tab, click **Local Traffic > Profiles > SSL > Server**.  
The SSL Server List screen opens.
2. Click **Create**.  
The New Server SSL Profile screen opens.
3. In the Name field, type `app1_server-ssl`.
4. In the Certificate setting, select the checkbox on the right. Then select app1.tls.nccoe.org-<value> in the pull-down.
5. In the Key setting, select the checkbox on the right. Then select app1.tls.nccoe.org-<value> in the pull-down.  
The Add SSL Certificate to Key Chain screen opens.
6. For **Certificate** pull-down, select app1.tls.nccoe.org-<value>.
7. For **Key** pull-down, select app1.tls.nccoe.org-<value>.
8. Click **Finished**.

#### 2.4.2.3.4 Create a Virtual Server to Manage https Traffic

A virtual server can be specified to be either a host virtual server or a network virtual server to manage https traffic.

1. On the Main tab, click **Local Traffic > Virtual Servers**.  
The Virtual Server List screen opens.
2. Click the **Create** button.  
The New Virtual Server screen opens.
3. In the Name field, type `app1_vs`.
4. In the Destination Address field, type `192.168.5.85`.

5. In the Service Port field, type 443.
6. In the HTTP Profile setting, select **http** in the pull-down.
7. In the SSL Profile (Client) setting, from the Available list, select **app1\_client-ssl**, and click the  button to move over to the Selected list.
8. In the SSL Profile (Server) setting, from the Available list, select **app1\_server-ssl**, and click the  button to move over to the Selected list.
9. In the Source Address Translation setting, select **Auto Map** in the pull-down.
10. In the Default Pool setting, select **app1\_pool** in the pull-down.
11. In the Default Persistence Profile setting, select **cookie** in the pull-down.
12. Click **Finished**.

The https virtual server appears in the Virtual Server List screen.

#### 2.4.2.3.5 Create Redirect Virtual Server from http to https

When a user types *http://<virtual server>* in the browser, this virtual server redirects the user to the secure site *https://<virtual server>*.

1. On the Main tab, click **Local Traffic > Virtual Servers**.  
The Virtual Server List screen opens.
2. Click the **Create** button.  
The New Virtual Server screen opens.
3. In the Name field, type `app1_redir_vs`.
4. In the Destination Address field, type `192.168.5.85`.
5. In the Service Port field, type 80.
6. In the HTTP Profile setting, select **http** in the pull-down.
7. In the iRules setting, select **\_sys\_https\_redirect** in Available, and click the  button to move over to the Enabled list.
8. Click **Finished**.

The http redirect virtual server appears in the Virtual Server List screen.

## 2.4.3 Day N: Ongoing Security Management and Maintenance

### 2.4.3.1 Software Updates

BIG-IP VE updates in the same major version are installed in a similar manner as updates to BIG-IP software already installed on BIG-IP hardware. There is no need to reinstall BIG-IP VE in the hypervisor guest environment to upgrade the system. To update a BIG-IP VE virtual machine, use the Software Management tool in the Configuration utility, or upgrade the software from the command line. The update procedure described in this guide uses the Software Management tool.

#### 2.4.3.1.1 Download the Latest Software

Software release notes contain instructions for that specific installation.

*To find the latest software version for an F5 product:*

1. Navigate to F5 Downloads ([downloads.f5.com](https://downloads.f5.com)).
2. Click **Find a Download**.
3. Find the product desired for download, and click the link for the appropriate version.
4. Find and click the link for the update to download.
5. Read and accept the End User Software license agreement.
6. Click the file name, choose a download location, and save the file to the computer.

#### 2.4.3.1.2 Upgrading BIG-IP Software

Before upgrading the BIG-IP software, we recommend reviewing the release notes on AskF5 ([support.f5.com](https://support.f5.com)) in the Documentation section of the product and version. In particular, verify the new version supports the hardware, and carefully review these items:

- known issues list
- behavior change section(s)
- upgrading from earlier versions section
- upgrading from earlier configurations section
- installation checklist

#### 2.4.3.1.3 Import a BIG-IP VE Software Update

To install an update, BIG-IP software needs access to the ISO file previously downloaded.

1. Open browser, and navigate to the BIG-IP address *https://192.168.3.85*
2. Log in as an admin.

3. On the **Main** tab, click **System > Software Management**.

The *Software Management Image List* screen opens.

4. At the right side of the screen, click **Import**.

The *New Image* screen opens.

5. Click **Browse** to navigate to the downloaded installation file.
6. When the image name appears in the Software Image field, click **Import** to begin the operation.

The system presents a progress indicator during the operation.

#### 2.4.3.1.4 Installing a BIG-IP VE update

After import the software image, initiate the installation operation.

1. On the **Main** tab of the navigation pane, click **System > Software Management**.

The *Software Management Image List* screen opens.

2. From the *Available Images* table, select the software image you want to install.

The image properties screen opens.

3. Click **Install**.

The *Install Software* screen opens.

4. Select the disk you want to install the image on, and type or select a volume name, and click **Install**.

The upgrade process installs the software on the inactive disk location that you specify. This process usually takes between three and ten minutes.

Tip: If a problem arises during installation, use log messages to troubleshoot a solution. The system stores the installation log file as */var/log/liveinstall.log*.

5. The software image is installed.

#### 2.4.3.1.5 Reboot BIG-IP VE to update

When the installation operation is complete, you can safely reboot into the newly installed volume or partition.

1. On the **Main** tab of the navigation pane, click **System > Software Management**.

The *Software Management Image List* screen opens.

2. On the menu bar, click **Boot Locations**.

The *Boot Locations* screen opens.

3. In the *Boot Location* column, click the link representing the boot location you want to activate.

The properties screen for the boot location opens.

4. Click **Activate**.

A confirmation screen opens.

5. Click **OK** to initiate the reboot operation.

The system presents progress messages during the restart operation.

When the BIG-IP VE system reboot is complete, the system presents the login screen. To configure the system, log in using an account that has administrative permissions.

### 2.4.3.2 License and Entitlement

If support is purchased from F5, it is associated with a particular BIG-IP system. A system with an active support contract is considered entitled until the contract expires. To continue receiving support, the contact must be renewed.

Licenses are also associated with modules purchased to run a specific system. Model licenses are considered add-ons to the main license for a system, and are automatically linked to the main BIG-IP system license and eligible for technical support if that system is entitled.

Major software upgrades are only supported for entitled systems and require relicensing of the BIG-IP system. Minor upgrades do not require relicensing.

#### 2.4.3.2.1 Viewing and verifying a BIG-IP system license

Test the validity of the BIG-IP software license by obtaining license information in any of the following ways:

- view license information at the command line
- request a product license profile from F5
- view license profile in BIG-IP iHealth®
- view license profile in the Configuration utility
- At the command line, type the following command: `tmsh show /sys license`

Output displays licensing information for the BIG-IP system should include a list of active modules. For a system with a valid license, output appears similar to the following example:

#### 2.4.3.2.2 Provisioning licenses

If a license is installed for an add-on module on a BIG-IP system, you must provision resources for the module.

Until provisioned, module function is limited in the following ways:

- the system does not perform the functions of the licensed module
- items related to the module do not appear in Configuration utility menus
- the TMOS Shell (tmsh) does not present or permit configuration of objects related to the module.
- the `bigstart status` command returns output similar to the following example for daemons related to the unprovisioned module: `<daemon_name> down, Not provisioned` For information on provisioning modules, refer to “Modules.”

When you upgrade a BIG-IP system, the install script verifies the Service Check Date with the license check date of the version being installed. If the service check date is missing or the verification process finds your license pre-dates the software’s release date, a line displays in the `/var/log/liveinstall.log` with a note about the service check date verification, and the installation of the software may continue.

#### 2.4.3.2.3 Reactivating a BIG-IP System License

F5 recommends reactivating the BIG-IP system license before conducting a software upgrade.

Follow these steps to reactivate a BIG-IP system license using the Configuration utility:

1. Navigate to System > License.
2. Click **Re-activate**.
3. In the Activation Method area, select **Automatic** (requires outbound connectivity).
4. Click **Next**.

#### 2.4.3.2.4 Moving a BIG-IP VE license

BIG-IP VE licenses are permanently associated with the virtual instance. To move a license, contact F5 Technical Support for assistance. However, with BIG-IP 12.1.3.3 and BIG-IP 13.1 and later, you can move the RegKey without contacting support by revoking the instance’s license from tmsh, the Configuration utility, and iControl/REST by using the ‘`tmsh revoke sys license`’ command on that virtual instance. This action revokes the license and unlocks the RegKey—enabling the user to activate a new virtual machine.

Call F5 Technical Support for assistance if the connection is lost and you want to move the license to the current VE, if hypervisor crashes, or if you can’t access the password or network address.

### 2.4.3.3 Backup and Data Recovery

BIG-IP software offers two supported methods for backing up and restoring the configuration: user configuration set (UCS) archives and single configuration files. This guide focuses on using the UCS archive only. To create, delete, upload, or download an archive, you must have either administrator or resource administrator role privileges.

#### 2.4.3.3.1 Backup Configuration Data to a UCS Archive

A UCS archive contains BIG-IP configuration data that can fully restore a BIG-IP system in the event of a failure or return material authorization.

Each time you back up the configuration data, the BIG-IP system creates a new UCS archive file in the `/var/local/ucs` directory. In addition to configuration data, each UCS file contains various configuration files necessary for the BIG-IP system to operate correctly.

A UCS archive contains the following types of BIG-IP system configuration data:

- system-specific configuration files (traffic management elements, system and network definitions, and others)
- product licenses
- user accounts and password information
- DNS
- zone files
- installed SSL keys and certificates

To easily identify the file, include the BIG-IP host name and current time stamp as part of the file name.

F5 recommends keeping a backup copy of the UCS archives on a secure remote server. To restore the BIG-IP system if you can't access the `/var/local/ucs` directory on the BIG-IP system, upload the backup file from the remote server, and use it to restore your system.

#### 2.4.3.3.2 To create a UCS archive using the Configuration utility

When creating a new archive, unless otherwise directed, the BIG-IP system automatically stores it in `/var/local/ucs` directory—a default location. You can create as many archives as you want, but each archive must have a unique file name.

All boot locations on a BIG-IP system use the same `/shared` directory, making it a good choice for a UCS save location. Saving an archive to the `/shared` directory allows you to boot to another boot location and access the archive, and can greatly simplify the recovery from a variety of issues.

1. Navigate to **System > Archives**.

2. Click **Create**.
3. Type a unique file name.
4. To encrypt the archive for Encryption, click **Enabled**.
5. To include private keys in the BIG-IP system, for Private Keys, click **Include**. If you choose to include private keys, store the archive file in a secure environment.
6. Click **Finished**.
7. Click **OK** after the data is backed up and the file is created.

#### 2.4.3.3.3 To download and copy an archive to another system using the Configuration utility

1. Navigate to **System > Archives**.
2. Click the UCS file name you want to download.
3. In Archive File, click Download <filename>.ucs.
4. Save the file.
5. Find the file in your computer's Downloads folder and copy it.

#### 2.4.3.3.4 Restoring Configuration Data from a UCS Archive

If the BIG-IP System configuration data becomes corrupted, you can restore the data from the archive currently stored in the directory `/var/local/ucs`.

When restoring configuration data, F5 recommends running the same version of the BIG-IP software on the BIG-IP system from which it was backed up.

F5 also recommends restoring a UCS file to another platform of the same model where the UCS file was created. Certain core hardware changes can cause a UCS to load properly on dissimilar hardware, requiring manual intervention to correct.

#### 2.4.3.3.5 To restore a configuration in a UCS archive using the Configuration utility

1. Navigate to **System > Archives**.
2. Click the name of the UCS archive you want to restore.
3. To initiate the UCS archive restore process, click **Restore**.

When the restoration process is completed, examine the status page for any reported errors before proceeding to the next step.

4. To return to the Archive List page, click **OK**.

If you receive activation errors after restoring a UCS archive on a different device, you must reactivate the BIG-IP system license. Restarting the system ensures that the configuration is fully loaded after relicensing,

#### 2.4.3.3.6 Downloading a UCS Archive to a Remote System

Downloading a copy of an existing archive to a remote system protects the configuration data should you need to restore your BIG-IP system and be unable to access the `/var/local/ucs` directory on the BIG-IP system.

To download an existing archive, first display the properties of the archive to specify the complete path name of the location where you want to save the archive copy.

1. Navigate to **System > Archives**.
2. Click the name of the archive that you want to view.

The General Properties for that archive display.

3. Click **Download**: `<ucs filename>`.
4. Click **Save**.

The BIG-IP system downloads a copy of the UCS file to the system from which you initiated the download.

#### 2.4.3.3.7 Uploading a UCS Archive from a Remote System

If a UCS archive on your BIG-IP system is unavailable or corrupted, upload a previously created archive copy from a remote or backup system to replace it.

1. Navigate to **System > Archives**.
2. Click **Upload**.
3. Type the complete path and file name of the archive that you want to upload onto the BIG-IP system.

If you do not know the path or file name, click **Browse** and navigate to the location.

4. Click **Upload**.

The specified archive uploads to the `/var/local/ucs` directory on the BIG-IP system.

#### 2.4.3.3.8 Deleting a UCS Archive

Use the Configuration utility to delete any archive on the BIG-IP system that is stored in the directory `/var/local/ucs`.

1. Navigate to **System > Archives**.
2. Select the check box next to the name of the file you want to delete.
3. Click **Delete**.
4. Click **Delete** again.

The archive is deleted from the `/var/local/ucs` directory on the BIG-IP system.

#### 2.4.3.4 *Log Files and Alerts*

This section provides context for our recommended procedures in the form of overviews and supplemental information, including the following topics:

- Config for Syslog
- Set up SMTP for email alerts

##### 2.4.3.4.1 *Managing Log files on a BIG-IP System*

Log files track usage or troubleshoot issues—if left unmanaged, they can grow to an unwieldy size. The BIG-IP system uses a utility called logrotate to manage local log files. The logrotate script deletes log files older than the number of days specified by the Logrotate.LogAge database variable. By default, the variable is set to eight. Therefore, the system is configured to delete archive copies that are older than eight days.

To modify the Logrotate.LogAge database variable:

1. Log in to tmsh at the command line by typing the following command: `tmsh`
2. Modify the age at which log files are eligible for deletion by using the following command  
syntax: `modify /sys db logrotate.logage value <value 0 - 100>`
3. Save the change by typing the following command: `save /sys config`

##### 2.4.3.4.2 *Audit Logging*

Audit logging is an optional way to log messages pertaining to configuration changes that users or services make to the BIG-IP system configuration. Audit logging is also known as master control program.

#### LOG FILES AND ALERTS—PROCEDURES

(MCP) Audit Logging. As an option, you set up audit logging for any tmsh commands that users type on the command line.

For MCP and tmsh audit logging, select a log level. The log levels will not affect the severity of the log messages but may affect the initiator of the audit event.

### 2.4.3.5 *Technical Support*

In addition to Support Centers around the world, there are many technical resources available to customers.

#### 2.4.3.5.1 Phone Support

Open a Case at any of the Network Support Centers:

- 1-888-882-7535 or (206) 272-6500
- International contact numbers: <http://www.f5.com/training-support/customer-support/contact/>

#### 2.4.3.5.2 AskF5 - Web Support

F5 self-support portal: <http://www.askf5.com>

#### 2.4.3.5.3 DevCentral - F5 User Community

More than 360,000 members—including F5 engineering resources—are actively contributing, sharing and assisting our peers.

<http://devcentral.f5.com>

#### 2.4.3.5.4 BIG-IP iHealth

BIG-IP iHealth comprises BIG-IP iHealth Diagnostics and BIG-IP iHealth Viewer. BIG-IP iHealth Diagnostics identifies common configuration problems and known software issues. It also provides solutions and links to more information. With BIG-IP iHealth Viewer, you can see the status of your system at-a-glance, drill down for details, and view your network configuration.

<https://ihealth.f5.com/>

#### 2.4.3.5.5 Subscribing to TechNews

AskF5 Publications Preference Center provides email publications to help keep administrators up-to-date on various F5 updates and other offerings:

- TechNews Weekly eNewsletter Up-to-date information about product and hotfix releases, new and updated articles, and new feature notices.
- TechNews Notifications Do you want to get release information, but not a weekly eNewsletter? Sign up to get an HTML notification email any time F5 releases a product or hotfix.
- Security Alerts Receive timely security updates and ASM attack signature updates from F5.

To subscribe to these updates:

1. Go to the Communications Preference Center (<https://interact.f5.com/F5-Preference-Center.html>).
2. Under My preferences click **Show**.
3. Select the updates you want to receive.
4. Click **Submit**.

#### 2.4.3.5.6 AskF5 recent additions and updates

You can subscribe to F5 RSS feeds to stay informed about new documents pertaining to your installed products or products of interest. The Recent additions and updates page on AskF5 provides an overview of all the documents recently added to AskF5.

New and updated articles are published over RSS. You can configure feeds that pertain to specific products, product versions, and/or document sets. You can also aggregate multiple feeds into your RSS reader to display one unified list of all selected document.

## 2.5 Symantec SSL Visibility Appliance

The Symantec SSL Visibility appliance is a high-performance transparent proxy for SSL network communications. It enables a variety of applications to access the plaintext (that is, the original unencrypted data) in SSL encrypted connections, and is designed for security and network appliance manufacturers, enterprise IT organizations, and system integrators. Without compromising any aspect of enterprise policies or government compliance, the SSL Visibility appliance permits network appliances to deploy with highly granular flow analysis while maintaining line rate performance.

### 2.5.1 Day-0: Install and Standard Configuration

#### 2.5.1.1 Prerequisites

- 120V or 220V Power Source
- computer with browser access to activate license and configure appliance
- putty or a terminal emulator
- four-post equipment rack with a depth of 27.75" to 37.00" with square mounting holes
- category 5E network cables or better (Category 6 or 6A)
- license key for SSL Visibility appliance
- Broadcom account

- DNS Server
- SSL VISIBILITY running version 3.X

### 2.5.1.2 *Unpacking the Appliance*

Before racking and configuring the SSL Visibility Appliance, ensure the following contents are included in the SSL Visibility shipping package:

|                                               | SV800 | SV1800 | SV2800 | SV3800 |
|-----------------------------------------------|-------|--------|--------|--------|
| External power supply with AC power cord      | √     |        |        |        |
| Two AC power cords                            |       | √      | √      | √      |
| Rack-mount rail kit                           |       | √      | √      | √      |
| Rack-mount ears with fasteners                |       | √      | √      | √      |
| <i>Safety and Regulatory Compliance Guide</i> | √     | √      | √      | √      |
| <i>Quick Start Guide</i> (this document)      | √     | √      | √      | √      |
| Software License Agreement                    | √     | √      | √      | √      |
| Hardware Warranty                             | √     | √      | √      | √      |

### 2.5.1.3 *Rack-Mount the Appliance*

The list below shows the requirements to install the SSL Visibility Appliance.

- At least 1U rack space (deep enough for a 27" device)—power and management ports at rear
- Phillips (cross head) screwdriver
- Weight Capacity: 28 lb (12.7 kg)
- Dimensions: 17.5" (W) x 19.5" (D) x 1.75" (H) (444.5 mm x 495.3 mm x 44.5 mm)
- Two available power outlets (110 VAC or 220-240 VAC)
- Two IEC-320 power cords (normal server/PC power cords) should the supplied power cords not be suitable for your environment
- Cooling for an appliance with two 450W power supply units

To see detailed instructions for installing the SSL Visibility in a rack, please refer to Symantec's Quick Start guide located at the below link:

[https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/ssl-visibility/4-5/Getting\\_Started/initial\\_config.html](https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/ssl-visibility/4-5/Getting_Started/initial_config.html)

### 2.5.1.4 *Connect Cables*

To connect the appliance's cables:

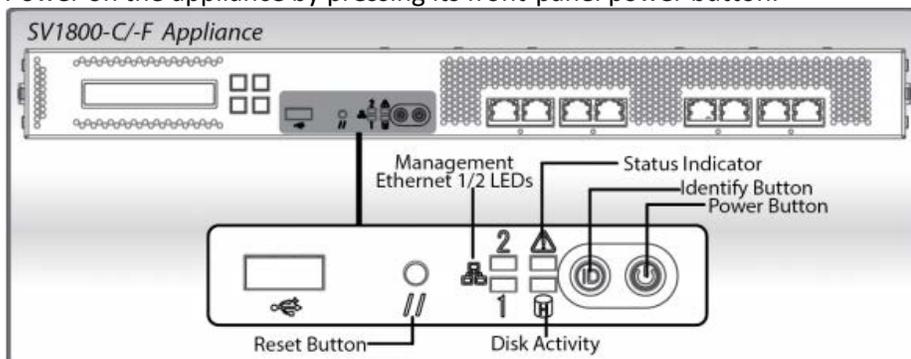
1. Connect a network cable between the **Management Ethernet 1** port, on the rear of the SSL VISIBILITY appliance, and Datacenter Secure network.

**Warning:** When deploying the SV1800, SV2800, and SV3800 appliances, do not connect to the Management Ethernet 2 port. This port is not functional.

2. Connect the two AC power cords to the appliance's AC power inlets on the rear panel. Two power supplies are provided for redundant operation.
3. Connect the other ends of the power cords to a 120 V or 220 V power source.

### 2.5.1.5 *Power on the Appliance and Verify LEDs*

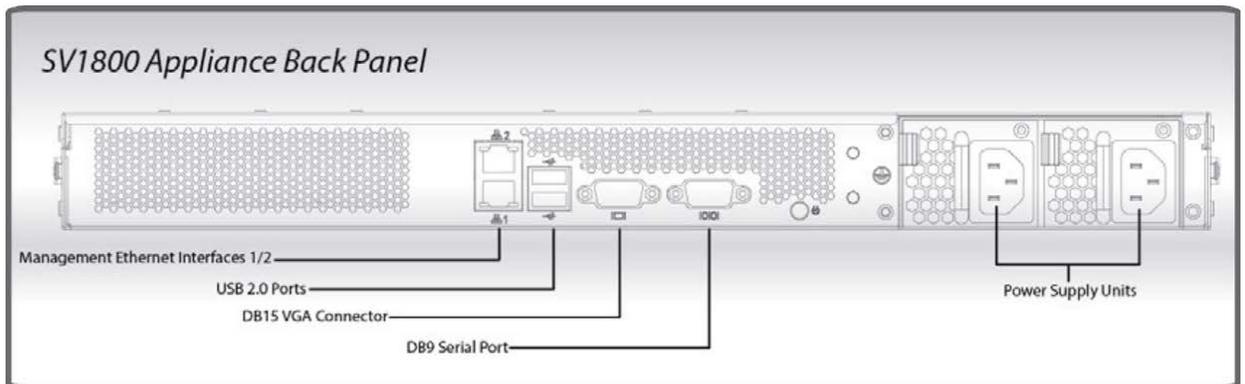
1. Confirm the appliance's power cord or power cords are securely connected to a 120 V or 220 V power source.
2. Power on the appliance by pressing its front-panel power button.



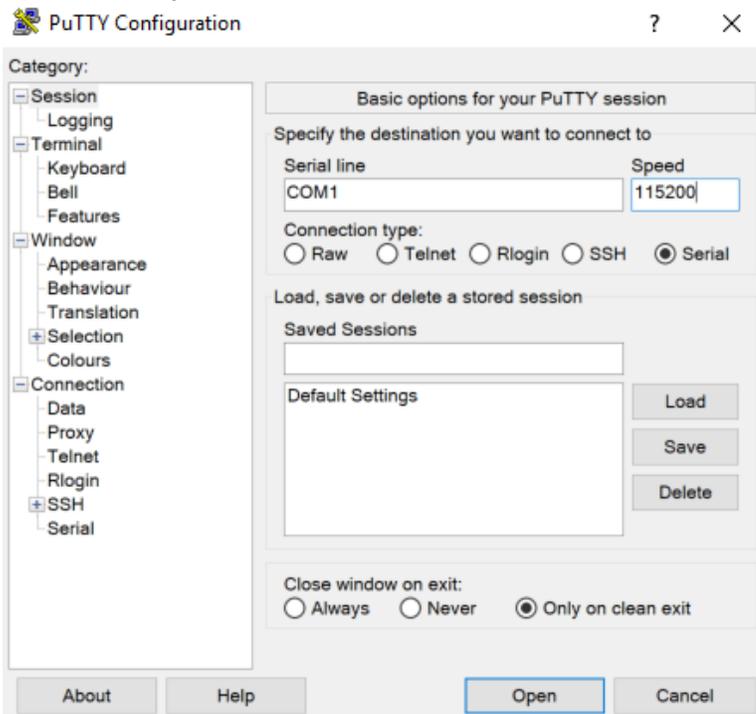
3. As the appliance boots verify the following:
  - The LCD displays startup messages while the appliance boots (Appliance Startup, Validating Firmware, Appliance Boot, etc.).
  - The System Status indicator for the SV1800 changes from red to off.
  - The LEDs for the Management Ethernet port (connected to a management workstation) light up.
  - When the boot process is complete, the LCD displays the appliance's model, software version, and the Up/Down arrows.

### 2.5.1.6 *Initial Appliance Configuration*

1. To perform initial configuration of the SSL Visibility Appliance, connect a serial cable to the **DB9 Serial port** on the rear of the Appliance.

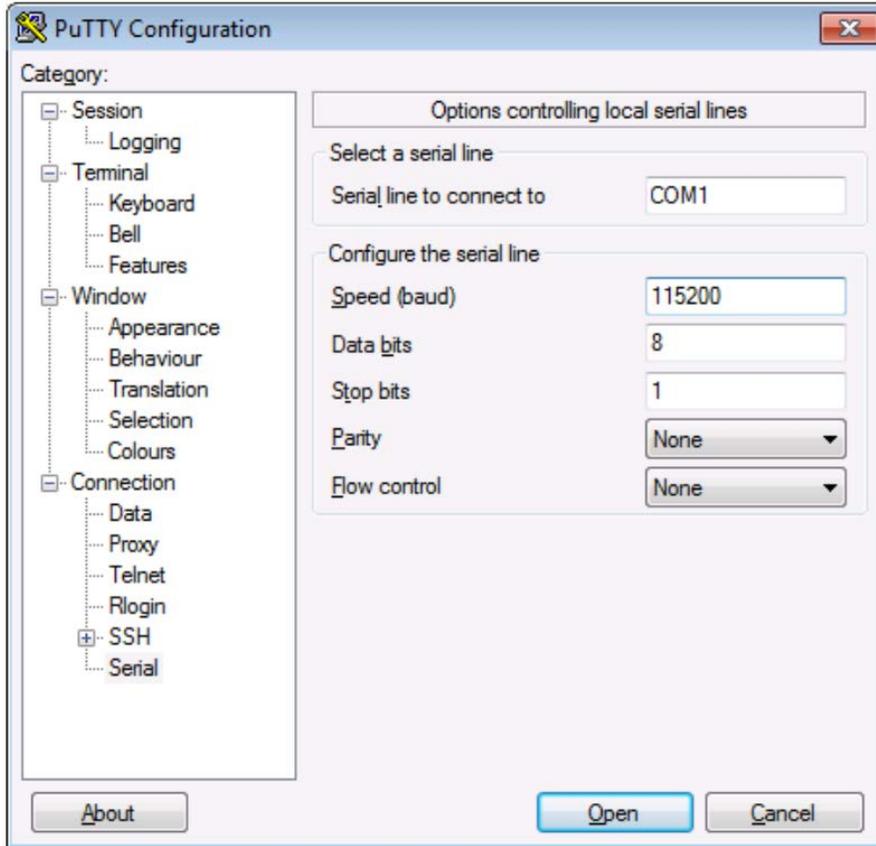


2. On the management laptop, open up the PuTTY Application and select a **Connection type** of **Serial** with a **Speed** of **115200**.

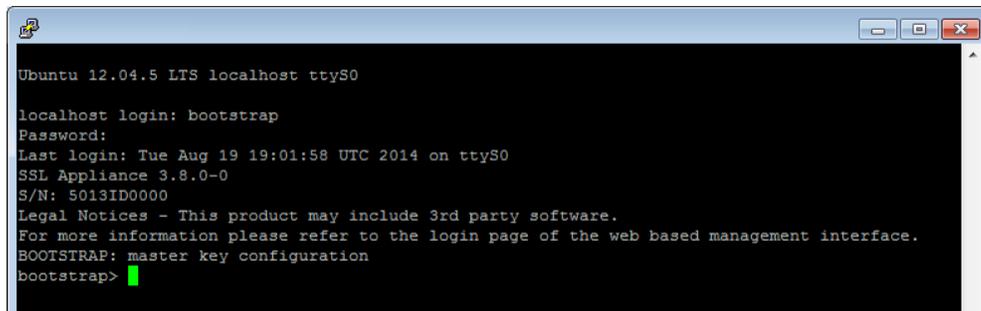


3. Navigate to the **Serial** Category on the bottom left side of the window.
4. Configure the serial connection to support the SSL Visibility Appliance's console speeds by selecting the following options:
  - **Speed (baud): 115200**
  - **Data bits: 8**
  - **Stop bits: 1**

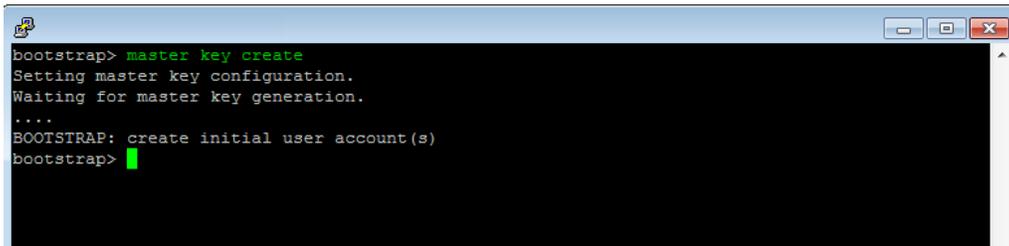
- **Parity: None**
- **Flow Control: None**



5. Login into the appliance by using the default credentials of:
  - **Username: bootstrap**
  - **Password: bootstrap**

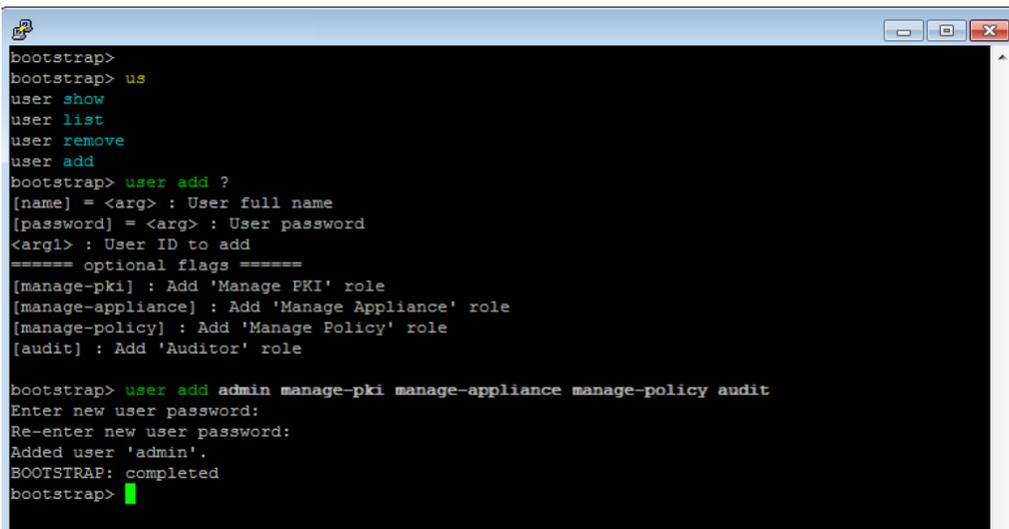


6. Next, create the master key by running the command:  
master key create



```
bootstrap> master key create
Setting master key configuration.
Waiting for master key generation.
....
BOOTSTRAP: create initial user account(s)
bootstrap>
```

7. Create a new user by running the command:  
user add admin manage-pki manage-appliance manage-policy audit



```
bootstrap>
bootstrap> us
user show
user list
user remove
user add
bootstrap> user add ?
[name] = <arg> : User full name
[password] = <arg> : User password
<arg1> : User ID to add
===== optional flags =====
[manage-pki] : Add 'Manage PKI' role
[manage-appliance] : Add 'Manage Appliance' role
[manage-policy] : Add 'Manage Policy' role
[audit] : Add 'Auditor' role

bootstrap> user add admin manage-pki manage-appliance manage-policy audit
Enter new user password:
Re-enter new user password:
Added user 'admin'.
BOOTSTRAP: completed
bootstrap>
```

Tip: This step created a single admin user account with all four roles allocated to it. The only requirements for completing the bootstrap phase are that there is a user account with the Manage Appliance role and a user account with the Manage PKI role. These may be the same or different accounts. In most cases, creating a single account with all four roles is the simplest approach.

8. Run the following command to configure the management network interface with a static IP address:  
network set ip 192.168.1.95 netmask 255.255.255.0 gateway 192.68.1.1
9. Reboot the system for the changes to take effect (confirm that you wish to reboot) with the following command: platform reboot

```
admin>
admin> platform reboot
Reboot appliance? (enter 'yes' to confirm): yes
```

10. On reboot, confirm that the “SSL Visibility startup stage 3: CONFIRMED” is displayed as shown below.

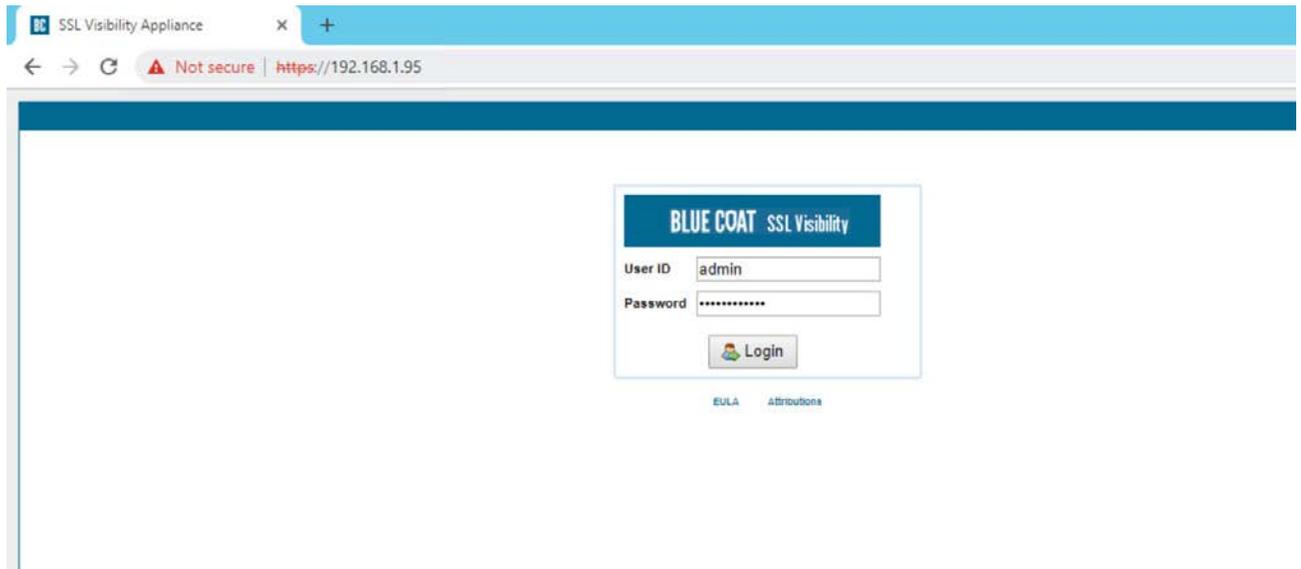
```
fscck from util-linux 2.20.1
data: clean, 60/3489792 files, 266044/13950976 blocks
The disk drive for /var/log is not ready yet or not present.
Continue to wait, or Press S to skip mounting or M for manual recovery
fscck from util-linux 2.20.1
fscck from util-linux 2.20.1
fscck from util-linux 2.20.1
coredump: clean, 11/436320 files, 63995/1743872 blocks
ui: clean, 208/65536 files, 17408/262144 blocks

log: clean, 41/262144 files, 51515/1048576 blocks
* Using makefile-style concurrent boot in runlevel S
* Using makefile-style concurrent boot in runlevel 2
SSLV startup stage 1: housekeeping
* Starting NTP server ntpd [OK]
* Loading cpufreq kernel modules... [OK]
* CPU0... * CPU1...
* CPU2...
* CPU3...
* CPUFreq Uti

lities: Setting ondemand CPUFreq governor... [OK]

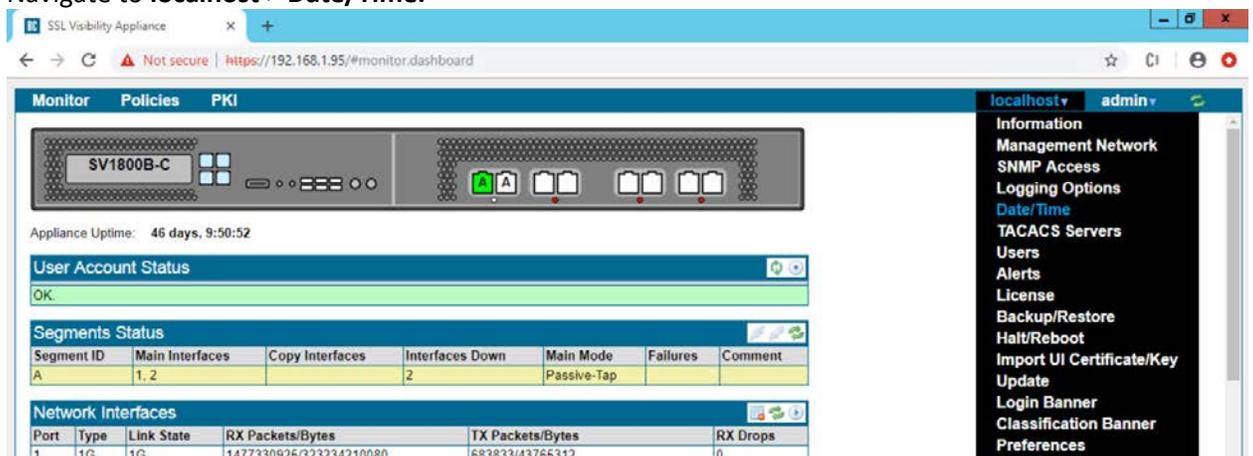
localhost login: Verified OK
Verified OK
Verified OK
Validating firmware...
NFE is up-to-date
BIOS is up-to-date
SSLV startup stage 3: CONFIRMED
```

11. Confirm you can log in to the appliance via your browser. Log in via a web browser, using the format `https://192.168.1.95`. Log in with the username and password you created.



### 2.5.1.7 Date and Time (NTP)

1. To configure Date and Time, login into the WebUI by browsing to <https://192.168.1.95>.
2. Navigate to **localhost > Date/Time**.



3. Click on the Add button  under NTP Servers.
4. In the server field type `time.nist.gov` and click **OK**.

5. Click **Apply Changes** to save the new NTP server.

### 2.5.1.8 Additional Configuration

To add a host name and DNS for the SSL Visibility Appliance, perform the following steps:

1. Log in to the SSL Visibility by opening a web browser and navigating to *https://192.168.1.95*.
2. From the **Dashboard** page navigate to **localhost > Management Network**.

| Segment ID | Main Interfaces | Copy Interfaces | Interfaces Down | Main Mode   | Failures | Comment |
|------------|-----------------|-----------------|-----------------|-------------|----------|---------|
| A          | 1, 2            |                 | 2               | Passive-Tap |          |         |

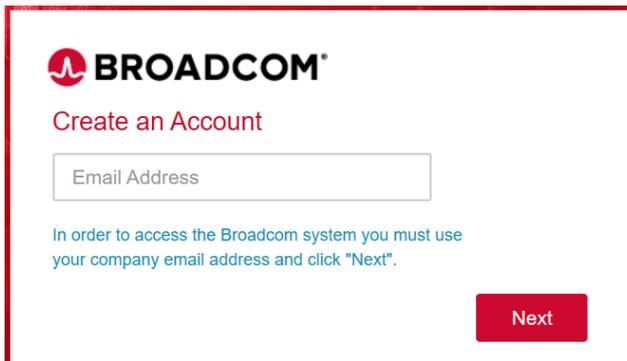
| Port | Type | Link State | RX Packets/Bytes        | TX Packets/Bytes        | RX Drops |
|------|------|------------|-------------------------|-------------------------|----------|
| 1    | 1G   | 1G         | 1477342332/323238764805 | 583835/43765440         | 0        |
| 2    | 1G   | Down       | 8589/551865             | 1485232670/316784587304 | 0        |
| 3    | 1G   | Unknown    | 0/0                     | 1280811088/236683069790 | 0        |
| 4    | 1G   | Unknown    | 0/0                     | 0/0                     | 0        |
| 5    | 1G   | Unknown    | 0/0                     | 0/0                     | 0        |

3. Click the **Edit** button  under the **Management Network** Field.
4. Enter the following information into the fields:
  - **MTU: 1500**
  - **Host Name: SSL Visibility.int-nccoe.org**
  - **Primary Nameserver: 192.168.1.6**

5. Click **Apply Changes**.
6. Click **Reboot** to restart the system and apply changes (required).

### 2.5.1.9 *Broadcom Account Creation*

1. To create a Broadcom Account, navigate to the following link:
2. <https://portal.broadcom.com/web/guest/registration?source=CA>
3. Enter the requested information and click **Next**.



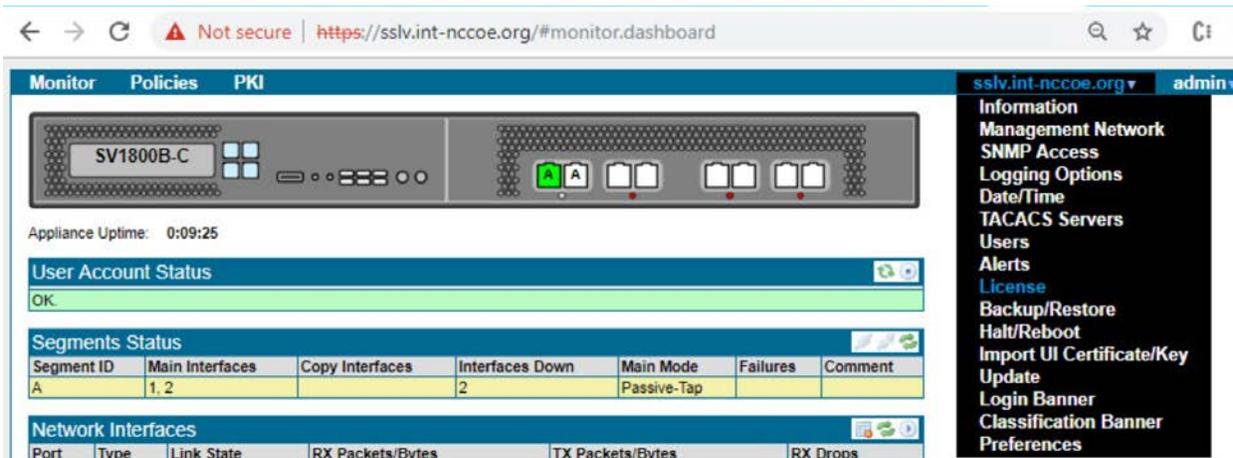
### 2.5.1.10 *License the SSL Visibility Appliance*

#### 2.5.1.10.1 Download a Blue Coat License

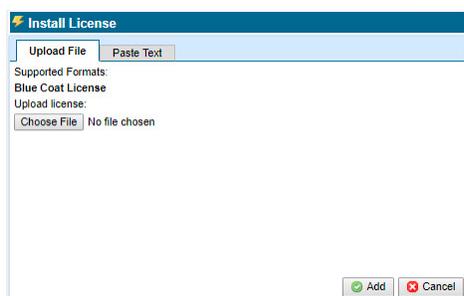
1. Using your BlueTouch Online account, log in to the Blue Coat Licensing Portal.  
([https://services.bluecoat.com/eservice\\_enu/licensing/register.cgi](https://services.bluecoat.com/eservice_enu/licensing/register.cgi)).
2. From the menu on the left side, select **SSL Visibility**, then select **License Download**.
3. When prompted, enter the serial number of your appliance, then press **Submit**.
4. Once the license is generated, press **Download License File** for the required SSL Visibility Appliance.

#### 2.5.1.10.2 Install a Blue Coat License

1. Select **SSL Visibility.int-nccoe.org > License**.



2. Click the **Add** button  in the **License** field.
3. On the **Upload File** tab, use the **Choose File** button to browse to the license file location.



4. Click **Add**. You will see a confirmation message and the specific appliance platform model. The license is now installed, and all standard SSL Visibility Appliance features are operational.

## 2.5.2 Day 1: Product Integration Configuration

### 2.5.2.1 Prerequisites

1. Install version 3.x on the SSL Visibility Appliance.
2. Complete initial configuration as outlined in the Day 0 Section [2.5.1](#) above.
3. Required Ports, Protocols and Services:  
SSL Visibility 3.x uses the following ports while operating—allow these ports when setting up SSL Visibility:  
Inbound Connection to SSL Visibility Appliance

Table 18

| Service                           | Port | Protocol | Configurable | Source         | Description                                         |
|-----------------------------------|------|----------|--------------|----------------|-----------------------------------------------------|
| WebUI Admin GUI                   | 443  | TCP      | No           | User client    | Management Interface WebUI service                  |
| SSH Admin CLI                     | 22   | TCP      | No           | User client    | SSH Admin CLI service                               |
| Symantec/Blue Coat License        | 443  | HTTPS    | No           | License server | Symantec/Blue Coat license service                  |
| SNMP management                   | 161  | UDP      | No           | User client    | SNMP agent for SNMP management access               |
| NTP                               | 123  | UDP      | No           | NTP server     | NTP time synchronization service                    |
| DHCP                              | 68   | UDP      | No           | DHCP server    | DHCP service                                        |
| Remote Diagnostics Facility (RDF) | 2024 | TCP      | No           | RDF            | Can be opened for support requests; normally closed |

Outbound Connections from SSL Visibility Appliance

Table 19

| Service          | Port                          | Protocol          | Configurable | Destination   | Description          |
|------------------|-------------------------------|-------------------|--------------|---------------|----------------------|
| SMTP/Secure SMTP | 25, 465, 587, 525, 2526 *     | TCP               | Yes          | SMTP server   | SMTP alerts          |
| Syslog           | 514, 601 *<br>6514 *<br>514 * | TCP<br>TLS<br>UDP | Yes          | Syslog server | Remote syslog server |

|                            |     |            |     |                    |                                                   |
|----------------------------|-----|------------|-----|--------------------|---------------------------------------------------|
| DNS                        | 53  | TCP<br>UDP | No  | DNS server         | Domain Name System service                        |
| SNMP Trap                  | 162 | UDP        | No  | SNMP Trap receiver | SNMP traps                                        |
| Host Categorization (BCWF) | 443 | HTTPS      | No  | Symantec           | Host categorization database                      |
| HSM                        | 443 | HTTPS      | No  | HSM appliance      | HSM authentication and requests                   |
| TACACS+                    | 49  | TCP        | Yes | TACACS server      | TACACS+ authentication                            |
| NTP                        | 123 | UDP        | No  | NTP server list    | Synchronization to customer-configured NTP server |
| DHCP                       | 67  | UDP        | No  | DHCP server        | DHCP service                                      |
| Diagnostics Upload         | 443 | HTTPS      | No  | Symantec           | Diagnostics upload service                        |

\*Common Values For this Port

Required URLs

Ensure connectivity from SSL Visibility to the following URLs:

Table 20

| URL                    | Port | Protocol     | Description                                    |
|------------------------|------|--------------|------------------------------------------------|
| abrca.bluecoat.com     | 443  | HTTPS<br>TCP | Symantec CA                                    |
| *.es.bluecoat.com      | 443  | HTTPS<br>TCP | License, validation, and subscription services |
| appliance.bluecoat.com | 443  | HTTPS<br>TCP | Trust package downloads                        |
| upload.bluecoat.com    | 443  | HTTPS<br>TCP | Upload diagnostic reports to Symantec support  |

### 2.5.2.2 Venafi Integration

Venafi TPP was used to copy known server key and certificates to the SSL Visibility appliance for TLS decryption.

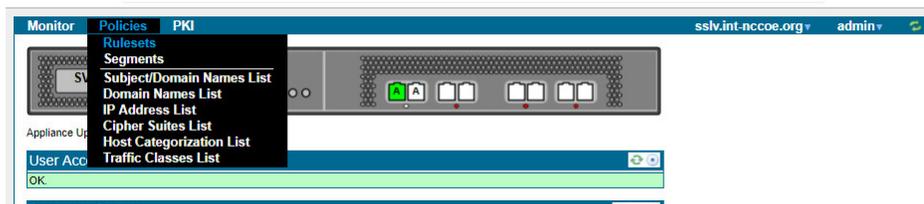
For information on integration with Venafi TPP, see Section: [2.6.13.9](#).

### 2.5.2.3 Ruleset Creation

To ensure your SSL Visibility Appliance is connected and configured properly, create a basic ruleset to test that traffic isn't getting blocked. To perform this test, create a ruleset with a Catch All Action of Cut Through.

Note: At least one rule must be added to the ruleset for SSL Visibility Appliance to start processing SSL traffic.

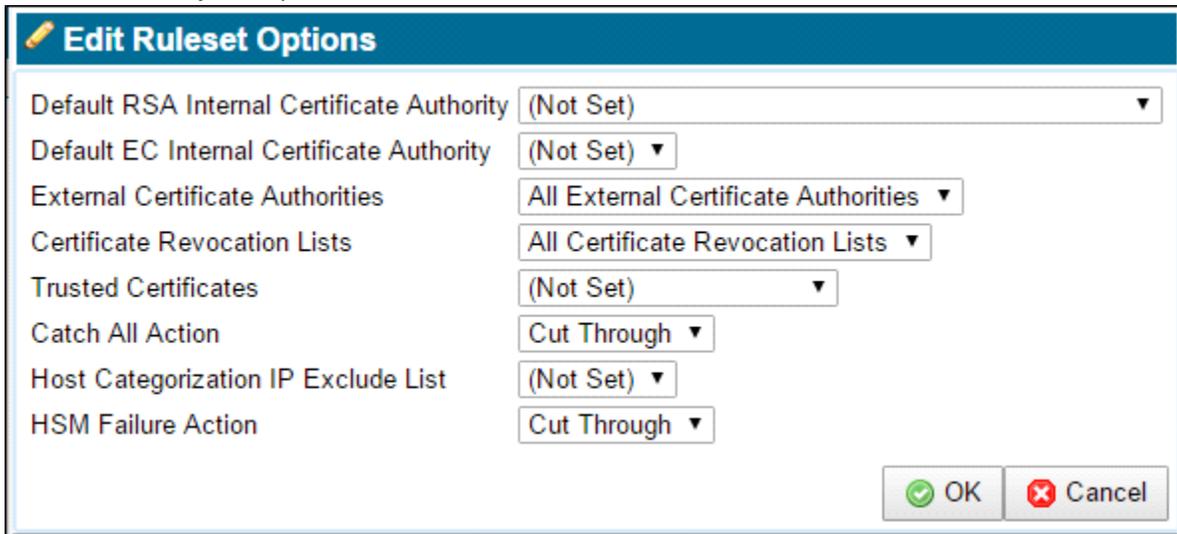
1. Select **Policies > Rulesets**.



2. In the **Rulesets** panel, click the **Add**  icon.
3. In the **Add Ruleset** window, enter a name for the ruleset and click **OK**.



- In the **Ruleset Options** panel, click the **Edit**  icon.

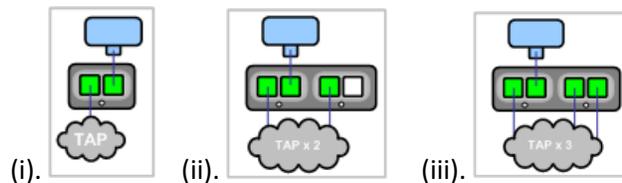


- Confirm the **Catch All Action** is **Cut Through**.
- Apply** the Policy Changes.

#### 2.5.2.4 *Segment Creation*

Note: Before creating the segment, determine your deployment mode and create a ruleset for the segment.

The following pictures demonstrate various passive tap deployment types:

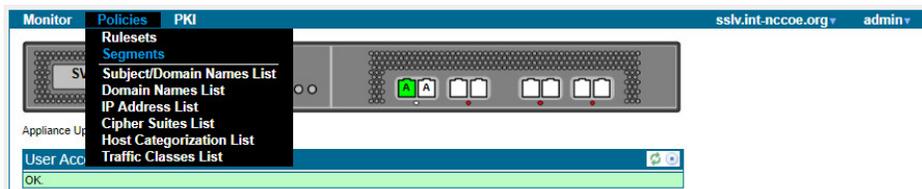


For purpose of this document we used (i).

Note: The latter two tap modes combine traffic from two or three network taps onto a single SSL Visibility Appliance segment. These ports are called *aggregation ports*.

#### 2.5.2.4.1 Add a Segment

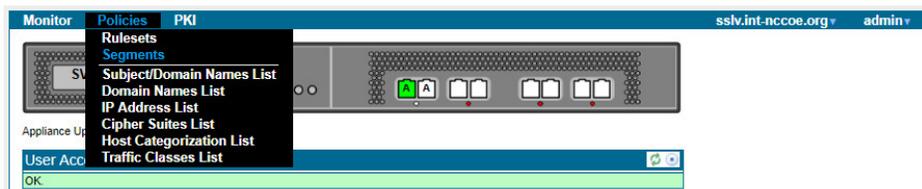
1. Select **Policies > Segments**.



2. Click the **Add**  icon in the **Segments** field.
3. Click **Edit** to select the Mode of Operation.
4. For Mode of Operation, choose  **Passive Tap** mode.
5. Click **OK**.
6. Select the **Ruleset** you previously created.
7. Choose the desired **Session Log Mode**.
8. Enter a brief description of the segment in the **Comments** box.
9. Click **OK**. The new segment appears in the *Segments* panel.
10. **Apply** the Policy Changes.

#### 2.5.2.4.2 Activate a Segment

1. Select **Policies > Segments**.



2. In the **Segments** panel, select the segment to activate.
3. Click the **Activate**  icon. The Segment Activation window displays.

Note: During segment activation, a series of screens appear that allow you to select the ports the segment will use, and any copy ports and modes where the copy ports will operate. Connect

any copy ports to your passive security devices (for example, Symantec DLP Network Monitor, Security Analytics, or an IDS).

- Follow the prompts. Once the segment is active, the system dashboard displays a green background for the segment, and there are entries under Main Interfaces and Copy Interfaces (if applicable to your deployment).
- Apply** the Policy Changes.

### 2.5.2.5 Verification

This section walks through verifying that the SSL Visibility is seeing SSL traffic without blocking it (cut through).

- To see a list of recent SSL sessions, select **Monitor > SSL Session Log**.
- Look for the domains of the servers that were accessed, and observe the value in the Action column. Since the initial rule you created cuts through all traffic, the Action should say **Cut Through** for all sessions.

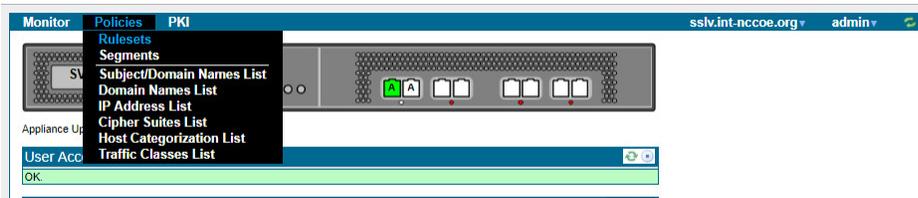
| Start Time          | Segment ID | SrcIP:Port           | DstIP:Port         | Domain Name      | Certificate Status | Cipher Suite                                | Action      | Status  |
|---------------------|------------|----------------------|--------------------|------------------|--------------------|---------------------------------------------|-------------|---------|
| Mar 18 22:37:07.723 | A          | 24.154.127.184:33387 | 23.210.249.115:443 | sb.monetate.net  | Valid              | TLS_RSA_WITH_AES_256_CBC_SHA                | Cut Through | Success |
| Mar 18 22:36:07.825 | A          | 24.154.127.184:51898 | 74.125.28.104:443  | Multiple domains | Valid              | TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 | Cut Through | Success |
| Mar 18 22:29:25.054 | A          | 24.154.127.184:33383 | 23.210.249.115:443 | Multiple domains | Valid              | TLS_RSA_WITH_AES_256_CBC_SHA                | Cut Through | Success |
| Mar 18 22:29:18.565 | A          | 24.154.127.184:33382 | 23.210.249.115:443 | Multiple domains | Valid              | TLS_RSA_WITH_AES_256_CBC_SHA                | Cut Through | Success |
| Mar 18 22:28:49.863 | A          | 24.154.127.184:33381 | 23.210.249.115:443 | Multiple domains | Valid              | TLS_RSA_WITH_AES_256_CBC_SHA                | Cut Through | Success |
| Mar 18 22:28:36.421 | A          | 24.154.127.184:51533 | 173.194.46.52:443  | Multiple domains | Valid              | TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 | Cut Through | Success |
| Mar 18 22:28:18.818 | A          | 24.154.127.184:33379 | 23.210.249.115:443 | Multiple domains | Valid              | TLS_RSA_WITH_AES_256_CBC_SHA                | Cut Through | Success |
| Mar 18 22:27:37.563 | A          | 24.154.127.184:51891 | 74.125.28.104:443  | Multiple domains | Valid              | TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 | Cut Through | Success |
| Mar 18 22:25:07.776 | A          | 24.154.127.184:52072 | 74.125.28.105:443  | Multiple domains | Valid              | TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 | Cut Through | Success |
| Mar 18 22:24:15.029 | A          | 24.154.127.184:59475 | 74.125.28.106:443  | Multiple domains | Valid              | TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 | Cut Through | Success |

#### 2.5.2.5.1 Create a Rule to Test Decryption

To test the SSL Visibility Appliance is decrypting SSL traffic, add a rule that decrypts everything from a specific source IP (e.g., your laptop).

Note: At least one rule must be added to the ruleset for SSL Visibility Appliance to start processing SSL traffic.

- Select **Policies > Rulesets**.



2. In the **Rulesets** panel, select the ruleset that was previously created.
3. In the **Rules** panel, click the **Insert**  icon to add a new rule. The **Insert Rule** dialog displays.
4. For Action, select **Decrypt (Certificate and Key Known)**.
5. Select one of the following:
  - If you imported one certificate, select **Known Certificate with Key**, and choose the certificate you imported.
  - If you imported multiple certificates, select **Known Certificates with Keys and All Known Certificates with Keys**.
6. For **Source IP**, enter the IP address of your computer.
7. Click **OK**.
8. **Apply** the Policy Changes.
9. Next Step: Use the SSL Session Log to verify that the SSL Visibility Appliance is decrypting properly.

#### 2.5.2.5.2 Verify Decryption

View the SSL Session log to test, and verify the SSL Visibility Appliance is decrypting traffic according to the rules you created.

1. Access a variety of websites or internal SSL servers. If you have created policies for specific host categories, domains, IP addresses, etc., visit websites that test these policies.
2. To see a list of recent SSL sessions, select **Monitor > SSL Session Log**.
3. Look for the domains of the websites/servers you visited, and observe the value in the Action column. Is the value you expected listed? For example, if you wanted the SSL Visibility Appliance *not* to decrypt a particular type of traffic, does the Action say Cut Through? For sessions designated as decrypted, does the Action say Decrypt? If unexpected values appear, review your policies.

Note: When a session is decrypted, the Action column will show either *Resign Certificate* (if the deployment is using the certificate resigning method) or *Certificate and Key Known* (if you have imported known certificates and keys).

| Start Time          | Segment ID | SrcIP:Port          | DstIP:Port       | Domain Name           | Certificate Status | Cipher Suite                          | Action                              | Status                       |
|---------------------|------------|---------------------|------------------|-----------------------|--------------------|---------------------------------------|-------------------------------------|------------------------------|
| Mar 12 18:11:11.084 | A          | 192.168.1.16:63463  | 192.168.3.87:443 | ws1.int-nccoe.org     | Valid              | TLS_RSA_WITH_AES_256_GCM_SHA384       | Decrypt (Certificate and Key known) | TCP queue processing timeout |
| Mar 12 18:11:09.816 | A          | 192.168.1.16:63475  | 192.168.3.87:443 | ws1.int-nccoe.org     | Valid              | TLS_RSA_WITH_AES_256_GCM_SHA384       | Decrypt (Certificate and Key known) | Success                      |
| Mar 12 18:11:05.078 | A          | 192.168.1.16:63463  | 192.168.3.87:443 | ws1.int-nccoe.org     | Valid              | TLS_RSA_WITH_AES_256_GCM_SHA384       | Decrypt (Certificate and Key known) | Success                      |
| Mar 12 18:10:56.372 | A          | 192.168.1.81:63892  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:56.286 | A          | 192.168.1.81:63891  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:56.274 | A          | 192.168.1.81:63890  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:56.264 | A          | 192.168.1.81:63889  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:56.257 | A          | 192.168.1.81:63888  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:56.243 | A          | 192.168.1.81:63887  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:56.233 | A          | 192.168.1.81:63886  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:52.484 | A          | 192.168.4.199:56169 | 192.168.3.88:443 | ws2.int-nccoe.org     | Valid              | TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 | Cut Through                         | Decrypt not possible         |
| Mar 12 18:10:39.083 | A          | 192.168.1.16:63430  | 192.168.3.87:443 | SN1.ws1.int-nccoe.org | Valid              | TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 | Drop                                | Success                      |
| Mar 12 18:10:32.485 | A          | 192.168.4.199:56133 | 192.168.3.88:443 | ws2.int-nccoe.org     | Valid              | TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 | Cut Through                         | Decrypt not possible         |
| Mar 12 18:10:26.375 | A          | 192.168.1.81:63838  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:26.296 | A          | 192.168.1.81:63837  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |
| Mar 12 18:10:26.283 | A          | 192.168.1.81:63836  | 192.168.1.95:443 | 192.168.1.95          | Self Signed        | TLS_RSA_WITH_AES_256_CBC_SHA          | Drop                                | Success                      |

### 2.5.2.5.3 Other Ways to Learn About this Deployment Method

Download a PDF (<https://www.broadcom.com/site-search?q=Visibility+SSL+First+Steps>)

View a video tutorial ([https://www.youtube.com/watch?v=qxSDDXhE\\_B8&feature=youtu.be](https://www.youtube.com/watch?v=qxSDDXhE_B8&feature=youtu.be))

## 2.5.3 Day N: Ongoing Security Management and Maintenance

### 2.5.3.1 Alerting & Monitoring

#### 2.5.3.1.1 Alerts

Use the Alerts panels to configure the email details the system will use to send out alerts, monitor events, and assess the conditions where an alert is generated. Click **Edit** to bring up the upper Edit Alert Mail Configuration window to construct details of the email system.

#### 2.5.3.1.2 SNMP Support

The SSL Visibility Appliance supports the more secure SNMP version 3, which maintains authentication and encryption for SNMP monitoring. Symantec recommends disabling SNMP versions 1 and 2c, and the default options of using AES for encryption, and SHA for authentication for SNMP version 3.

For more details, see the SSL Visibility Appliance 3.x Administration & Deployment Guide

<https://techdocs.broadcom.com/content/broadcom/techdocs/us/en/symantec-security-software/web-and-network-security/ssl-visibility/4-5/ssl-visibility-appliance-admin-deployment.html>

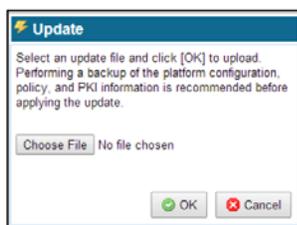
#### 2.5.3.1.3 Logging Options

Use **Platform Management (SSL Visibility-int.nccoe.org) > Logging Options** to enable or disable WebUI TLS logging and to configure remote syslog servers.

Use Logging Options to include Web UI TLS trusted channel establishment and termination logs in the System Log. These events are not included in the System Log by default.

### 2.5.3.2 Software Update

Use the **Update** menu item to load and apply a file that will update the system software. Update files are digitally signed and checked before being applied to the system. An invalid update file will not be applied.



Click **Choose File** to open a window where you browse the system and select the update file to use. Click **OK**, and the file is checked; if valid, it is copied to the system and applied.

## 2.6 Venafi Trust Protection Platform (TPP)

### 2.6.1 Prerequisites

Venafi TPP requires the following in order to be installed:

- Windows Server
- Microsoft SQL Server Database
- Hardware Security Module (if one will be used)
- Microsoft .NET Framework

### 2.6.2 Installation

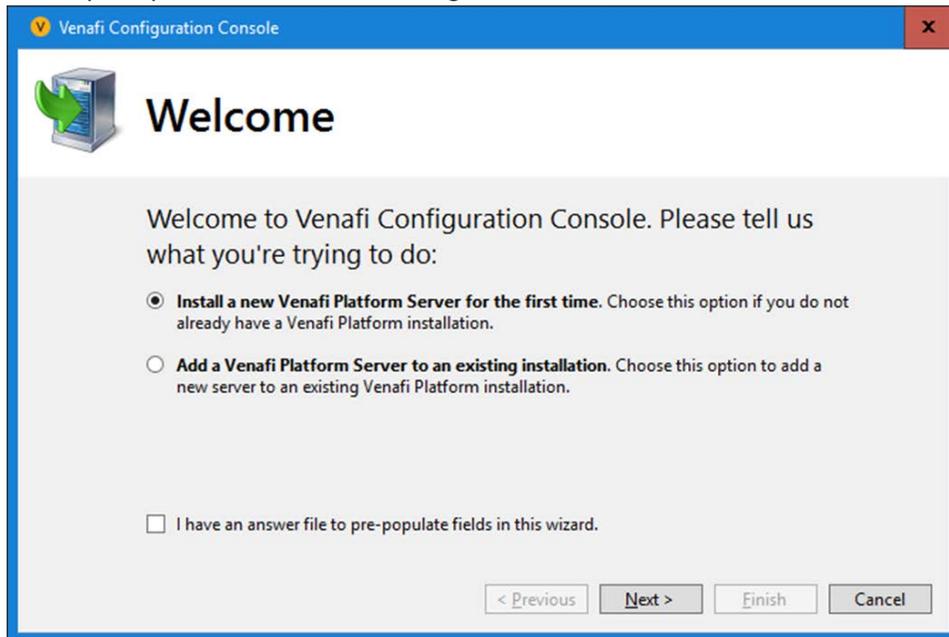
We installed Venafi TPP on Microsoft Windows Server 2012. Before starting the Venafi TPP installation, make sure you have configured your database and HSM.

The installation can be automated via a configuration file or manually performed with an installation wizard. The automated installation configuration file for installation into the production environment is typically created based on the Venafi TPP deployment in the DEV testing environment and placed in the user acceptance environment to formally test it. We recommend using the automated installation to reduce the possibility of errors during the installation into the production environment.

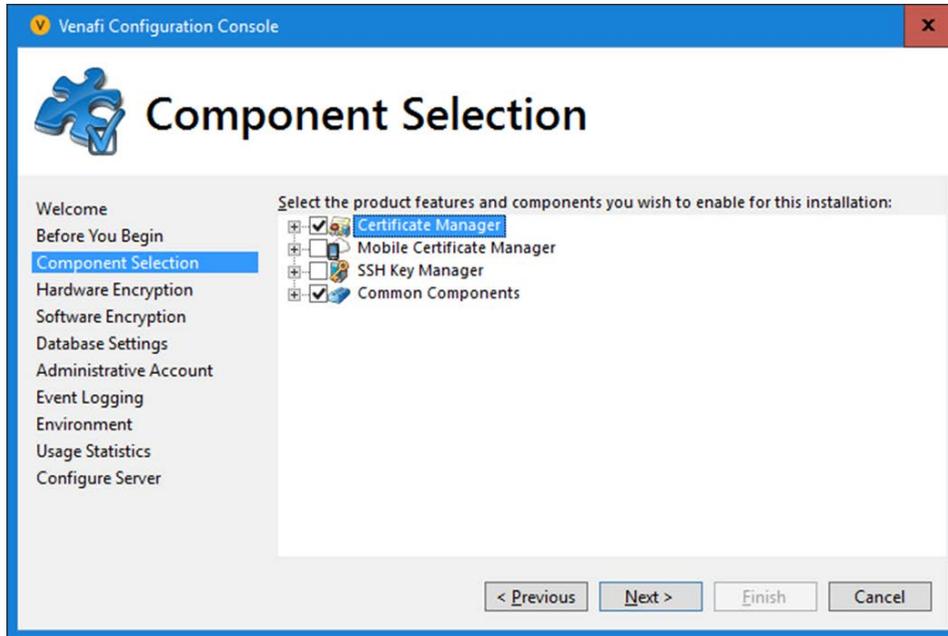
Because we were only configuring a single server in our lab environment, we manually installed and configured the product using the wizard. To install the Venafi TPP binaries and supporting files using the wizard, follow steps 1-7 in the *Venafi Trust Protection Platform Installation Guide* chapter titled “Installing using the Venafi Configuration Console wizard.”

Following step 7, the Venafi Configuration Console is automatically launched and is explained in steps 8-22 where specific integrations with the HSM and database are performed. We performed the following steps in our implementation:

1. At the prompt for first time or existing installation, select “first-time installation.”



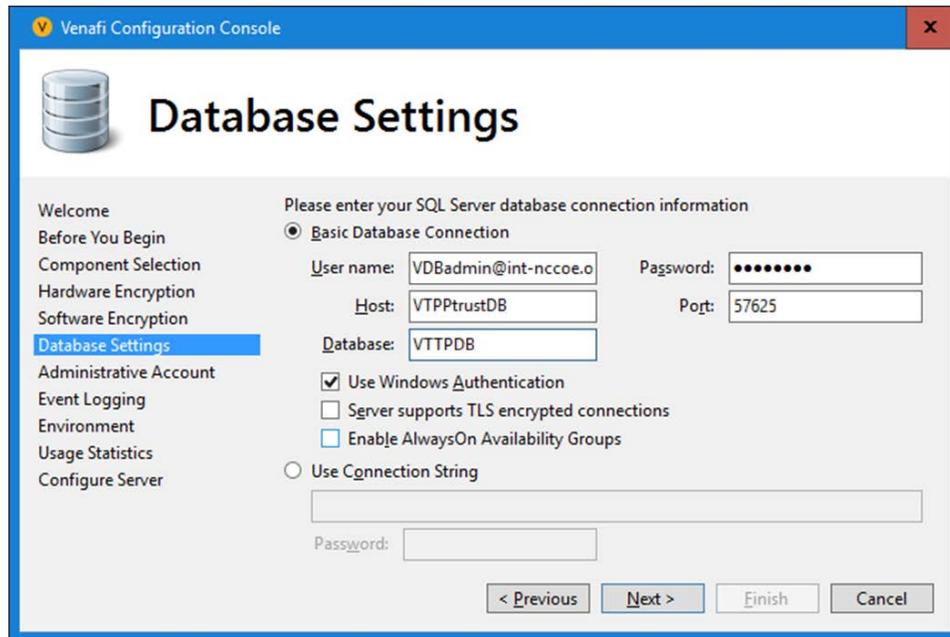
2. The Venafi Certificate Manager manages TLS server certificates, so it was selected. The Mobile Certificate and SSH Key Managers were not enabled.



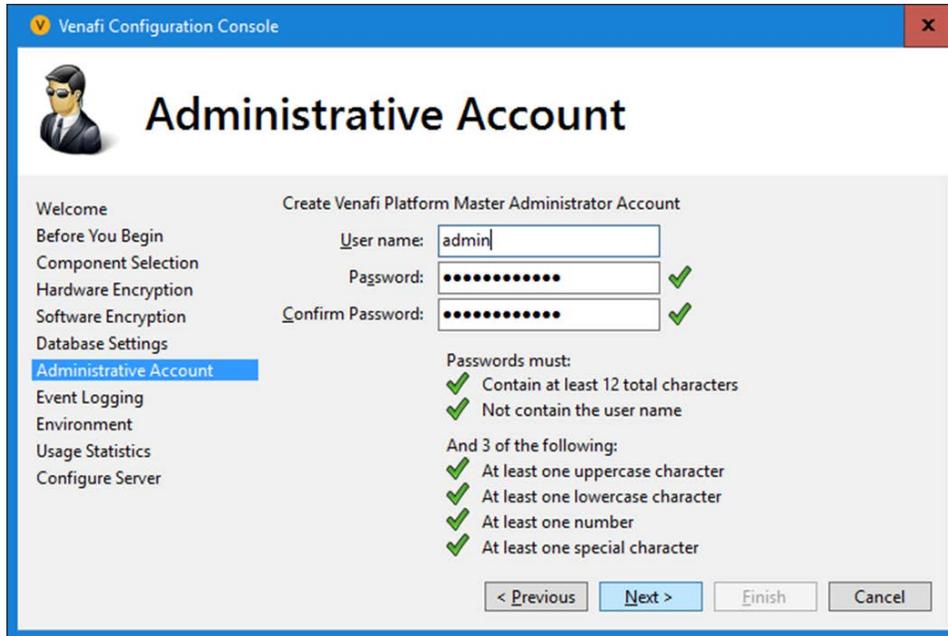
3. We recommend using an HSM with Venafi TPP to protect the symmetric key that encrypts private keys and credentials in the Venafi TPP database. In our implementation, we integrated with the Thales TCT HSM. We entered the following configuration:



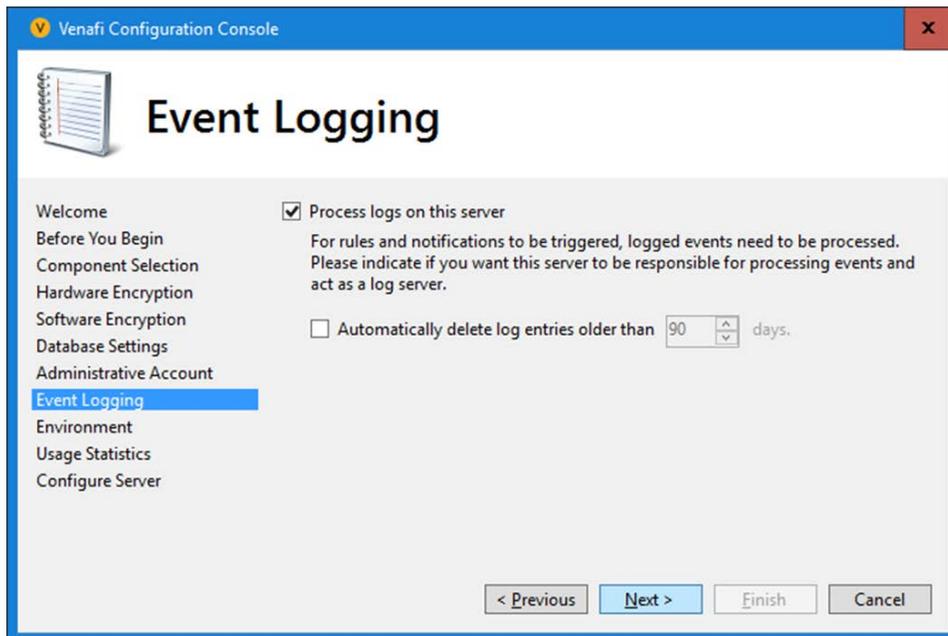
4. Windows authentication was used to authenticate to Microsoft SQL Server from Venafi TPP. Windows authentication is recommended, because it consolidates user account management, including control of password rules, failed logins, etc.



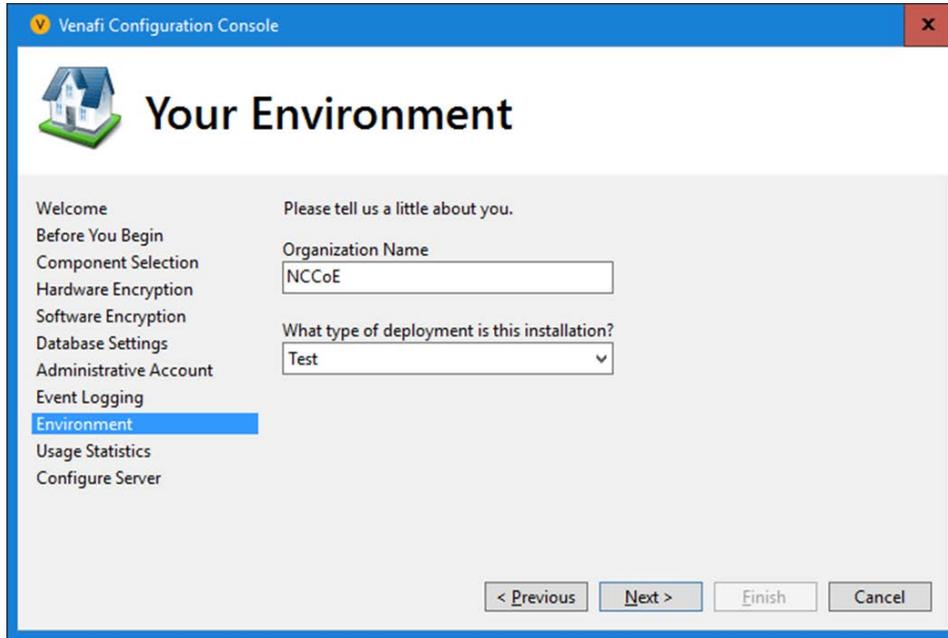
5. The initial Master Administrator account username was set to “admin,” and the password was also set.



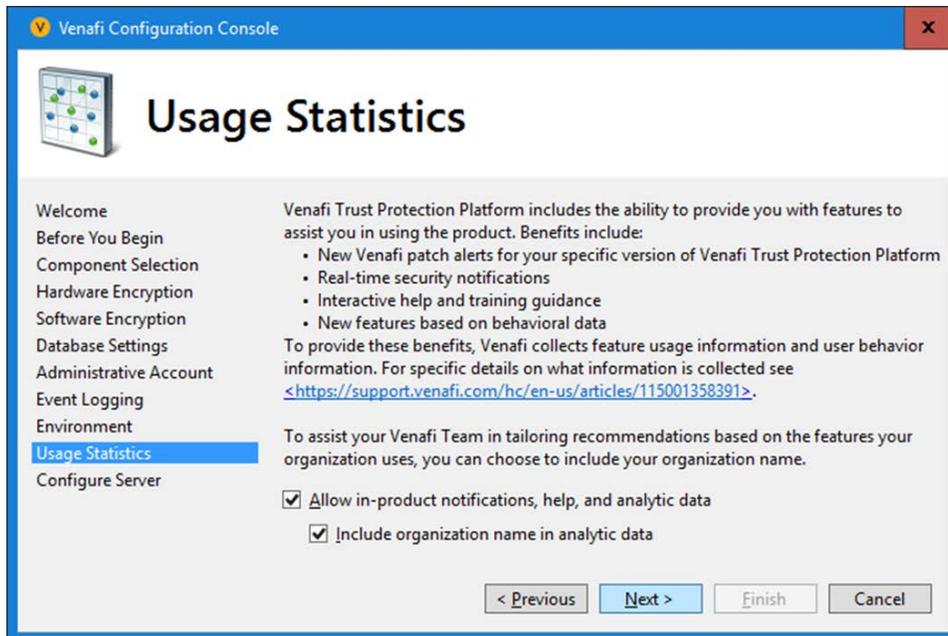
6. The Venafi TPP server was configured to process logs, as it was the only server in the environment.



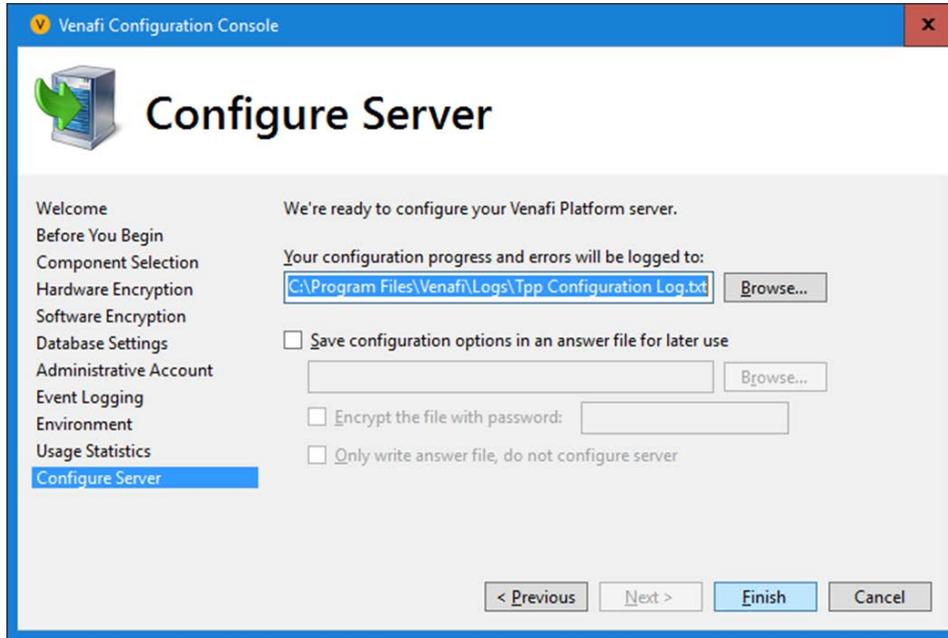
7. The organization name was set to “NCCoE”; the environment was set to “Test.”



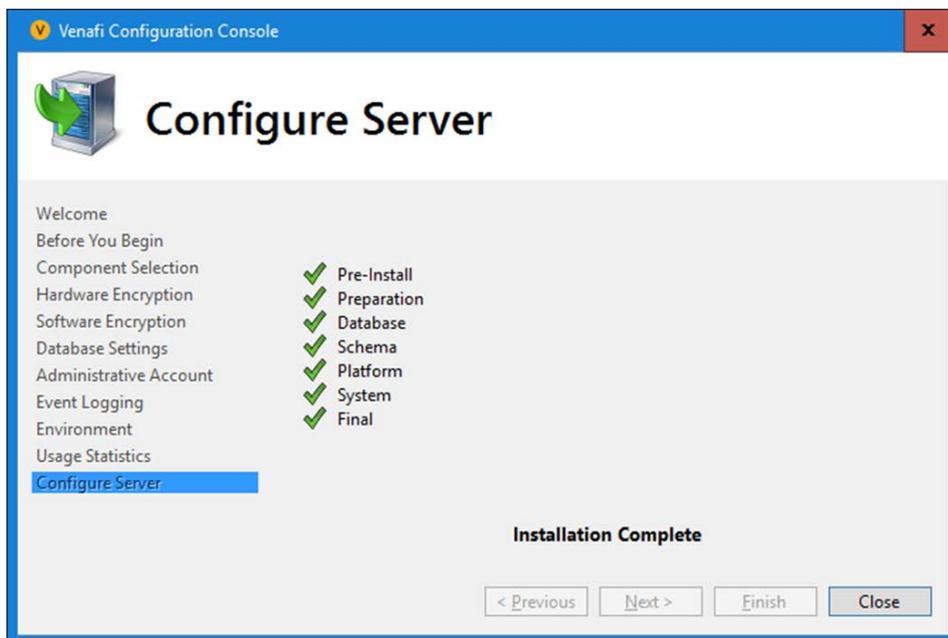
8. The collection of usage statistics was enabled.



9. The default log file location was used.



10. The Finish button was selected, and the configuration of the Venafi TPP server was completed successfully.



## 2.6.3 CA Integration

In our implementation, we integrated Venafi TPP with two CAs: DigiCert was used for publicly trusted certificates, and Active Directory Certificate Services for internally trusted certificates.

### 2.6.3.1 DigiCert

To configure integration with DigiCert so that Venafi TPP can automatically enroll for and retrieve certificates, follow the instructions in the “DigiCert CertCentral” section of the *Venafi Trust Protection Platform Certificate Authority and Hosting Platform Integration Guide*.

In our implementation, we used DigiCert Multi-SAN SSL certificates. The following configuration was used:

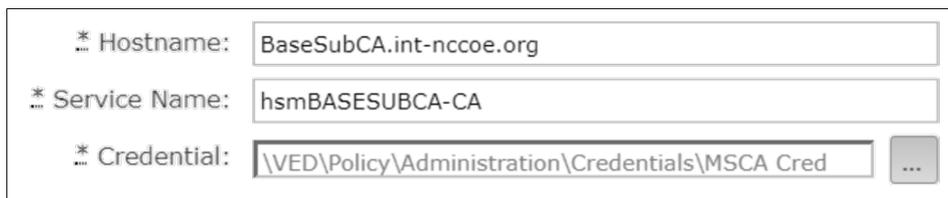
|                                  |                                               |
|----------------------------------|-----------------------------------------------|
| * Product Name:                  | Standard SSL ▼                                |
| * Organization:                  | National Cybersecurity Center of Excellence ▼ |
| Manual Approval:                 | <input type="checkbox"/>                      |
| Subject Alt Name Enabled:        | <input checked="" type="checkbox"/>           |
| Signature Algorithm:             | SHA256 ▼                                      |
| Organizational Unit Override:    | <input type="text"/>                          |
| Allow Reissuance:                | <input checked="" type="checkbox"/>           |
| Renewal Window (days):           | 90                                            |
| Certificate Transparency:        | Send certificates to a CT log server ▼        |
| * Validity Period:               | 1 year ▼                                      |
| Allow Users to Specify End Date: | <input type="checkbox"/>                      |

### 2.6.3.2 Active Directory Certificate Services

We used Microsoft ADCS to issue certificates to TLS servers inside the lab firewall. To configure integration with ADCS so Venafi can automatically enroll for and retrieve certificates, follow the instructions in the “Microsoft Active Directory Certificate Services (ADCS) - Enterprise and Standalone—

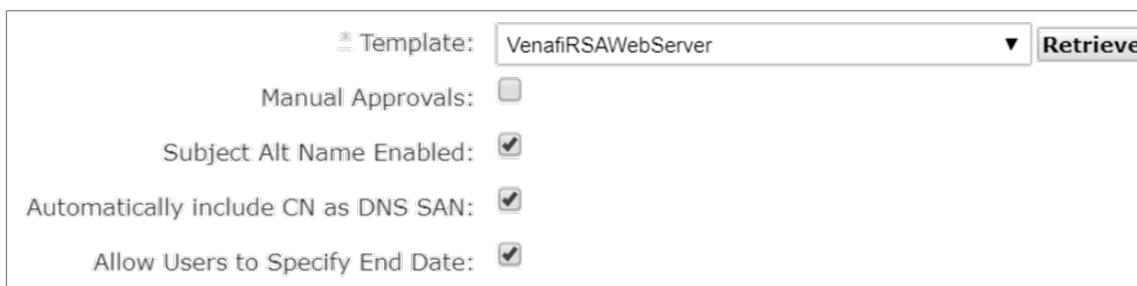
CA template configuration” section of the *Venafi Trust Protection Platform Certificate Authority and Hosting Platform Integration Guide*.

In our implementation, we configured the host name, service name, and credential information in Venafi TPP to access the ADCS Issuing CA:



A screenshot of a configuration form with three rows. Each row starts with a small icon and an asterisk. The first row is 'Hostname:' with a text box containing 'BaseSubCA.int-nccoe.org'. The second row is 'Service Name:' with a text box containing 'hsmBASESUBCA-CA'. The third row is 'Credential:' with a text box containing '\\VED\Policy\Administration\Credentials\MSCA Cred' and a small grey button with three dots to its right.

In our implementation, a certificate template named “VenafiRSAWebServer” was configured in ADCS to issue TLS server certificates. The CA template object we used in Venafi TPP to request certificates pointed to this template in ADCS and had the following configuration:



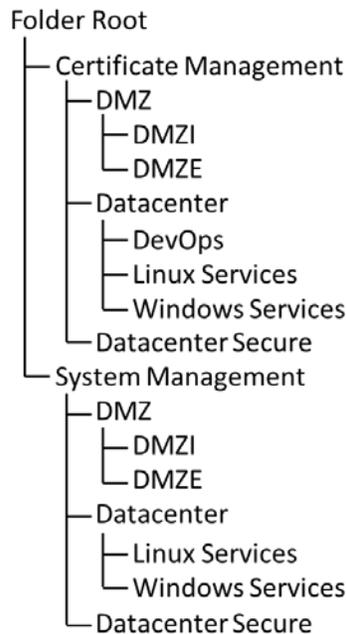
A screenshot of a configuration form for a certificate template. It has a 'Template:' dropdown menu set to 'VenafiRSAWebServer' and a 'Retrieve' button. Below are four checkboxes: 'Manual Approvals:' (unchecked), 'Subject Alt Name Enabled:' (checked), 'Automatically include CN as DNS SAN:' (checked), and 'Allow Users to Specify End Date:' (checked).

We recommend enabling “Subject Alt Name Enabled” and “Automatically include CN as DNS SAN,” as SANs in lieu of using CNs. Including a CN and SAN in certificates ensures backward compatibility with older clients that only support CNs and compatibility with newer clients that require SANs.

## 2.6.4 Folder Creation

To create a folder hierarchy for organizing certificate, application, and device objects, refer to the section titled “Managing your policies (folders)” in the *Venafi Trust Protection Platform Administration*

*Guide.* The following folder structure was created in our implementation of Venafi TPP to match the three fictitious departments of certificate owners in the lab:



### 2.6.5 Custom Fields

Follow the instructions in the section titled “Working with Custom Fields” in the *Venafi Trust Protection Platform Administration Guide* to define additional metadata fields for certificates and other objects. Two custom fields were defined in our Venafi TPP implementation: Biz Owner and Cost Center.

We configured the Biz Owner custom field with a field type of “Identity” to allow the selection of user identities in AD.

The Cost Center custom field was configured with a “String” field type, including a regex to validate that the cost centers that were entered matched the pattern of two letters, one dash, and four numbers.

(e.g., AB-1234). A custom error message displays if a cost center doesn't match the regex pattern entered by a user.

The screenshot shows a configuration form for a certificate field. The field is named "Cost Center" and is of type "String". It has a validation template set to "Custom" with a regular expression of `\b[a-zA-Z]{2}\b-\b[0-9]{4}\b`. The field is required and applies to certificates. A custom error message is defined: "Cost centers must include two letters, a dash, and four numbers (e.g., AB-1234)".

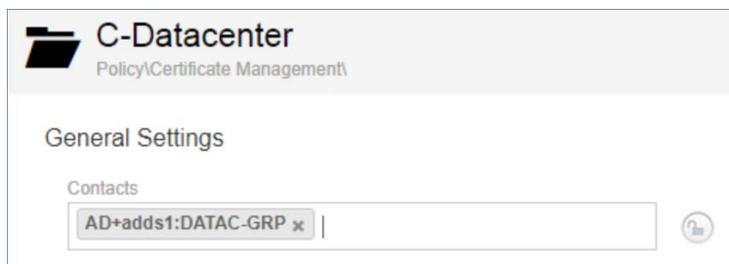
|                                                                    |                                                                                 |
|--------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Name *                                                             | Field Type *                                                                    |
| Cost Center                                                        | String                                                                          |
| Make field...                                                      | Validation Template                                                             |
| <input type="checkbox"/> Required                                  | Custom                                                                          |
| <input type="checkbox"/> Hidden                                    | Validation Regular Expression                                                   |
| <input type="checkbox"/> Controlled by Policy                      | <code>\b[a-zA-Z]{2}\b-\b[0-9]{4}\b</code>                                       |
| <input type="checkbox"/> Read-only                                 | Validate Sample Entry ⓘ                                                         |
| Apply to... *                                                      | Customizable Error Message                                                      |
| <input checked="" type="checkbox"/> Certificates                   | Cost centers must include two letters, a dash, and four numbers (e.g., AB-1234) |
| <input type="checkbox"/> Devices                                   |                                                                                 |
| Customizable Help Text                                             |                                                                                 |
| Please provide the cost center for this certificate (e.g. WR-3201) |                                                                                 |

## 2.6.6 Assigning Certificate Owners

The assignment of certificate owners was done with AD groups Venafi TPP folders in our implementation, to ensure new certificates automatically had the correct owner assigned. The AD groups were created to represent the certificate owners in the four fictitious departments in our implementation. These groups were assigned as contacts and granted permissions at the folder level.

### 2.6.6.1 Contacts

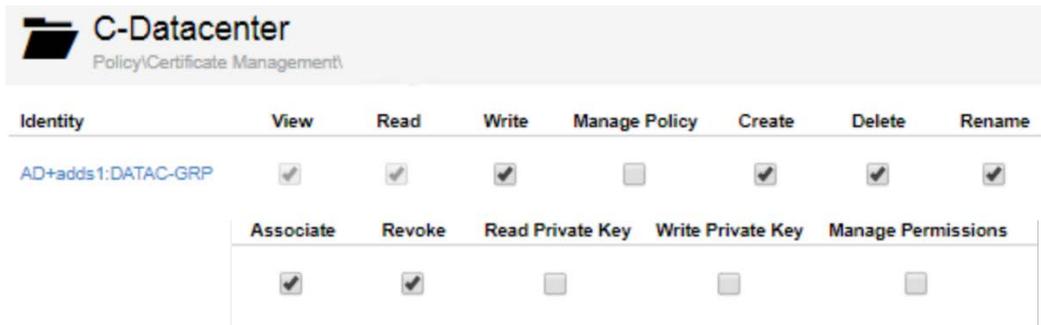
For information about assigning Contacts to folders in Venafi TPP, refer to the section titled "General configuration options" in the *Venafi Trust Protection Platform Administration Guide*. Each certificate owner AD group was assigned as a contact to their respective Venafi TPP folder, so they would receive notifications (e.g., impending expirations, errors, etc.).



### 2.6.6.2 Permissions

For instructions on assigning permissions in Venafi TPP, refer to the section titled “Assigning permissions to objects in Aperture” in the *Venafi Trust Protection Platform Administration Guide*. In our implementation, we assigned each group representing a certificate owner View, Read, Write, Create, Delete, Rename, Associate, and Revoke.

For example, the DATAC-GRP was assigned the following privileges to the C-Datacenter folder in our implementation of Venafi TPP.



The screenshot shows a table of permissions for the 'C-Datacenter' folder. The table has two rows of permissions and one identity listed. The first row contains: Identity, View, Read, Write, Manage Policy, Create, Delete, and Rename. The second row contains: Associate, Revoke, Read Private Key, Write Private Key, and Manage Permissions. The identity 'AD+adds1:DATAC-GRP' is listed with checkboxes for View, Read, Write, Create, Delete, and Rename, all of which are checked. The checkboxes for Associate, Revoke, Read Private Key, Write Private Key, and Manage Permissions are unchecked.

| Identity           | View                                | Read                                | Write                               | Manage Policy            | Create                              | Delete                              | Rename                              |
|--------------------|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| AD+adds1:DATAC-GRP | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
|                    | Associate                           | Revoke                              | Read Private Key                    | Write Private Key        | Manage Permissions                  |                                     |                                     |
|                    | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/>            | <input type="checkbox"/> | <input type="checkbox"/>            |                                     |                                     |

### 2.6.7 Setting Policies

For information about defining policies on folders in Venafi TPP, refer to the chapter titled “Using policies to manage encryption assets” in the *Venafi Trust Protection Platform Administration Guide*.

In our Venafi TPP implementation, the following policies were set:

- The Organization, City/Locality, State/Province, and Country fields within Subject DNs were locked on a top-level folder, so that those values were required in certificates across all groups.

Subject DN

Organizational Units

Organization

City/Locality

State/Province

Country

- Specific domains were placed on an allowlist. See [Section 2.6.8](#), the Establishing a Domain allowlist, of this document for more information.
- Approvers were assigned and locked at the folder level. See the “Workflow – RA Reviews” [Section 2.6.9](#) of this document for more information.
- The key length was set to 2048 on the Certificate Management folder and locked.

Key Size

- The following policies for certificate authorities were configured:
  - The internal Issuing CA was enforced on the following folders to ensure only internally issued certificates could be used:
    - DMZI
    - Datacenter
    - Datacenter Secure

CA Template

- The publicly trusted DigiCert Multi-SAN CA was enforced on the DMZE folder to ensure only publicly trusted EV certificates could be provisioned to the public facing interfaces of the F5 LTM.



## 2.6.8 Establishing a Domain Allowlist

To limit security exposure, control the domains for which certificates can be issued. For instructions on configuring the domains for which certificates can be requested in Venafi TPP (establishing a domain allowlist), refer to the section titled “To configure certificate policy on a folder” in the *Venafi Trust Protection Platform Certificate Management Guide*.

In our implementation, we allowed two internal domains (int-nccoe.org and ext-nccoe.org) for all folders that contained internal resources in Venafi TPP.



In the DMZE folder containing all the external resources, we also allowed the externally accessible domain (tls.nccoe.org).



## 2.6.9 Workflow – RA Reviews

For instructions on configuring workflow gates in Venafi TPP, refer to the section titled “Creating a certificate workflow” in the *Venafi Trust Protection Platform Certificate Management Guide*. In our implementation, we established a workflow gate for the Datacenter Secure zone. To do so, perform the following steps:

1. Create a workflow object. Assign the stage to “0.” Select “Approver assigned to object” for Request Approval From.

\* If Stage is:

If Application or Trust Store is:

Inject Commands:

Commands:

*Commands will be evaluated for macros. If the command includes a single "\$", and is not intended to be used as a macro, then "\$" should be replaced with "\$\$."*

Request Approval:

Request Approval From:  Approver assigned to object  
 Specified approver  
 Specify approver via macro

Specified Approver(s):

Approver Macro:

Approval Reason Code:

2. Assign the workflow to the Datacenter Secure folder policy.



The screenshot shows a configuration window with two sections: "Applied Workflows:" and "Blocked Workflows:". The "Applied Workflows:" section contains a text box with the path "\VED\Policy\Administration\Workflows\Stage 0 Approval" and a small grey button with three dots to its right. The "Blocked Workflows:" section is currently empty and also has a small grey button with three dots to its right.

3. Assign the appropriate AD group (datacs\_apprvr) to the **Approver(s)** for certificates on the Datacenter Secure folder.



The screenshot shows a field labeled "Approver(s)". Inside the field, the text "AD+adds1:datacs\_apprvr" is entered, followed by a small 'x' icon. To the right of the field is a circular icon containing a padlock, indicating that the field is locked.

### 2.6.10 CA Import

Once folder structure, policies, certificate owners, and other configurations are completed, begin building the inventory of certificates—start by importing certificates from the ADCS-issuing CA.

For instructions on configuring imports from ADCS, refer to the chapter titled “Importing certificates from a certificate authority” in *Venafi Trust Protection Platform Administration Guide*.

In our implementation, we configured Venafi TPP to import certificates from a particular ADCS template named, “WebBulkCertTemplate.” We included expired—not revoked—certificates. We chose not to define any placement rules and placed all certificates into a single folder named **ADCS Import**.

**CA Configuration**

CA Type  
Microsoft CA

**Get templates from Microsoft CA**

Hostname or IP Address  
BaseSubCA.int-nccoe.org

Credentials  
\VED\Policy\Administration\Credentials\MSCA Cred

Service Name  
hsmBASESUBCA-CA **Get Templates**

Select templates to import  Import all templates

| CA Templates Found | Selected for this Import |
|--------------------|--------------------------|
|                    | WebBulkCertTemplate      |

Include:  Expired certificates  Revoked certificates

**Placement Rules** [+ Add New Rule](#)

There are currently no placement rules

If no rule(s) apply,

- put certificates in: \VED\Policy\Certificate Management\ADCS I
- ignore certificates and do not place them in a policy

Automatically place certificates into policy when importing?

- Yes
- No, let me preview first in Summary

A total of 523 certificates were imported from the ADCS issuing CA.

## 2.6.11 Network Discovery

It's possible to accomplish network discovery scanning for TLS server certificates in several ways, including using existing vulnerability assessment tools or the certificate management solution. In our implementation, we used Venafi TPP to perform network discovery scans using two different methods: scanning using Venafi TPP servers and the Scanafi utility.

### *Venafi TPP Server*

In our implementation, we used Venafi TPP servers to perform network discovery scans in the Datacenter and Datacenter-Secure network zones. For instructions on performing network discoveries with Venafi TPP servers, see the chapter titled "Discovering certificates and keys" in the *Venafi Trust Protection Platform Certificate Management Guide*.

#### 2.6.11.1 Scanafi

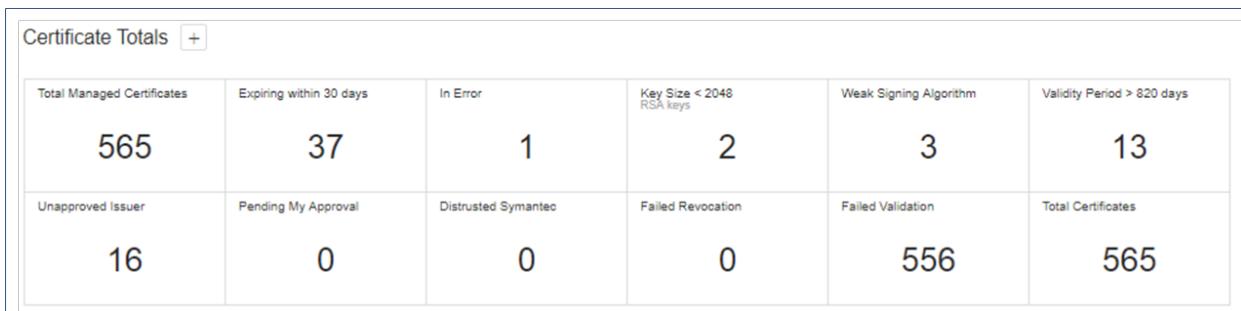
For information on using Scanafi to perform network discovery scans, refer to the section titled "Automatically calling Discovery/Import from Scanafi" in *Venafi Trust Protection Platform Web SDK Developer's Guide*.

In our implementation, we installed Scanafi on a Fedora Linux system in the DMZ network zone. The following command was used to execute a network discovery scan.

```
./scanafi_linux_x64 --tppurl=https://venafil.int-nccoe.org \
--tppuser=vscanuser --tpppass=***** --range=192.168.4.0/23 \
--zone="//VED\\Policy\\Certificate Management\\UNKNOWN ORIGIN" \
--certsonly
```

## 2.6.12 Identify Certificate Risks/Vulnerabilities

Following the import of certificates from the ADCS-issuing CA and the network discovery scans, we used the Venafi TPP dashboard to identify certificate risks and vulnerabilities. The following shows the dashboard micro-widgets for our implementation.



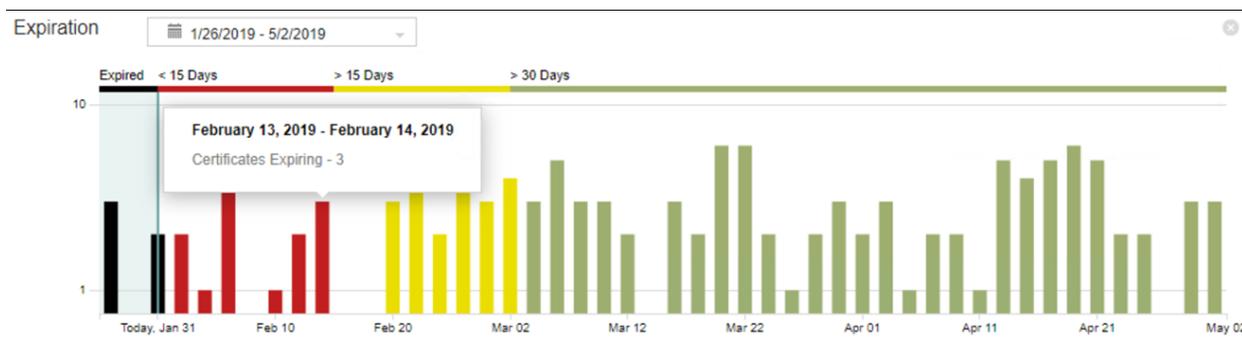
The screenshot displays a dashboard titled "Certificate Totals" with a plus sign in a box. Below the title is a table with two rows and six columns. The first row contains: "Total Managed Certificates" (565), "Expiring within 30 days" (37), "In Error" (1), "Key Size < 2048 RSA keys" (2), "Weak Signing Algorithm" (3), and "Validity Period > 820 days" (13). The second row contains: "Unapproved Issuer" (16), "Pending My Approval" (0), "Distrusted Symantec" (0), "Failed Revocation" (0), "Failed Validation" (556), and "Total Certificates" (565).

| Total Managed Certificates | Expiring within 30 days | In Error            | Key Size < 2048 RSA keys | Weak Signing Algorithm | Validity Period > 820 days |
|----------------------------|-------------------------|---------------------|--------------------------|------------------------|----------------------------|
| 565                        | 37                      | 1                   | 2                        | 3                      | 13                         |
| Unapproved Issuer          | Pending My Approval     | Distrusted Symantec | Failed Revocation        | Failed Validation      | Total Certificates         |
| 16                         | 0                       | 0                   | 0                        | 556                    | 565                        |

We used this information to identify certificates not compliant with policy (e.g., certificates issued by unapproved CAs or with weak lengths), so they could be replaced.

The dashboard was also used to identify outage risks related to certificate expirations. The following figure displays the Expiration widget of the dashboard that shows the expiration profile for certificates in our implementation.

Figure 2-2 Venafi Dashboard Expiration Widget showing the Certificate Expiration Profile



## 2.6.13 Automate Management

### 2.6.13.1 F5 BIG-IP LTM

#### 2.6.13.1.1 Discover Existing F5 Certificates and Manage

Venafi TPP can automatically discover existing certificates and configuration through its Onboard Discovery feature. Because most organizations have F5 systems with existing certificates installed, this is a common process for F5 systems we used in our implementation, which included the following steps:

1. Create an Onboard discovery job to discover certificates on F5 systems. For instructions on how to create Onboard Discovery jobs, refer to the section titled "Using Onboard Discovery" in the *Venafi Trust Protection Platform Certificate Management Guide*.
2. Create a device object in Venafi TPP with the address and credentials for the F5 device on which you want to discover and manage certificates.

|                              |                                                                            |
|------------------------------|----------------------------------------------------------------------------|
| Hostname/Address:            | <input type="text" value="192.168.3.85"/>                                  |
| Provisioning Mode:           | <input type="text" value="Agentless"/>                                     |
| Concurrent Connection Limit: | <input type="text" value="1"/>                                             |
| Device Credential:           | <input type="text" value="VED\Policy\System Management\A-Credentials\F5"/> |

3. Run the F5 Onboard Discovery job by clicking **Run Now**.

| Job Name             | Description                                 | Next Run | Last Run             | Type      | Results       | Status   |         |
|----------------------|---------------------------------------------|----------|----------------------|-----------|---------------|----------|---------|
| F5 Onboard Discovery | Discover                                    | Manual   | 1/31/2019            | Onboard   | Certificates: | Complete | Run Now |
| F5 LTM Advanced      | certs and configuration on F5 Big-IP in DMZ |          | 1:02 PM (-05:00 UTC) | Discovery | 1             |          |         |

4. Ensure the discovered certificate(s) are set to automatically renew when they are nearing expiration.

Automatic Renewal?\*

Yes

5. With this discovered configuration, including the certificate, Venafi TPP was set to automatically replace the existing certificate with a new certificate prior to expiration.

#### 2.6.13.1.2 Install a New Certificate on F5

In our implementation, Venafi TPP was used to enroll for and install a new certificate on the F5 LTM in the DMZ. The following steps were used to perform these operations:

1. Create a new certificate object in the Venafi TPP Aperture console.

Create a New Certificate

2. Select the appropriate folder.

Certificate Folder\* ?

Policy \ Certificate Management \ C-DMZ \ DMZE

3. Select a name for the certificate.

Nickname\* ?

app1.tls.nccoe.org

4. Select the “Provisioning” Management Type to configure the certificate for automated management.

Management Type\* <sup>?</sup>

Provisioning ▼

5. Enter the CN for the certificate.

Common Name <sup>?</sup>

app1.tls.nccoe.org

6. Enter the SANs for the certificate.

Subject Alternative Names (DNS)

app1.tls.nccoe.org x |

7. Configure the certificate for automatic renewal and installation when it is nearing expiration.

Automatic Renewal?\*

Yes ▼

8. Add a new installation for the certificate, and indicate that management will be automated for that installation.

**Track, validate, and automate installation of this certificate**

9. Select the F5 device where the certificate will be installed.

Find Existing Device [Create New Device](#)

Policy \ System Management \ S-DMZ \ DMZE \ F5LB1 ▼

10. Indicate that the Installation Type is “F5 BIG-IP Local Traffic Manager.”

Installation Type

F5 BIG-IP Local Traffic Manager ▼

11. The certificate we were installing was not for securing the administrative interface to the F5 LTM, therefore, we selected “No” for the Device Certificate.

Device Certificate  Yes  No

12. We indicated that Venafi TPP should update the profile when the new certificate was installed. This ensures the configuration was properly set up to use the new certificate.

Force Profile Update  Yes  No

13. We instructed Venafi TPP to install the CA certificates with the new certificate—enabling clients connecting to the F5 to validate the certificate signature with the chain.

Install Chain  Yes  No

14. We chose to have Venafi TPP bundle the CA certificates with the new certificate (in the same file on the F5 device).

Bundle Certificates  Yes  No

15. An HSM was not installed on the F5 device we were using, so we indicated this to Venafi TPP.

Use FIPS  Yes  No

16. We instructed Venafi TPP to overwrite the existing certificate each time it installed a new certificate (prior to expiration).

Overwrite Certificate and Key  Yes  No

17. We instructed Venafi TPP to delete the existing certificate when the new certificate was installed.

Delete Previous Cert and Key  Yes  No

18. To ensure the certificate was associated with the correct SSL profile on the F5 LTM, we configured the following:

### SSL Profile Settings

|                    |                                              |
|--------------------|----------------------------------------------|
| SSL Profile*       | <input type="text" value="app1_client-ssl"/> |
| SSL Profile Type   | <input type="text" value="Client"/>          |
| Parent SSL Profile | <input type="text" value="clientssl"/>       |
| SSL Partition      | <input type="text" value="Common"/>          |

19. We provided Venafi TPP information about the virtual server where the certificate should be associated.

### Virtual Server Settings

|                          |                                      |
|--------------------------|--------------------------------------|
| Virtual Server*          | <input type="text" value="app1_vs"/> |
| Virtual Server Partition | <input type="text" value="Common"/>  |

20. We indicated to Venafi TPP that we did not use mutual authentication or other advanced features on the F5 LTM.

### Advanced Settings

Use Advanced Settings  Yes  No

21. After configuring these settings, we clicked **Save**.

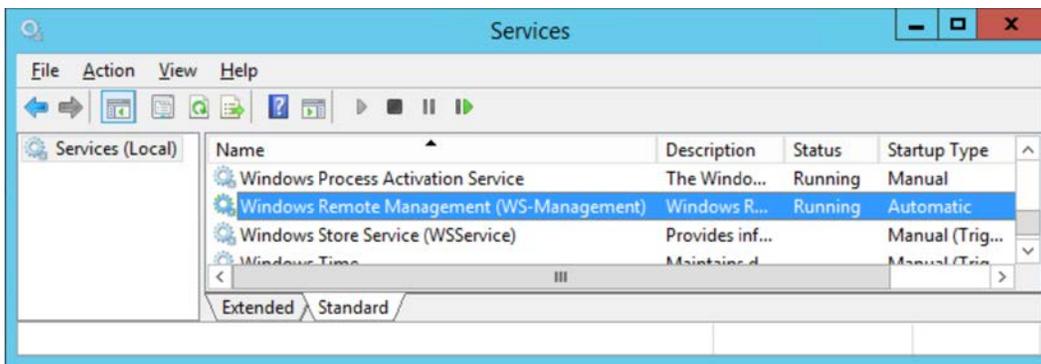
**Save**

22. Click **Renew Now** on the certificate to start to enroll a new certificate and to install it on the F5 LTM with these configuration settings.

### 2.6.13.2 Microsoft IIS – Agentless

The Microsoft IIS system we used in our implementation to demonstrate automated management had an existing certificate. Venafi TPP can automatically discover existing certificates and configuration through its Onboard Discovery feature. Consequently, the following process was used:

1. Create an Onboard discovery job to discover certificates on Microsoft IIS systems. For instructions on how to create Onboard Discovery jobs, refer to the section titled “Using Onboard Discovery” in the *Venafi Trust Protection Platform Certificate Management Guide*.
2. Confirm Windows Remote Management (WinRM) service was running on the Windows server hosting IIS.



3. Enable WinRM at the command line.

```
C:\>winrm quickconfig
```

4. Create a device object in Venafi TPP with the address of the Windows server hosting IIS and a credential for Venafi TPP to authenticate to the system.

|                              |                                                                                    |
|------------------------------|------------------------------------------------------------------------------------|
| Hostname/Address:            | <input type="text" value="192.168.3.5"/>                                           |
| Provisioning Mode:           | <input type="text" value="Agentless"/>                                             |
| Concurrent Connection Limit: | <input type="text" value="1"/>                                                     |
| Device Credential:           | <input type="text" value="\\VED\Policy\System Management\A-Credentials\IIS2"/> ... |

- Execute the IIS Onboard Discovery job that applied to the folder where the device was located. The certificate and binding configuration on IIS were discovered.

| Job Name                   | Next Run | Last Run                       | Type              | Results         | Status   |
|----------------------------|----------|--------------------------------|-------------------|-----------------|----------|
| IIS<br>CAPI (IIS Bindings) | Manual   | 1/27/2019 8:09 PM (+00:00 UTC) | Onboard Discovery | Certificates: 1 | Complete |

- The certificate is discovered.

The screenshot shows the Venafi TPP interface for a discovered server certificate. The certificate is for **iis2.int-nccoe.org** and was discovered on 1/27/2019. The details include:

- Template:** Venafi RSA Web Server
- Issuer:** hsmBASESUBCA-CA
- Common Name:** iis2.int-nccoe.org
- Organization:** NCCOE
- City/Locality:** Gaithersburg
- State/Province:** Maryland
- Country:** US
- Key Size:** 2048
- Key Usage:** Digital Signature, Key Encipherment (a0)
- Enhanced Key Usage:** Server Authentication (1.3.6.1.5.5.7.3.1)

- In addition, IIS binding information is discovered, so that all the necessary configuration for automated management is populated in Venafi TPP.

The screenshot shows the Venafi TPP interface displaying the discovered IIS binding information for **iis2.int-nccoe.org**. The binding details are as follows:

| Installation Type                                | Device             | Contacts        | Installation Status                                                                | SSL/TLS Validation Port |
|--------------------------------------------------|--------------------|-----------------|------------------------------------------------------------------------------------|-------------------------|
| iis2.int-nccoe.org (443_iis2.int-nccoe.org) CAPI | iis2.int-nccoe.org | local:VTTPAdmin | Installation Validation Successful<br>Last Checked: 4/22/2019 1:00 AM (-04:00 UTC) | 443                     |

- To ensure the certificate automatically renews and is replaced when nearing expiration, confirm the certificate was set to automatically renew prior to expiration.

Automatic Renewal?\*

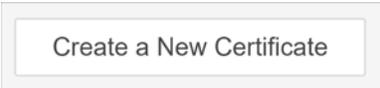
Yes ▼

### 2.6.13.3 Microsoft IIS with Thales TCT HSM – Agentless

The Venafi TPP server was used to remotely trigger the generation of a key pair and CSR on the Thales TCT HSM. The HSM is connected to the Microsoft IIS server in the Datacenter Secure zone and can enroll a certificate using the generated CSR. It can also install the certificate in the Windows server with the

proper configuration for the Microsoft IIS server. The following steps are used to perform these operations:

1. Ensure the Thales TCT HSM client is installed and configured on a Windows server hosting Microsoft IIS. See Section [2.2.2.4](#) for instructions.
2. Create a new certificate object in the Venafi TPP Aperture console.



3. Select the appropriate folder.

Certificate Folder\* (?)

Policy \ Certificate Management \ C-Datacenter Secure x ▾

4. Select a name for the certificate.

Nickname\* (?)

IIS-SafeNet-HSM

5. Select the “Provisioning” Management Type to configure the certificate for automated management.

Management Type\* (?)

Provisioning ▾

6. Enter the CN for the certificate.

Common Name (?)

hrhsm.int-nccoe.org

7. Enter the SANs for the certificate.

Subject Alternative Names (DNS)

hrhsm.int-nccoe.org x

8. Configure the certificate for automatic renewal and installation when it is nearing expiration.

Automatic Renewal?\*

Yes ▼

9. Add a new installation for the certificate and indicate that management is automated for that installation.

**Track, validate, and automate installation of this certificate**

10. Enter the address for the device where the certificate will be installed.

Device Address [Find Existing Device](#)

hrhsm.int-nccoe.org

11. Select the folder where the device object should be created.

Choose Device Folder

Policy \ System Management \ S-Datacenter Secure ▼

12. Indicate that the application type for the installation is “Windows CAPI & IIS.”

Installation Type

Windows CAPI & IIS ▼

13. Select the credential to authenticate to the system for management operations.

Device Credential

Policy \ System Management \ A-Credentials \ HRhsm credential x ▼

14. Enter a CAPI-friendly name for the certificate to be installed.

Friendly Name\*

HRhsm.int-nccoe.org

15. Click **Renew Now** on the certificate to start generating a new key pair on the HSM and to start getting a new corresponding certificate.

#### 2.6.13.4 Apache – Agentless

1. Create a new certificate object in the Venafi TPP Aperture console. For instructions on creating a new certificate, refer to “Creating a new certificate in Aperture” in *Venafi Trust Protection Platform Working with Certificates*.
2. Add an installation location for the certificate for the Apache where the certificate will be installed. For instructions on adding an Apache installation in Aperture, refer to the section titled “Creating an Apache application object” in the *Venafi Trust Protection Platform Certificate Authority and Hosting Platform Configuration Guide*. Notable configuration information that we used in our implementation, includes:
  - a. Set the private-key file location to correspond to the Virtual Host configuration on the Apache server.

|                   |                                                               |
|-------------------|---------------------------------------------------------------|
| Private Key File* | <input type="text" value="/etc/pki/tls/private/private.key"/> |
|-------------------|---------------------------------------------------------------|

- b. Set the certificate file location to correspond to the Virtual Host configuration on the Apache server.

|                   |                                                          |
|-------------------|----------------------------------------------------------|
| Certificate File* | <input type="text" value="/etc/pki/tls/certs/cert.crt"/> |
|-------------------|----------------------------------------------------------|

- c. Set the CA certificate chain file location to correspond to the Virtual Host configuration on the Apache server.

|                        |                                                              |
|------------------------|--------------------------------------------------------------|
| Certificate Chain File | <input type="text" value="/etc/pki/tls/certs/ca-chain.crt"/> |
|------------------------|--------------------------------------------------------------|

- d. Instruct Venafi TPP to update the CA chain.

|                          |                                      |                          |
|--------------------------|--------------------------------------|--------------------------|
| Overwrite Existing Chain | <input checked="" type="radio"/> Yes | <input type="radio"/> No |
|--------------------------|--------------------------------------|--------------------------|

3. Click **Install** in the Actions menu to deploy the certificate to the Apache system.

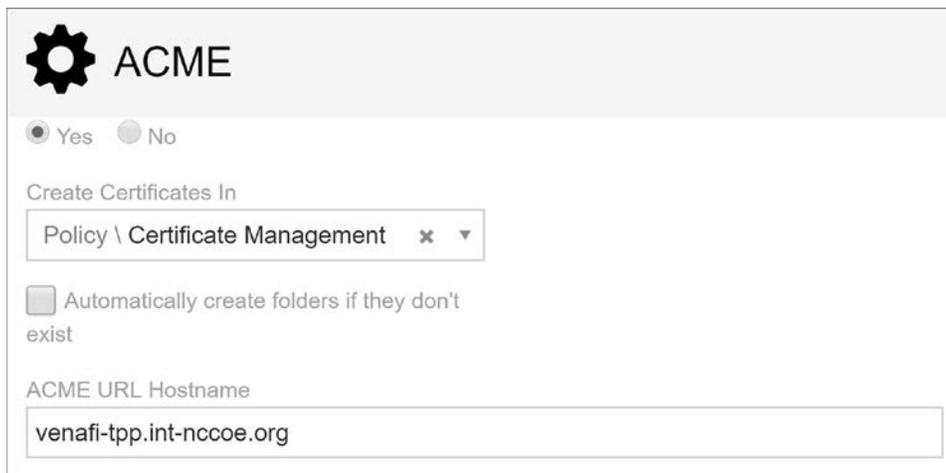
#### 2.6.13.5 Apache – ACME

Venafi TPP was configured as an ACME server in our implementation to support ACME-based requests from internal systems. For instructions on using ACME with Venafi TPP, refer to the section titled “ACME integration with Trust Protection Platform” in the *Venafi Trust Protection Platform Certificate Management Guide*.

### 2.6.13.6 *Configuring Venafi TPP for ACME*

The following steps are needed for configuring Venafi TPP to request certificates using an ACME client.

1. Configure Venafi TPP to enable the ACME server.
  - a. The ACME server is not enabled by default in Venafi TPP.
  - b. When ACME is enabled, select the folder where ACME-enrolled certificates are placed.
  - c. Enter the address of the Venafi TPP server that will service ACME clients.



The screenshot shows the ACME configuration window in Venafi TPP. At the top, there is a gear icon and the text 'ACME'. Below this, there are two radio buttons: 'Yes' (selected) and 'No'. Underneath, there is a section titled 'Create Certificates In' with a dropdown menu showing 'Policy \ Certificate Management'. Below the dropdown is a checkbox labeled 'Automatically create folders if they don't exist', which is currently unchecked. At the bottom, there is a text input field labeled 'ACME URL Hostname' containing the text 'venafi-tpp.int-nccoe.org'.

2. Assign an email address to the requesting account. The ACME protocol requires an email address be provided during the registration process. Venafi TPP must be able to find the entered email address in the local Venafi TPP identity directory or AD (depending on which directory is used).

### 2.6.13.7 *Configuring Certbot for Apache*

Certbot is the standard client use for ACME on many systems. Find instructions on installing certbot at the following address: <https://certbot.eff.org/>. We installed certbot on a Fedora Linux system to automate certificate requests and installation for Apache.

We performed the following steps in our implementation.

1. Ensure the virtual host is configured in Apache.
2. Install certbot for Apache.

```
sudo dnf install certbot certbot-apache
```

3. The root certificate for the CA that issued the Venafi TPP server's certificate must be trusted on the system where certbot is run. This is done by adding it to one of the following files depending on the OS:

```
/etc/ssl/certs/ca-certificates.crt", // Debian/Ubuntu/Gentoo etc.
/etc/pki/tls/certs/ca-bundle.crt", // Fedora/RHEL 6
/etc/ssl/ca-bundle.pem", // OpenSUSE
/etc/pki/tls/cacert.pem", // OpenELEC
/etc/pki/ca-trust/extracted/pem/tls-ca-bundle.pem", // CentOS/RHEL 7
```

#### 4. Run certbot to request a certificate. A certificate was installed on the Apache system.

```
certbot certonly \
--server "https://venafil.int-nccoe.org/vacme/v1/directory" \
--cert-name apache1 --domains apache1.int-nccoe.org \
--apache --email acmeuser@int-nccoe.org --no-eff-email
```

### 2.6.13.8 *Kubernetes*

Instructions for installing, configuring, and using Kubernetes are available on <https://kubernetes.io/>.

We installed a three-node Kubernetes cluster on three CentOS Linux systems in the Datacenter network zone in our implementation. We installed the following for the Kubernetes deployment:

- Docker version 18.09.3, build 774a1f4
- kubelet, kubeadm, and kubectl v1.13.4
- Weave (as our overlay network)

Once these components were installed, we installed and configured cert-manager in Kubernetes to automatically request certificates for ingresses in Kubernetes. We performed the following steps:

1. Verified a user account with Venafi TPP WebSDK access and permissions to the folder(s) where certificates are being requested from cert-manager (see the definition of the issuer below). We created a user named “vapirequester” in AD for this purpose. The account was granted Create, Write, Read, and View permissions to a folder named DevOps. We also granted that account WebSDK access.

Allow WebSDK Access:

2. Verified Jetstack Cert-Manager was installed with the necessary components to request certificates from Venafi TPP. This automatically creates a namespace named “cert-manager,” which we used for the rest of our configuration.

```
[ec2-user@kubemaster ~]$ kubectl describe deployment cert-manager -n cert-manager
Name: cert-manager
Namespace: cert-manager
CreationTimestamp: Wed, 06 Mar 2019 03:15:23 +0000
Labels: app=cert-manager
 chart=cert-manager-v0.6.0-venafi.0
 heritage=Tiller
 release=cert-manager
Annotations: deployment.kubernetes.io/revision: 2
 kubectl.kubernetes.io/last-applied-configuration:
 {"apiVersion":"apps/v1beta1","kind":"Deployment","metadata":
Selector: app=cert-manager,release=cert-manager
Replicas: 1 desired | 1 updated | 1 total | 1 available | 0 unavailable
StrategyType: RollingUpdate
MinReadySeconds: 0
RollingUpdateStrategy: 25% max unavailable, 25% max surge
Pod Template:
 Labels: app=cert-manager
 release=cert-manager
 Service Account: cert-manager
 Containers:
 cert-manager:
 Image: quay.io/jetstack/cert-manager-controller:venafi-0
 Port: <none>
 Host Port: <none>
 Args:
 --cluster-resource-namespace=$(POD_NAMESPACE)
 --leader-election-namespace=$(POD_NAMESPACE)
 Requests:
 cpu: 10m
 memory: 32Mi
 Environment:
 POD_NAMESPACE: (v1:metadata.namespace)
 Mounts: <none>
 Volumes: <none>
 Conditions:
 Type Status Reason
 ---- -
 Progressing True NewReplicaSetAvailable
 Available True MinimumReplicasAvailable
OldReplicaSets: <none>
NewReplicaSet: cert-manager-7d9f97d789 (1/1 replicas created)
Events: <none>
[ec2-user@kubemaster ~]$
```

```
kubectl apply -f https://raw.githubusercontent.com/jetstack \
/cert-manager/venafi/contrib/manifests/cert-manager/with-rbac.yaml
```

3. Created Kubernetes secret for authenticating to Venafi TPP.

```
kubectl create secret generic tppsecret \
--from-literal=username='vapirequester' \
--from-literal=password='*****' \
```

```
--namespace cert-manager
```

4. Copied the Root CA certificate that the certificate on the Venafi TPP chains up to (this is used by cert-manager to validate the Venafi TPP certificate). This was copied to a file named *rootca.pem*.
5. Generated a base64 representation of the Root CA certificate.

```
cat rootca.pem | base64 | tr -d '\n'
```

6. Created a yaml file (*tppvenafiissuer.yaml*) for the configuration for a cert-manager issuer that points to Venafi TPP. Note that the base64 representation of the Root CA certificate is placed after “caBundle:” with a single space separating (there is no carriage return). The “zone” sets the folder where the requested certificate will be placed.

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
 name: tppvenafiissuer
 namespace: cert-manager
spec:
 venafi:
 zone: 'Certificate Management\C-Datacenter\DevOps'
 tpp:
 url: https://venafil.int-nccoe.org/vedsdk
 credentialsRef:
 name: tppsecret
 caBundle:
```

```
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURSB0tLS0tCk1JSUMvVENDQWVXZ0F3SUJBZ01RSnBydys5NUMyNnhKd2FEeXF5WUhxekFOQmdrcWhraUc5dzBCQVZfZkRkFEQVlKTVE4d0RRWURWUWVFERXdaU1QwOVVRMEV3SGhjTklUZ3d0eKe1TWpNME1EUTVXaGNOTWpBd056QTVNak0xTURRNAPxakFSTVE4d0RRWURWUWVFERXdaU1QwOVVRMEV3Z2dFau1BMEdDU3FHU01lM0RRRUJBUVVBQTRJQkR3QXdnZ0VLCkFvSUJBURaaH2xUXk3ckZrTnlWenZxSW5GeE4ydVBLTEJRdz1lMk5kb1NmTXhMTVU5T1B4UUcwOVNyT1V1SSsKYmhhkckJNeEtFbStzMm5PTUNT Y3g2SDNldGp0UmtWU2pxQVZkYnQrVkn0TmtQWlZYTlRkaWlkOFVlTmRYYY1dDMQp jMkM5RUVBNDVUOG94eG10TEkvd010N2RamHpwVldxSitvT1VLVGFIZWpRTFcvUxYwKivU3AvZzFuUmFOMXhqCjFZV1lRQ2dCMWxVZ01GQ3lXUzJJSmWvQXMrRjN6ckFOazg1K0krYlBCQ050ZUFYVTNkS0xTU0N2WmxqdVZlYncKa2QwVzhzMDRPRmdCR2lCM2o2MXBydEZZc1N5WlZKYjNKVDRFVWpTM1NBbXlHZlFteVFheEpJWC9RbmIzSGp5NwpHa0ViaVFqT1FLNE9mYlZiU2tKcTh5bHdmNkhEQWdnNQkFBR2pVVEJQTUFzR0ExVWREd1FFQXdJQmhqQVBCZ05WckhSTUJBJjhFQlRBREFRSC9NQjBHQTFVZERnUUVdCQlRZkZBtL3dWR EptaEdmUCTxbHJQcUI2M0t5akRBURUJna3IKQmdFRUFZSTNGUUVFQXdJQkFEQU5CZ2txaGtpRz13MEJBUXNGQUFpQ0FRRUFGZk5EeWVlK1ZSSGhrUEx1Y1pGeQpmTlNEb0d0alZQck15Q2J3aXMyQUFOL0xYV2JMVz1YUG1YOWVwSFJQ3Zla1Rfa0RQam1OVWxFd0cwTGUwbnBycmM3bTVrbDh jYTBNaHhkMUhUrm1XbmtydjdmRy80dmt6eUhxR0FwekNTcFlyUEhsS01EaisxU1pmY1VrQ2lWVWQKb2RjL3V3K1A1RTNHa1NJZHdaK0RoODRFRVURhQ0JHc1I1MzZOMnlaMUR jekRTUWg5SHBP aTh6b3dYcnFWbzdkcApCYVpsUUNRUG1 jN0hRaE0rS0VLM1Vha1J4U1Z2ciszoEJRvysOS9zbUFET1QxN2o0MmxEcHFpdjRBTWd4cUxWCmdXMFrsclpwK1FHRnU1TEExjSnVqS311T09nM2NYanI3S11wU0FoOVpWNzFpcFRzL2Q4Nzdi1dWdPYURkL2Yrd1kKSFE9PQotLS0tLUVORCBDRVJUSUZJQ0FURSB0tLS0tCgo=
```

7. Created the issuer in Kubernetes using the newly created file.

```
kubectl apply -f tppvenafiissuer.yaml
```

- Created a yaml file for the ingress to the nginx service. Note the annotation 'certmanager.k8s.io/issuer: "tppvenafiissuer"' in the yaml file. This tells Jetstack Cert-Manager that it should automatically request and install a certificate from this ingress using the issuer we defined earlier. Cert-manager uses the host name under **tls** and **hosts** (kube-ingress.int-nccoe.org) for the CN and SAN it submits in the certificate request to Venafi TPP.

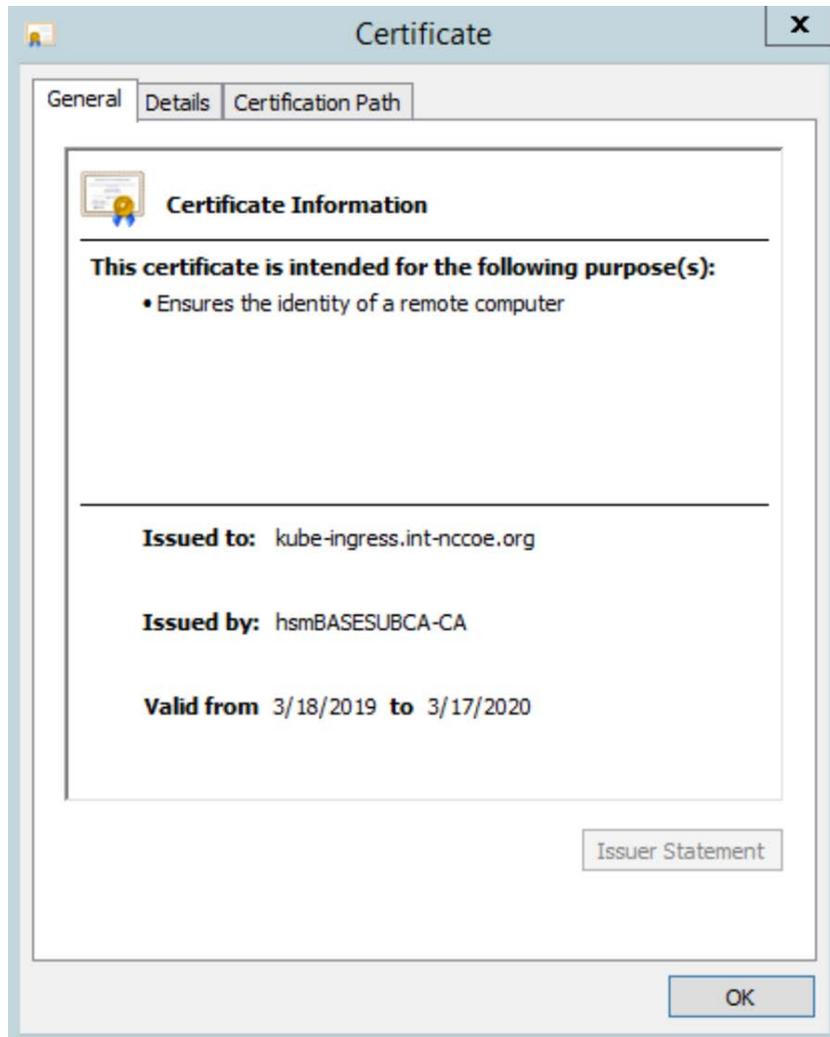
```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
 name: nginx-ingress
 namespace: cert-manager
 annotations:
 kubernetes.io/ingress.class: "nginx"
 certmanager.k8s.io/issuer: "tppvenafiissuer"

spec:
 tls:
 - hosts:
 - kube-ingress.int-nccoe.org
 secretName: nginx-cert
 rules:
 - host: kube-ingress.int-nccoe.org
 http:
 paths:
 - path: /
 backend:
 serviceName: nginx
 servicePort: 80
```

- Created the ingress.

```
kubectl create -f nginx-ingress.yaml
```

10. Once the ingress was created, connected with a browser kube-ingress.int-nccoe.org to confirm that a certificate was properly issued through Venafi TPP and installed for the ingress.



### 2.6.13.9 Symantec SSL Visibility

In our implementation, we configured Venafi TPP to automatically install TLS certificates and private keys used on several of the TLS servers—including IIS and Apache—onto the Symantec SSL Visibility to inspect traffic going to those servers.

1. Device object was created in Venafi TPP with the address and credentials for the Symantec SSL Visibility. For instructions on adding a device object, refer to the section titled “Adding Objects” in the *Venafi Trust Protection Platform Administration Guide*.

2. To ensure all required certificates and private keys are copied to the TLS inspection device, Venafi includes a feature called Bulk Provisioning. We created a bulk provisioning job.



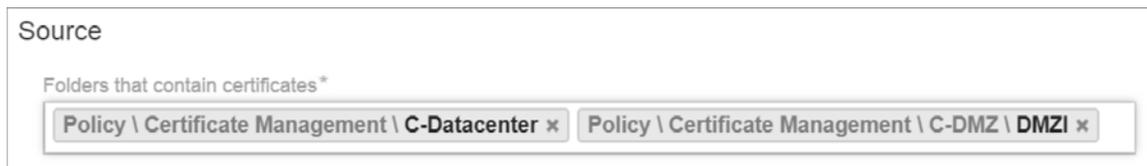
3. We named the job to distinguish it from other bulk provisioning jobs.



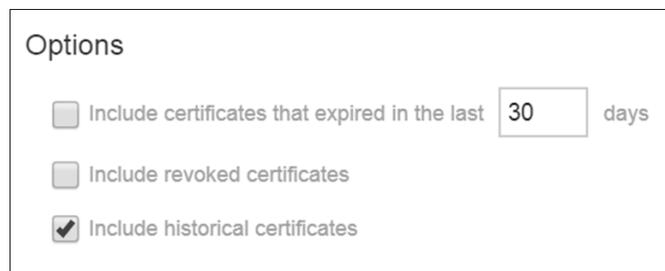
4. We selected the device object created above for the Symantec SSL Visibility Appliance as the target to which private keys would be provisioned.



5. Venafi TPP was instructed to provision private keys associated with certificates in two folders:



6. The default options excluded expired and revoked certificates and included historical certificates. Historical certificates are certificates that Venafi replaced by Venafi TPP. These certificates are still valid (not expired) and active on certain systems, though a new certificate was issued. Consequently, it is important to provision them to the TLS inspection appliance to ensure all traffic can be decrypted.



7. The bulk provisioning job was configured to run every Sunday at midnight to ensure new certificates and private keys are deployed to the TLS inspection device.

Run Time (All times are local)

Frequency \*

On Days \*

Start Time \*

- Venafi TPP uses an adaptable framework for bulk provisioning, so these jobs can be customized based on the environment's requirements. To support bulk provisioning to the Symantec SSL Visibility, the bulk provisioning script has the Venafi TPP copied into the *C:\Program Files\Venafi\Scripts\AdaptableBulk* directory. The bulk provisioning job was configured to use this script.

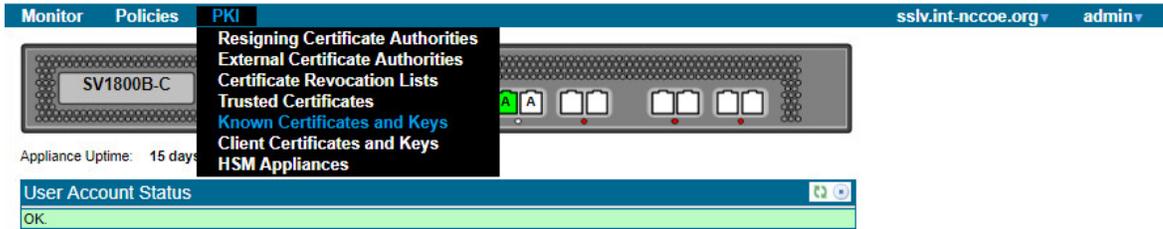
Settings

PowerShell Script\*

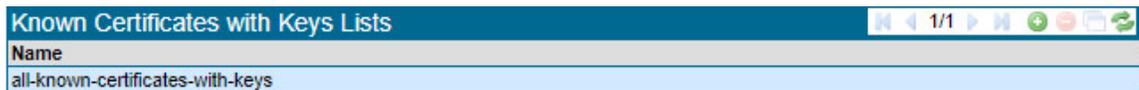
List Name

- The bulk provisioning job will run once it is saved. The private keys were confirmed to be on the device.
- To check if keys are saved in the SSL VISIBILITY, login to the SSL VISIBILITY WebUI by going to <https://192.168.1.95>

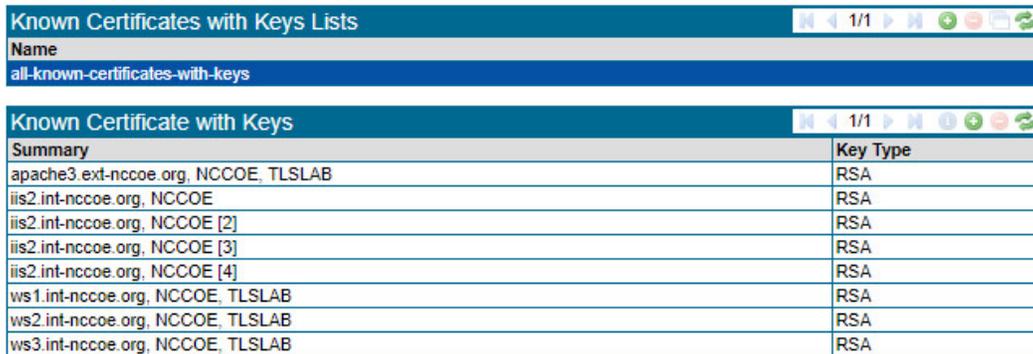
- Go to **PKI > Known Certificates and Keys**.



- In the **Known Certificates with Keys** Lists field, click on the **all-known-certificates-with-keys** field.



- The imported certificates and keys are then shown under the Known Certificate with Keys field.



## 2.6.14 Continuous Monitoring

Venafi TPP provides several tools that can continuously monitor TLS certificates within an enterprise, including scheduled network discovery scanning, monitoring certificates for expiration, and monitoring the operational status of known certificates.

### 2.6.14.1 Regular Network Scanning

In the lab, Venafi TPP was configured to perform weekly network discovery scans of the Datacenter and Datacenter Secure networks zones from the Venafi TPP server. The scans were scheduled to run at 2:00 a.m. each Sunday. The lab network was small enough for network scans to complete within a few minutes. Nonetheless, blackout periods were configured from 6:00 a.m. to 7:00 p.m. weekdays to ensure network scans were not performed during “normal business hours.”

A notification rule was defined to send an alert to the certificate services team upon discovery of either new certificates or previously unknown certificates (indicating they may have been issued and installed outside of standard processes) installations.

### 2.6.14.2 Certificate Expiration Monitoring

Significant application outages can occur when a certificate expires while in use. Consequently, it is critical that certificate owners track certificate expiration dates and replace them. The certificate services team can help certificate owners by implementing automated processes that monitor certificate expiration dates and notify the owners.

We used Venafi TPP in the lab to monitor certificate expiration dates and notify certificate owners. The methodology used in the lab followed the recommendations in *SP 1800-16 Volume B*. A weekly expiration report was scheduled giving certificate owners a list of certificates set to expire within the next 120 days. The following shows an example expiration report from the lab environment. The top of the report summarizes the status of certificates associated with a particular certificate owner.

## EXPIRATION REPORT

This report contains details about the upcoming expiration dates of your certificates. Expiration dates are displayed from most urgent to least urgent, as defined when the report was generated.

Please see Appendix for source details and other information regarding this report.

| Status    | Range      | Certificates (135) | Percentage of Total |
|-----------|------------|--------------------|---------------------|
| Expired   | 0-0 Days   | 5                  | 3.7 %               |
| Immediate | 0-5 Days   | 9                  | 6.7 %               |
| Near Term | 5-30 Days  | 35                 | 25.9 %              |
| Long Term | 30-90 Days | 86                 | 63.7 %              |

The expiration report lists all of the applicable certificates.

| Common Name                                                        | Valid To  | Contact        | Issuer          | Type | Days |
|--------------------------------------------------------------------|-----------|----------------|-----------------|------|------|
| <a href="http://9cka1wpk.tls.nccoe.org">9cka1wpk.tls.nccoe.org</a> | 2/28/2019 | Administrators | hsmBASESUBCA-CA | Prov | 0    |
| <a href="http://ck0jb30u.tls.nccoe.org">ck0jb30u.tls.nccoe.org</a> | 2/28/2019 | Administrators | hsmBASESUBCA-CA | Prov | 0    |
| <a href="http://nlc1wv8.tls.nccoe.org">nlc1wv8.tls.nccoe.org</a>   | 2/28/2019 | Administrators | hsmBASESUBCA-CA | Prov | 0    |
| <a href="http://4tpbc539.int.nccoe.org">4tpbc539.int.nccoe.org</a> | 3/1/2019  | Administrators | hsmBASESUBCA-CA | Prov | 0    |
| <a href="http://-m7pgw09.int.nccoe.org">-m7pgw09.int.nccoe.org</a> | 3/1/2019  | Administrators | hsmBASESUBCA-CA | Prov | 0    |
| <a href="http://i-8r4ol9.ext.nccoe.org">i-8r4ol9.ext.nccoe.org</a> | 3/2/2019  | Administrators | hsmBASESUBCA-CA | Prov | 1    |
| <a href="http://wdw7yww7.ext.nccoe.org">wdw7yww7.ext.nccoe.org</a> | 3/2/2019  | Administrators | hsmBASESUBCA-CA | Prov | 1    |
| <a href="http://owg82h5z.tls.nccoe.org">owg82h5z.tls.nccoe.org</a> | 3/3/2019  | Administrators | hsmBASESUBCA-CA | Prov | 2    |
| <a href="http://axz8jof2.int.nccoe.org">axz8jof2.int.nccoe.org</a> | 3/4/2019  | Administrators | hsmBASESUBCA-CA | Prov | 3    |

In addition to the reports, notification rules were configured to send emails to the owners of certificates expiring within 30 days. These notifications were configured to send daily, until the certificate was replaced. For any certificate expiring in less than 20 days, a notification rule was configured to send an additional email to escalation contacts, including the person identified as the Biz Owner and an incident response team. The objective was to minimize the amount of email that certificate owners received if all of their certificates were replaced in a timely fashion—ensuring sufficient alerts were sent for those certificates that still needed replacement.

### *2.6.14.3 Certificate Operation Monitoring*

Network discovery scans provide insight into newly installed certificates, however, it's equally important to monitor the operational state of known certificates. For example, a certificate owner may get a replacement certificate for an installed certificate set to expire. If the certificate isn't installed prior to its expiration date, an outage can result. They may install the new certificate on several but not all of the systems where the existing certificate is installed, causing the systems that were not updated to fail when the existing certificate expires. Finally, they may install the new certificate in all necessary locations, but not reset the application so the new certificate is read and use by the application, resulting in an outage, because the application is continuing to use the existing certificate that expires.

Venafi TPP provides a service call network certificate validation that automatically checks deployed certificates to ensure the correct certificate is installed and operational, thereby addressing the issues described above. If a certificate issue is detected, the certificate owner is notified. Network certificate validation was enabled on Venafi TPP in the lab.

### *2.6.14.4 Logging of Certificate-related Security Events*

Venafi TPP logs all management operations performed on certificates, including changes that administrators make within the user interfaces, changes via API, and all automated operations that are performed. Errors are also logged. All logged events are automatically stored in the Venafi TPP database. These events can be reviewed in the Venafi TPP console. It also is possible to sort, filter, and export the log events.

The following provides an example of several administrative events logged in our implementation, created by filtering on specific types of administrative events focused on configuration changes:

| Client Time            | Sev... | Event                            | Description                                                                           |
|------------------------|--------|----------------------------------|---------------------------------------------------------------------------------------|
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Object Updated        | X509 Server Certificate \VED\Policy\Certificate Management\C-DMZ\DMZE\app1.tls...     |
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Configuration Changed | User AD+adds1.pturner changed attribute X509 SubjectAltName DNS on object \...        |
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Renew Now             | Certificate renewal for \VED\Policy\Certificate Management\C-DMZ\DMZE\app1.tls...     |
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Configuration Changed | User AD+adds1.pturner changed attribute {842c5c55-d408-4904-8c26-582bce12f...         |
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Configuration Changed | User AD+adds1.pturner changed attribute Certificate Authority on object \VED\Polic... |
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Configuration Changed | User AD+adds1.pturner changed attribute Organizational Unit on object \VED\Polic...   |
| 05/01/2019 01:46:42 pm | Info   | Admin UI - Configuration Changed | User AD+adds1.pturner changed attribute X509 Subject on object \VED\Policy\Cert...    |

In addition to manually reviewing events within the console, it is possible to configure rules that will automatically send events. These events can be sent via a variety of different channels, including via email, to Splunk, to a syslog server, to an SNMP server, to a file, or to a database. Rules can be defined to send events based on specific criteria. For example, it is possible to send alerts prior to certificate expiration based on a configured set of days prior to expiration.

In our implementation, we configured Venafi TPP to send all events to the syslog server described in [Section 1.5.5.6](#).

A syslog channel was created that pointed to the syslog server.

\* Target Host:

Facility:

A rule was created to send a range of events from a severity of emergency to debug to the syslog channel.

Rules

IF Severity is between Emergency AND Debug

Target Channels

Target Channel: \VED\Logging\Channels\TLS\_LAB\_SYSLOG\_SERVERS

This approach to sending certificate-related events to an external security information and event management (SIEM) system enables all security-related events to be centralized and analyzed cohesively.

## Appendix A Passive Inspection

The example implementation demonstrates the ability to perform passive inspection of encrypted TLS connections. The question of whether or not to perform such an inspection is complex. There are important tradeoffs between traffic security and traffic visibility that each organization should consider. Some organizations prefer to decrypt internal TLS traffic, so it can be inspected to detect attacks that may be hiding within encrypted connections. Such inspection can detect intrusion, malware, and fraud, and can conduct troubleshooting, forensics, and performance monitoring. For these organizations, TLS inspection may serve as both a standard practice and a critical component of their threat detection and service assurance strategies.

The example implementation uses Symantec's SSL Visibility to perform passive inspection and is one example of how to accomplish passive inspection. The implementation demonstrates how to securely copy private keys from several different TLS servers to the SSL Visibility Appliance. The SSL Visibility Appliance can also securely replace expiring keys on servers—and immediately copy those keys to the SSL Visibility Appliance before expiration—manually and via standardized automated certificate installation.

This appendix discusses how the SSL Visibility Appliance was configured to support passive inspection. The goal was to demonstrate how to provision and revoke TLS certificates in an enterprise environment. To verify this is being done, analysis of the traffic between the TLS clients and the TLS servers was executed. The SSL Visibility Appliance can inspect traffic while located in line between the TLS clients and TLS servers on the network, or it can perform passive observation of all the network traffic between all the clients and servers mirrored to a port accessible to the server. The TLS lab configured its switching fabric to support passive monitoring of traffic utilizing traffic mirroring.

Mirroring the traffic from the virtual TLS lab environment to its physical appliances presented a few challenges. The TLS lab environment is housed within a larger VMWare and physical networking architecture. VMWare's Virtual Distributed Switch Virtual Distributed Switch (VDS) provides a centralized interface for the virtual machines' access switching in the larger NCCoE environment where the TLS lab lives as a resident. The TLS lab also has its own physical switching connections several routing hops away from the NCCoE datacenter where VMWare resides. The VDS can route traffic internally between multiple labs and virtual machines within each lab. However, VDS does not mirror VMWare's local east-west traffic between virtual machines to other physical systems outside of the VDS environment. This design limits the traffic that can be mirrored from TLS' virtual machines that live on VMWare to physical switches in the TLS lab.

To remediate this issue, the NCCoE IT team worked with VMWare senior engineers on a solution. VMWare advised the NCCoE IT team to configure remote SPAN (RSPAN) on the VDS. The IT team mapped the traffic to a RSPAN port that resided in a VLAN on an external switch. This external switch connects all the VMWare TLS hosts to the physical TLS lab. An additional RSPAN instance was configured

on the TLS lab external switch, which is a physical NCCoE-managed and controlled device connected to all the TLS team-managed and controlled physical internal switches. The external switch was configured to carry the RSPAN traffic to the internal physical access switch in the TLS lab. A SPAN was created on the internal access switch in the TLS lab and configured as source from the RSPAN VLAN. The destination was set to the physical interface connected to the SSL Visibility Appliance.

Network packets captured from VMWare vSphere workloads must be forwarded to the physical remote monitoring appliance; the packet must traverse the switch fabric between the VMWare ESXi cluster and the physical remote monitoring appliance. Two factors must be considered from a solution feasibility perspective:

- **Low end switches**—Have limitations on how many Remote SPAN sessions can be configured to run concurrently. The switch fabric must establish a Remote SPAN Session between the VMWare ESXi cluster and physical remote monitoring appliance. An alternative solution is to deploy a robust network physical tap in lieu of leveraging the switch fabric between the VMWare ESXi cluster and physical remote monitoring appliance.
- **VMWare vSphere workloads**—VMWare High Availability Features move from one ESXi host to another, as computer resources are monitored and workloads are rescheduled. This requires the ESXi cluster to automatically re-route the path that captured packets will take from a given VM workload, as it moves from one ESXi host to another when migrated or when rescheduled by Distributed Resource Scheduler to run on another host. The captured packets must egress the ESXi cluster from the specific ESXi host on which the VM workload is running.

Successful deployment of this use case requires selection of the appropriate VMWare vSphere 6.x Port Mirroring configuration option. VMWare vSphere 6.x offers 5 options:

- Distributed Port Mirroring
- Remote Mirroring Source
- Remote Mirroring Destination
- Encapsulated Remote Mirroring (L3) Source
- Distributed Port Mirroring (Legacy)

This use case that depends on the switch fabric having a Remote SPAN configured to pass traffic between the VMWare ESXi cluster and the physical remote monitoring appliance, option 2, Remote Mirroring Source, is the appropriate choice. When configured, this option will establish a Remote SPAN VLAN that will span the VMWare distributed switch. It also utilizes the physical switch fabric and leverages a distributed port group mapped to a pre-selected/pre-configured NIC on each ESXi host in the ESXi cluster. Packets are automatically re-routed from captured VM workloads that are transient between the ESXi hosts in a VMWare vSphere ESXi cluster. When a VM workload moves, vSphere will

note the change of the networking state of the VM and automatically re-establish an egress path for captured packets on the NIC of the ESXi host on which the VM is running.

## Appendix B Hardening Guidance

Hardening secures systems to reduce their vulnerabilities and minimizes the attack surface, which improves security. To harden the systems, the TLS team implemented the Defense Information Agency's Security Technical Implementation Guides (STIGs). STIGs are technical configurations applied to systems to maintain their security posture. This hardening guidance provides the baseline standard for a variety of Operating Systems—see the link below to download the STIG guidance:

<https://public.cyber.mil/stigs/>

NIST's Security Content Automation Protocol (SCAP) is used to generate compliance reports of the security health of systems. To further strengthen security of systems, use SCAP in conjunction with STIGs. Nessus is another option that can scan for vulnerabilities and misconfigurations.

STIGs are implemented through GPOs that define policy settings for computer and user settings across the network. Configure GPOs in AD to comply with STIGs. Refer to the link below to download the current DISA STIG GPO Package and select those applicable to your environment.

<https://public.cyber.mil/stigs/gpo/>

Follow the steps below to implement STIGs using GPOs in AD:

1. Open Group Policy Management Console (GPMC):
  - a. Go to **Start > Administrative Tools > Group Policy Management**.
2. Create an OU in the domain:
  - a. Go to **GPMC >** right-click on the **<YOUR DOMAIN>** > click **New Organizational Unit**.
  - b. In the Name box on the New OU dialog box, type a descriptive name for the OU > click **OK**.
3. Create a GPO in the domain:
  - a. Go to **GPMC > <YOUR DOMAIN>** > right-click **Group Policy Objects** > click **New**.
  - b. In **New GPO** dialog box enter a descriptive name > click **OK**.
4. Import DISA GPOs:
  - a. Go to **GPMC > <YOUR DOMAIN>** > **Group Policy Objects** > right-click on the GPO to edit > click **Import Settings**.
  - b. The **Import Settings Wizard** appears > click **Next** > select the folder location of the DISA GPO being used. The TLS lab used GPOs for MS Computer, MS User, DC Computer and DC User.

Note: To apply desired security configurations edit settings in the specific GPO.

5. Edit a GPO in the domain, an OU, or the Group Policy objects folder:
  - a. Go to **GPMC > <YOUR DOMAIN>** > select **Group Policy Objects** to display all GPOs in the domain.
  - b. Right-click the desired GPO > click **Edit** > the GPO will open in the Group Policy Management Editor (GPME).
  - c. In the GPME, edit the Group Policy settings as preferred.
6. Link a GPO to a domain or OU:
  - a. Go to **GPMC**> right-click **<YOUR DOMAIN>** or OU to link to the GPO > click **Link an Existing GPO**.
  - b. The **Select GPO** dialog box appears - > select the GPO you want linked to the domain or OU > click **OK**.

\*Shortcut: Drag the GPO from the Group Policy Objects folder and drop it onto the OU you want it linked to.
7. Optional:
  - Unlink a GPO from a domain or OU:
    - a. Go to **GPMC** > click **<YOUR DOMAIN>** or OU containing the GPO you want to unlink.
    - b. Right-click the **GPO** > click **Delete**.
    - c. In the Group Policy Management dialog box, confirm deletion and click **OK**.

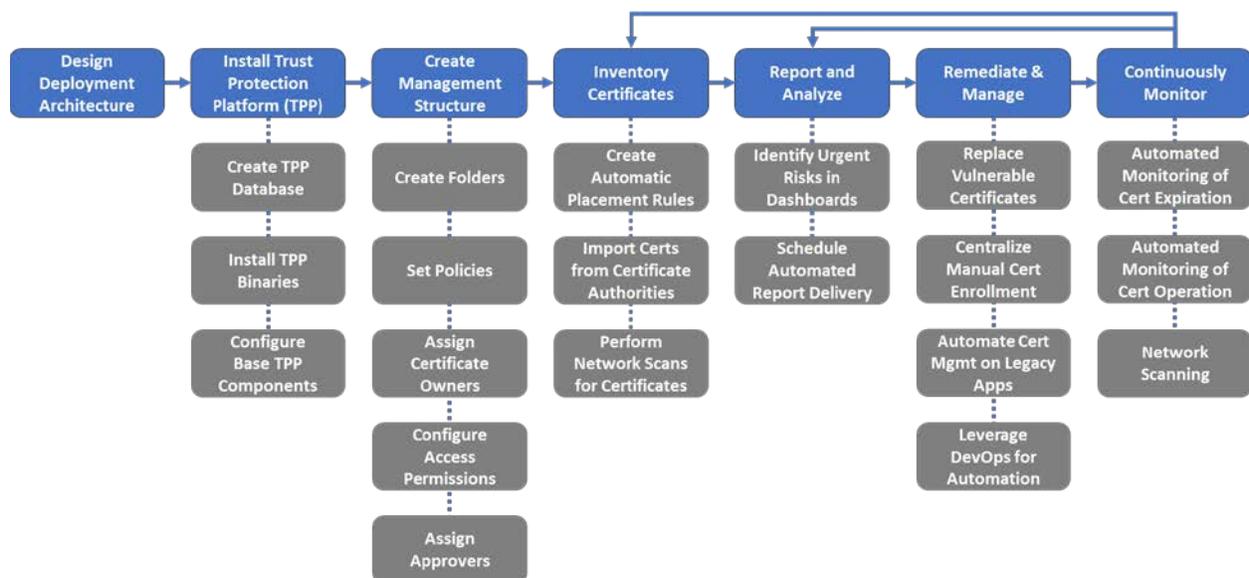
Note: Unlink a GPO when it no longer applies. Unlinking a GPO from a domain or OU does not delete the GPO—it deletes the link. After unlinking the GPO, you can still find it in the Group Policy Objects folder.
  - Add computer to OU:
    - a. Go to **Start > Administrative Tools > Active Directory Users and Computers**.
    - b. Click on **<YOUR DOMAIN>** > refresh. The newly added OU will appear.
    - c. Go to **Computers** > right-click the desired computer > click **Move**.
    - d. Select the desired OU to move the computer to > **click OK**.
    - e. To apply new settings > log out and log back in.

## Appendix C Venafi Underlying Concepts

The following background information may help users better understand some of the configurations we made in the configuration management databases (CMDBs) implementation of Venafi TPP.

Venafi TPP is one machine identity protection platform that enables enterprises to address TLS server certificate security and operational risks. Venafi TPP served as the certificate management platform for the TLS lab.

The following diagram illustrates the process of architecting, deploying, configuring, and using Venafi TPP to manage certificates and keys in enterprises.



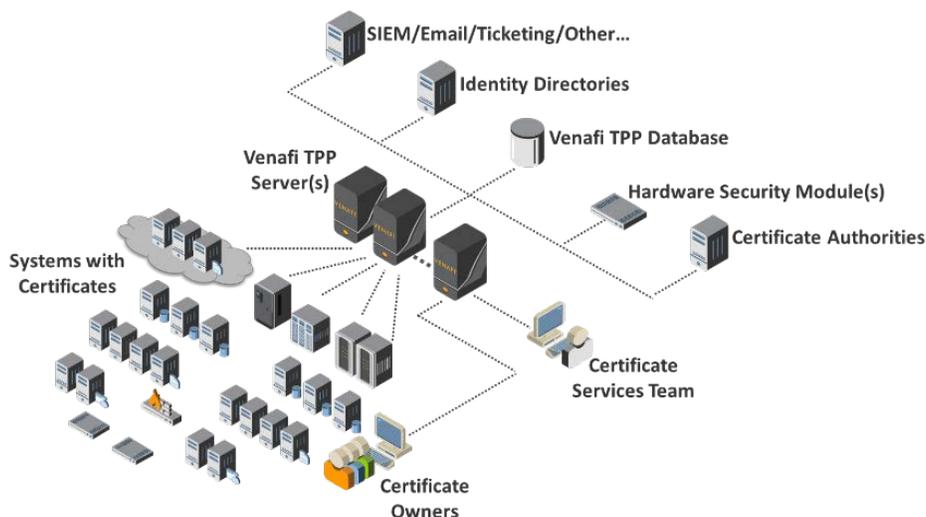
Venafi TPP interfaces with a variety of different types of systems and people/groups, including:

1. **Venafi TPP Database:** Venafi TPP requires a database to store certificates, private keys, and configuration information (all private keys and credentials are encrypted prior to storage in the database). Venafi TPP supports the use of Microsoft SQL Server to host its database.
2. **HSM:** Stores and protects the symmetric key used to encrypt private keys and credentials in the Venafi TPP database.
3. **Identity Directory:** Venafi TPP integrates with identify management systems such as AD, LDAP directories, or proprietary directories, and enables the use of existing user accounts and groups.
4. **CAs:** Venafi TPP integrates supports direct integration with over two dozen public and private CAs for the automated enrollment, renewal, and revocation of certificates.
5. **SIEM/Email/Ticketing:** Venafi TPP integrates with SIEM systems to pass certificate and cryptographic key event information. It integrates with ticketing systems for the automated

creation of change tickets and approvals and with email systems for the notifications to certificate owners for impending expirations or errors.

6. **Other Enterprise Systems:** Venafi TPP can be integrated with a variety of other enterprise systems, such as CMDBs, enterprise dashboards, and custom applications.
7. **Systems with Certificates:** Venafi TPP communicates directly with systems with certificates to automatically discover and manage those certificates.
8. **Certificate Services Team:** This team manages the Venafi TPP servers and supports Certificate Owners.
9. **Certificate Owners:** These are groups and individuals responsible for systems where certificates are deployed using Venafi TPP for automating a variety of functions, including scanning, inventory, enrollments, and installation of certificates.

The following diagram is a high-level view of these components.



Depending on an organization's needs, it's possible to deploy one or more Venafi TPP servers centrally or distributed in different network zones as well as different geographies. The number and placement of Venafi TPP servers is an important step to create an effective certificate management solution that supports the environmental and operational needs of an enterprise. The criteria driving the number and placement of Venafi TPP servers includes:

1. **Venafi TPP Services:** Each Venafi TPP can host one or more services, including network discovery scanning, certificate enrollment, certificate installation, administrative UI, etc. Depending on the size and structure of an organization, these services can be deployed on a single Venafi TPP server or, more likely, across multiple servers. The services that a Venafi TPP server can be configured to perform include:
  - a. Hosting administrative and user interfaces

- b. Network discovery scanning
- c. Onboard discovery
- d. CA import
- e. Certificate expiration monitoring
- f. Certificate operation monitoring (validation)
- g. Automated certificate enrollment
- h. Agentless certificate installation
- i. Agent management
- j. CRL expiration monitoring
- k. Revocation status monitoring
- l. Report generation
- m. Venafi TPP REST API access
- n. Log event management and notifications
- o. Trust store management

2. **Load and Performance Requirements:** The number of certificates and systems that must be managed by Venafi TPP plays an important part in the choice of how many Venafi TPP servers to deploy. Venafi TPP is based on a load-balanced architecture that enables multiple servers to share in the processing of work.
3. **Fault Tolerance:** Due to the critical role of certificate management, deployment architectures may include multiple Venafi TPP servers deployed across primary and disaster recovery sites to ensure continuous availability of certificate management services.
4. **Network Zones and Boundaries:** Network architectures often place limits on the type of traffic that can traverse between network zones (across firewalls). For example, a firewall may limit the allowed ports between two network zones, necessitating the placement of a Venafi TPP server directly inside a network zone to enable network discovery scans to run.
5. **Geographic Distribution:** Organizations are often distributed across multiple cities, states, countries, and continents. Ensuring that network latencies do not negatively impact the performance of certificate management services at each geographic location often involves distributing Venafi TPP servers near the systems and certificates being managed.

## C.1 Venafi TPP Object Model

To understand how Venafi TPP maintains inventory information, first review the Venafi TPP data model. Venafi TPP uses an object-based storage model where configuration information for certificates, associated devices, and applications are stored as objects and attributes in the Venafi TPP database. Several different object types exist in Venafi TPP—each of which includes associated attributes that store data relevant to the object. For example, a certificate object includes attributes for issuer, key length, common name, organization, etc.

The object types in Venafi TPP include:

1. **Folder:** Folders are containers that facilitate the hierarchical organization certificates, devices, applications, and other objects within Venafi TPP.
2. **Certificate:** These objects hold configuration data for certificates managed by Venafi TPP, including certificate authority (CA), key length, certificate owner, approver, and other information. A certificate object can have one or more applications objects—each indicating a location where the certificate is installed.
3. **Device:** These objects hold configuration information about the systems where certificates are deployed, including the network address and port, authentication credentials, and other information for the system.
4. **Application:** These objects hold information about the specific application (e.g., Apache, F5, Java, etc.) that uses a certificate on a device. Each device may have one or more applications that use certificates. The attributes and information stored in an application object depends on the type of application. For example, an F5 application object stores information such as the SSL profile, virtual server, and partition for the associated certificate on the F5 device.
5. **Workflow:** Workflow objects store the rules that are enforced for workflow gates within Venafi TPP. They include the stage of the certificate lifecycle where approval is needed, the required approvers, and even actions that may be automatically perform when the workflow gate is triggered.
6. **CA Template:** These objects store information about CAs from which Venafi TPP requests certificates and the specific certificate templates that the CAs will use.
7. **Credential:** These objects hold credential information that Venafi TPP uses to authenticate to other systems, including CAs, systems where certificates are managed via agentless management, etc. Passwords and private keys used in credentials are stored in encrypted form in the Venafi TPP database.

## C.2 Certificate Metadata in Venafi TPP

Certificates are stored in Venafi TPP in binary form (i.e., the DER encoded version of the certificate). In addition, the individual X.509 fields and extensions of each certificate are parsed and stored in unique database fields, to enable rapid searching and filtering. The certificate fields parsed and stored for rapid searching in Venafi TPP include:

- **X.509 Version:** V1, V2, or V3
- **Serial Number:** A unique identifier assigned by the issuing certificate authority
- **Issuer Distinguished Name:** The full X.500 distinguished name of the issuing-CA.
- **Valid From:** The date and time from which the certificate was issued. This is commonly referred to as an issue date.

- **Valid To:** The date and time after which the certificate should no longer be considered valid. This is commonly referred to as the expiration date.
- **Subject Distinguished Name (SAN):** The full X.500 distinguished name for the subject of the certificate (the entity to which the certificate was issued)—for example: “CN = iis2.int-nccoe.org, O = NCCOE, L = Gaithersburg, S = Maryland, C = US”.
- **Subject Alternative Names:** One or more identifiers for the subject of the certificate (the entity to which the certificate was issued). There could be additional DNS host names (e.g., server1.int-nccoe.org), IP address, or other types of identifiers.
- **Signature Algorithm:** The asymmetric and hashing algorithms that sign the certificate (e.g., sha256RSA).
- **Subject Key Identifier:** A unique identifier for the public key within the certificate. Because the public and private key are inextricably associated, this identifier applies to both of them.
- **Authority Key Identifier:** A unique identifier for the public/private key that the certificate authority uses to sign the certificate.
- **CRL Distribution Points:** One or more addresses where the CRL for the CA that issued the certificate can be retrieved.
- **AIA:** The location(s) where information and services, such as where to retrieve the CA certificate chain or access online certificate status protocol for the CA that issued the certificate.
- **Key Usage:** Defines the purposes for which the key within the certificate can be used, including digital signature, key encipherment, and key agreement.
- **Enhanced Key Usage:** Defines the purposes for which the certified public key within the certificate may be used, including server authentication, client authentication, and code signing.
- **Basic Constraints:** Defines whether the subject of the certificate is a CA and the maximum depth of certification path (number of CAs below this CA allowed).
- **Policy:** Policies defined within the certificate.
- **Key Size:** The length of the public key in the certificate.

In addition to certificate field and extension information, Venafi TPP stores other metadata relevant to each certificate, including:

- **Certificate Owner(s):** Groups and/or individual assigned to manage and receive notifications (e.g., expiration notices, processing errors, etc.) for the certificate
- **Approver(s):** Groups and/or individuals assigned to approve operations for the certificate
- **Processing Status:** Indicates whether the certificate processing is proceeding normally, is in error, or has completed

- **Processing Stage:** The current stage of processing (e.g., creating CSR, retrieving certificate from CA, installing certificate) for the certificate
- **Last Network Validation Time & Date:** The last date and time a network validation was performed to determine the operational status of the certificate
- **Network Validation Status:** The result of last network validation
- **Installation Location(s):** The devices and applications where the certificate is installed
- **CA Chain:** The chain of CA certificates from the root to the TLS server certificate
- **Management Method:** Determines if the certificate should be automatically enrolled and installed, or manually enrolled and installed
- **Log Information:** Logs of all administrative changes and automated operations performed on the certificate via Venafi TPP

### C.3 Custom Fields

With thousands of certificates, it is critical that organizationally-relevant information—such as cost center, application identifiers, business unit, and applicable regulations—can be associated with certificates. As a result, searches and reporting can return the certificates most relevant to a particular group or business function. Venafi TPP supports the definition of “custom fields” that can be assigned to certificates. The value of the custom fields (e.g., Cost Center = “B123”) can be assigned to individual certificates or folders, thereby flowing down and applying to all subordinate certificates. It should be noted that custom fields can be assigned to other assets such as devices associated with certificates.

#### C.3.1 Organizing Certificate Inventory

Many large enterprises have thousands or tens of thousands of certificates, often with hundreds of certificate owners across many different groups. To help effectively manage certificates across these broad environments, Venafi TPP enables the creation of a hierarchical folder structure where certificates and associated system configuration information can be placed.

The design of a Venafi TPP folder hierarchy for the organization of certificates is dependent on the needs and requirements of an enterprise—similar to having multiple approaches to create folder hierarchies when organizing files. However, through experience in working with many large enterprises, Venafi professional services has developed a set of guidelines, including:

- **Certificate Ownership:** The primary factor for designing a Venafi TPP hierarchy is based on the organization of certificate owners. Once a folder is assigned to a certificate owner, certificates and other assets placed within the folder automatically inherit the permissions, contacts, and approvers, so that ownership does not need to be managed on individual certificates (though ownership information can be managed on individual certificates in Venafi TPP, if necessary).

- **Policies:** Policies such as allowed key lengths, signing algorithms, and CAs are an important consideration in the organization of Venafi TPP folders.
- **Workflow and Approvals:** Workflow rules are assigned at the folder level in Venafi TPP. If an enterprise applies different workflow rules across their organizational groups, the design of the folder hierarchy may be adjusted to easily assign those rules as needed.

### C.3.2 Policy Enforcement

Venafi TPP supports the enforcement of written policies through the assignment of policies to any folder within the hierarchy. It is possible to define Venafi TPP policies for a broad set of areas, including allowed CAs, allowable domains, certificate contents (e.g., key length), approvers, and application configurations.

Policies set on a folder flow down to subordinate folders and objects within the folders. This makes it possible to configure group-specific policies on folders assigned to those groups and policies with broader applicability to higher level folders, so that they apply to all certificates, devices, applications across subordinate folders. Policies can be set as suggested, to provide a default value that users are able to change if desired, or enforced, where users are required to use the set value.

## C.4 The Domain Allowlist

Because certificates serve as trusted credentials, they should only be issued for authorized domains. To aid in this, Venafi TPP supports establishing allowlists of domains that can be used in certificates. For example, it is possible to only allow common names (CNs) and subject alternative names (SANs) that have the suffix “.int-nccoe.org”, which only allow CNs and SANs such as server1.int-nccoe.org and server2.ops.int-nccoe.org.

### C.4.1 Certificate Owner Assignment

The assignment and maintenance of certificate ownership is critical to prevent outages and respond to security incidents. Depending on the size of groups and the number certificates they manage, certificate management responsibilities may be assigned to one person or distributed among several different individuals. For larger groups managing greater numbers of certificates across a broad set of systems, the roles may vary for each team member. For example, a core group of technical people may be responsible for managing the configuration of certificates. That same group plus a manager may need to receive alerts and reports. To accommodate these differences in roles, Venafi TPP enables the assignment of permissions and contact information (for sending alerts) at the certificate or folder level.

### C.4.2 Permissions

In Venafi TPP, groups and individual users can be granted permissions to folders and individual objects (e.g., certificates). Venafi TPP can assign the following permissions:

- **View:** See an object in a folder and select it (but not see its configuration parameters). For example, an administrator with view rights to an application can associate that application to a certificate for which they are responsible.
- **Read:** Read an object’s configuration parameters and status.
- **Write:** Edit an object’s configuration parameters.
- **Create:** Create new objects under the object to which the Create permission is assigned. Applies only to objects that contain other objects.
- **Delete:** Delete the specified object or objects contained within it (unless blocked below).
- **Rename:** Rename the object.
- **Revoke:** Revoke a certificate. This only applies to certificates only but can be set on policies, devices, or applications for any certificates contained under them.
- **Associate:** Associate a certificate to one or more applications from within that certificate object.
- **Admin:** Grant users or groups permissions to the object.
- **Private-Key Read:** Retrieve the private-key for a certificate only applies to certificates but can be set on policies, devices, or applications for any certificates contained under them.
- **Private-Key Write:** Upload or overwrite the private-key for a certificate. This only applies to certificates but can be set on policies, devices, or applications for any certificates contained within them. The private-key write privilege is required for an administrator to extract a private-key and certificate from an application to be stored in the Venafi TPP database.
- **Permissions:** Permissions assigned to a folder are inherited subordinate objects and folders. Wherever possible, it’s a best practice to assign permissions to groups to quickly grant a new team member the needed permissions simply by being added to the group. It is also best to assign permissions at the folder level, applying to all subordinate certificates. When a new system and certificate are needed, they can be added within the folder and the permissions automatically apply.

### C.4.3 Contacts

Effectively managing certificates in an enterprise requires the ability to automatically notify the certificate owners of impending expirations, errors, or other events that affect their certificates. It’s possible to assign one or more groups or individuals as “contacts” to folders or individual objects in Venafi TPP. Contact assignment to folders are inherited by the objects below them.

## Appendix D List of Acronyms

|               |                                                                                                                                         |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>ACME</b>   | Automated Certificate Management Environment                                                                                            |
| <b>AD</b>     | Active Directory                                                                                                                        |
| <b>ADCS</b>   | Active Directory Certificate Services                                                                                                   |
| <b>ADS</b>    | Active Directory Services                                                                                                               |
| <b>AIA</b>    | Authority Information Access                                                                                                            |
| <b>API</b>    | Application Programming Interface                                                                                                       |
| <b>CA</b>     | Certificate Authority                                                                                                                   |
| <b>CAPI</b>   | Cryptographic Application Programming Interface (also known variously as CryptoAPI, Microsoft Cryptography API, MS-CAPI or simply CAPI) |
| <b>CDP</b>    | CRL Distribution Point                                                                                                                  |
| <b>CEP</b>    | Certificate Enrollment Policy                                                                                                           |
| <b>CES</b>    | Certificate Enrollment Service                                                                                                          |
| <b>CMDB</b>   | Configuration Management Database                                                                                                       |
| <b>CN</b>     | Common Name                                                                                                                             |
| <b>CNG</b>    | Cryptography API: Next Generation                                                                                                       |
| <b>CPU</b>    | Central Processing Units                                                                                                                |
| <b>CRL</b>    | Certificate Revocation List                                                                                                             |
| <b>CSR</b>    | Certificate Signing Request                                                                                                             |
| <b>DB</b>     | Database                                                                                                                                |
| <b>DC</b>     | Domain Controller                                                                                                                       |
| <b>DevOps</b> | Development Operations                                                                                                                  |
| <b>DMZ</b>    | Demilitarized Zone                                                                                                                      |
| <b>DNS</b>    | Domain Name System                                                                                                                      |
| <b>EULA</b>   | End User License Agreement                                                                                                              |

|              |                                                 |
|--------------|-------------------------------------------------|
| <b>EV</b>    | Extended Validation                             |
| <b>FIPS</b>  | Federal Information Processing Standards        |
| <b>FQDN</b>  | Fully Qualified Domain Name                     |
| <b>GPMC</b>  | Group Policy Management Console                 |
| <b>GPO</b>   | Group Policies Objects                          |
| <b>HSM</b>   | Hardware Security Module                        |
| <b>HTML</b>  | Hypertext Markup Language                       |
| <b>http</b>  | Hypertext Transfer Protocol                     |
| <b>https</b> | Hypertext Transfer Protocol Secure              |
| <b>IdP</b>   | Identity Provider                               |
| <b>IETF</b>  | Internet Engineering Task Force                 |
| <b>IIS</b>   | Internet Information Server (Microsoft Windows) |
| <b>IMAP</b>  | Internet Message Access Protocol                |
| <b>IP</b>    | Internet Protocol                               |
| <b>IT</b>    | Information Technology                          |
| <b>ITL</b>   | Information Technology Laboratory               |
| <b>KSP</b>   | Key Storage Provider                            |
| <b>LDAP</b>  | Lightweight Directory Access Protocol           |
| <b>LTM</b>   | Local Traffic Manager (F5)                      |
| <b>MSQL</b>  | Microsoft SQL                                   |
| <b>MTA</b>   | Mail Transfer Agent                             |
| <b>MUA</b>   | Mail User Agent                                 |
| <b>NAT</b>   | Network Address Translation                     |
| <b>NCCoE</b> | National Cybersecurity Center of Excellence     |
| <b>NIST</b>  | National Institute of Standards and Technology  |

|                   |                                                             |
|-------------------|-------------------------------------------------------------|
| <b>NTL</b>        | Network Trust Link                                          |
| <b>NTLS</b>       | Network Trust Link Service                                  |
| <b>OS</b>         | Operating System                                            |
| <b>OVA</b>        | Open Virtualization Appliance                               |
| <b>OVF</b>        | Open Virtualization Format                                  |
| <b>PCI-DSS</b>    | Payment Card Industry Data Security Standard                |
| <b>PED</b>        | PIN Entry Device                                            |
| <b>PIN</b>        | Personal Identification Number                              |
| <b>PKI</b>        | Public Key Infrastructure                                   |
| <b>PSCP</b>       | PuTTY Secure Copy Protocol                                  |
| <b>RA</b>         | Registration Authority                                      |
| <b>RAM</b>        | Random Access Memory                                        |
| <b>REST</b>       | Representational State Transfer (API)                       |
| <b>RHEL</b>       | Red Hat Enterprise Linux                                    |
| <b>RMF</b>        | Risk Management Framework                                   |
| <b>RSA</b>        | Rivest, Shamir, & Adleman (public key encryption algorithm) |
| <b>RSPAN</b>      | Remote Switched Port Analyzer                               |
| <b>Thales TCT</b> | Thales Trusted Cyber Technologies                           |
| <b>SAN</b>        | Subject Alternative Name                                    |
| <b>SCAP</b>       | Security Content Automation Protocol                        |
| <b>SCEP</b>       | Simple Certificate Enrollment Protocol                      |
| <b>SCP</b>        | Secure Copy Protocol                                        |
| <b>SIEM</b>       | Security Information and Event Management                   |
| <b>SMTP</b>       | Simple Mail Transfer Protocol                               |
| <b>SOAP</b>       | Simple Object Access Protocol                               |

|                       |                                          |
|-----------------------|------------------------------------------|
| <b>SP</b>             | Special Publication                      |
| <b>SPAN</b>           | Switched Port Analyzer                   |
| <b>SQL</b>            | Structured Query Language                |
| <b>SSL</b>            | Secure Socket Layer (protocol)           |
| <b>SSL VISIBILITY</b> | SSL Visibility (Symantec Appliance)      |
| <b>STIGs</b>          | Security Technical Implementation Guides |
| <b>TCP</b>            | Transmission Control Protocol            |
| <b>TLS</b>            | Transport Layer Security (protocol)      |
| <b>TMSH</b>           | Traffic Management Shell                 |
| <b>TPP</b>            | Trust Protection Platform (Venafi)       |
| <b>UCS</b>            | User Configuration Set                   |
| <b>UDP</b>            | User Datagram Protocol                   |
| <b>UPN</b>            | User Principal Name                      |
| <b>URL</b>            | Uniform Resource Locator                 |
| <b>VDS</b>            | Virtual Distributed Switch               |
| <b>VE</b>             | Virtual Edition                          |
| <b>VLAN</b>           | Virtual Local Area Network               |
| <b>WinRM</b>          | Windows Remote Management                |

## Appendix E Glossary

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Active Directory</b>                             | A Microsoft directory service for the management of identities in Windows domain networks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Application</b>                                  | <ol style="list-style-type: none"><li>1. The system, functional area, or problem to which information technology (IT) is applied. The application includes related manual procedures as well as automated procedures. Payroll, accounting, and management information systems are examples of applications. (<a href="#">NIST SP 800-16</a>)</li><li>2. A software program hosted by an information system. (<a href="#">NIST SP 800-137</a>)</li></ol>                                                                                                             |
| <b>Authentication</b>                               | Verifying the identity of a user, process, or device, often as a prerequisite to allowing access to a system's resources. ( <a href="#">NIST SP 800-63-3</a> )                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Automated Certificate Management Environment</b> | A protocol defined in IETF RFC 8555 that provides for the automated enrollment of certificates.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Certificate</b>                                  | A set of data that uniquely identifies an entity, contains the entity's public key and possibly other information, and is digitally signed by a trusted party, thereby binding the public key to the entity. Additional information in the certificate could specify how the key is used and its validity period. ( <a href="#">NIST SP 800-57 Part 1 Rev. 4 [3]</a> under Public-key certificate) (Certificates in this practice guide are based on ( <a href="#">IETF RFC 5280</a> .)                                                                             |
| <b>Certificate Authority</b>                        | A trusted entity that issues and revokes public key certificates. ( <a href="#">NISTIR 8149</a> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Certificate Chain</b>                            | An ordered list of certificates that starts with an end-entity certificate, includes one or more certificate authority (CA) certificates, and ends with the end-entity certificate's Root CA certificate, where each certificate in the chain is the certificate of the CA that issued the previous certificate. By checking to see if each certificate in the chain was issued by a trusted CA, the receiver of an end-user certificate can determine whether it should trust the end-entity certificate by verifying the signatures in the chain of certificates. |

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Certificate Management</b>      | Process whereby certificates (as defined above) are generated, stored, protected, transferred, loaded, used, and destroyed. ( <a href="#">CNSSI 4009-2015</a> ) (In the context of this practice guide, it also includes inventory, monitoring, enrolling, installing, and revoking.)                                                                                                                                                                                                                                                                                                                                              |
| <b>Certificate Revocation List</b> | A list of digital certificates that have been revoked by an issuing CA before their scheduled expiration date and should no longer be trusted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Certificate Signing Request</b> | A request sent from a certificate requester to a CA to apply for a digital identity certificate. The certificate signing request contains the public key as well as other information to be included in the certificate and is signed by the private key corresponding to the public key.                                                                                                                                                                                                                                                                                                                                          |
| <b>Client</b>                      | <ol style="list-style-type: none"> <li>1. A machine or software application that accesses a cloud over a network connection, perhaps on behalf of a consumer. (<a href="#">NIST SP 800-146</a>)</li> <li>2. A function that uses the PKI to obtain certificates and validate certificates and signatures. Client functions are present in CAs and end entities. Client functions may also be present in entities that are not certificate holders. That is, a system or user that verifies signatures and validation paths is a client, even if it does not hold a certificate itself. (<a href="#">NIST SP 800-15</a>)</li> </ol> |
| <b>Cloud Computing</b>             | A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. ( <a href="#">NIST SP 800-145</a> )                                                                                                                                                                                                                                                                                         |
| <b>Common Name</b>                 | An attribute type commonly found within a Subject Distinguished Name in an X.500 directory information tree. When identifying machines, it is composed of a fully qualified domain name or IP address.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Configuration Management</b>    | A collection of activities focused on establishing and maintaining the integrity of IT products and information systems, through control of processes for initializing, changing, and monitoring the configurations of those products and systems throughout the system development life cycle. ( <a href="#">NIST SP 800-53 Rev. 4</a> )                                                                                                                                                                                                                                                                                          |

|                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Container</b>                                       | A method for packaging and securely running an application within an application virtualization environment. Also known as an application container or a server application container. ( <a href="#">NIST SP 800-190</a> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Cryptographic Application Programming Interface</b> | An application programming interface (API) included with Microsoft Windows operating systems that provides services to enable developers to secure Windows-based applications using cryptography. While providing a consistent API for applications, the Cryptographic Application Programming Interface (CAPI) allows for specialized cryptographic modules (cryptographic service providers) to be provided by third parties, such as Hardware Security Module (HSM) manufacturers. This enables applications to leverage the additional security of HSMs while using the same APIs they use to access built-in Windows cryptographic service providers. (Also known variously as CryptoAPI, Microsoft Cryptography API, MS-CAPI or simply CAPI) |
| <b>Cryptography API: Next Generation</b>               | The long-term replacement for the CAPI.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Demilitarized Zone</b>                              | A perimeter network or screened subnet separating a more-trusted internal network from a less-trusted external network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Development Operations (DevOps)</b>                 | A set of practices for automating the processes between software development and IT operations teams, so they can build, test, and release software faster and more reliably. The goal is to shorten the systems development life cycle and improve reliability while delivering features, fixes, and updates frequently in close alignment with business objectives.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Digital Certificate</b>                             | Certificate (as defined above).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Digital Signature</b>                               | The result of a cryptographic transformation of data that, when properly implemented, provides origin authentication, assurance of data integrity and signatory non-repudiation. ( <a href="#">NIST SP 800-133</a> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Digital Signature Algorithm</b>                     | A Federal Information Processing Standard for digital signatures, based on the mathematical concept of modular exponentiations and the discrete logarithm problem. ( <a href="#">FIPS 186-4</a> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Directory Service</b>                               | A distributed database service capable of storing information, such as certificates and CRLs, in various nodes or servers distributed across a network. ( <a href="#">NIST SP 800-15</a> ) (In the context of this practice                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

|                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                        | guide, a directory services stores identity information and enables the authentication and identification of people and machines.)                                                                                                                                                                                                                                                                                                                             |
| <b>Distinguished Name</b>                              | An identifier that uniquely represents an object in the X.500 directory information tree. ( <a href="#">RFC 4949 Ver 2</a> )                                                                                                                                                                                                                                                                                                                                   |
| <b>Domain</b>                                          | A distinct group of computers under a central administration or authority.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Domain Name</b>                                     | A label that identifies a network domain using the Domain Naming System.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Domain Name System</b>                              | The system by which Internet domain names and addresses are tracked and regulated as defined by <a href="#">IETF RFC 1034</a> and other related RFCs.                                                                                                                                                                                                                                                                                                          |
| <b>Extended Validation (EV) Certificate</b>            | A certificate used for https websites and software that includes identity information, subjected to an identity verification process standardized by the CA Browser Forum in its <a href="#">Baseline Requirements</a> which verifies the identified owner of the website for which the certificate has been issued has exclusive rights to use the domain; exists legally, operationally, and physically; and has authorized the issuance of the certificate. |
| <b>Federal Information Processing Standards (FIPS)</b> | A standard for adoption and used by federal departments and agencies that has been developed within the Information Technology Laboratory (ITL) and published by the National Institute of Standards and Technology, a part of the U.S. Department of Commerce. A FIPS covers some topic in IT to achieve a common level of quality or some level of interoperability. ( <a href="#">NIST SP 800-161</a> )                                                     |
| <b>Hardware Security Module (HSM)</b>                  | A physical computing device that provides tamper-evident and intrusion-resistant safeguarding and management of digital keys and other secrets, as well as crypto-processing. ( <a href="#">FIPS 140-2</a> ) specifies requirements for HSMs.                                                                                                                                                                                                                  |
| <b>Host Name</b>                                       | Host names are most commonly defined and used in the context of DNS. The host name of a system typically refers to the fully qualified DNS domain name of that system.                                                                                                                                                                                                                                                                                         |
| <b>Hypertext Transfer Protocol (HTTP)</b>              | A standard method for communication between clients and Web servers. ( <a href="#">NISTIR 7387</a> )                                                                                                                                                                                                                                                                                                                                                           |

|                                                     |                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Internet Engineering Task Force (IETF)</b>       | The internet standards organization made up of network designers, operators, vendors, and researchers that defines protocol standards (e.g., IP, TCP, DNS) through process of collaboration and consensus.                                                                                                                                                    |
| <b>Internet Message Access Protocol</b>             | A method of communication used to read electronic mail stored in a remote server. ( <a href="#">NISTIR 7387</a> )                                                                                                                                                                                                                                             |
| <b>Internet Protocol (IP)</b>                       | The IP, as defined in <a href="#">IETF RFC 6864</a> , is the principal communications protocol in the IETF Internet protocol suite for specifying system address information when relaying datagrams across network boundaries.                                                                                                                               |
| <b>Lightweight Directory Access Protocol (LDAP)</b> | The LDAP is a directory access protocol. In this document, LDAP refers to the protocol defined by RFC 1777, which is also known as LDAP V2. LDAP V2 describes unauthenticated retrieval mechanisms. ( <a href="#">NIST SP 800-15</a> )                                                                                                                        |
| <b>Microservice</b>                                 | A set of containers that work together to compose an application. ( <a href="#">NIST SP 800-190</a> )                                                                                                                                                                                                                                                         |
| <b>Organization</b>                                 | An entity of any size, complexity, or positioning within an organizational structure (e.g., a federal agency or, as appropriate, any of its operational elements). ( <a href="#">NIST SP 800-39</a> ) This publication is intended to provide recommendations for organizations that manage their own networks (e.g., that have a chief information officer). |
| <b>Outage</b>                                       | A period when a service or an application is not available or when equipment is not operational.                                                                                                                                                                                                                                                              |
| <b>Payment Card Industry Data Security Standard</b> | An information security standard administered by the <a href="#">Payment Card Industry Security Standards Council</a> that is for organizations that handle branded credit cards from the major card schemes.                                                                                                                                                 |
| <b>PIN Entry Device</b>                             | An electronic device used in a debit, credit or smart card-based transaction to accept and encrypt the cardholder's personal identification number.                                                                                                                                                                                                           |
| <b>Post Office Protocol</b>                         | A mailbox access protocol defined by IETF RFC 1939. POP is one of the most commonly used mailbox access protocols. ( <a href="#">NIST SP 800-45 Version 2</a> )                                                                                                                                                                                               |

|                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Private Key</b>                            | The secret part of an asymmetric key pair that is used to digitally sign or decrypt data. ( <a href="#">NIST SP 800-63-3</a> )                                                                                                                                                                                                                                                                                                                                               |
| <b>Public CA</b>                              | A trusted third party that issues certificates as defined in IETF RFC 5280. A CA is considered public if its root certificate is included in browsers and other applications by the developers of those browsers and applications. The CA/Browser Forum defines the requirements public CAs must follow in their operations.                                                                                                                                                 |
| <b>Public Key</b>                             | The public part of an asymmetric key pair that is used to verify signatures or encrypt data. ( <a href="#">NIST SP 800-63-3</a> )                                                                                                                                                                                                                                                                                                                                            |
| <b>Public Key Cryptography</b>                | Cryptography that uses separate keys for encryption and decryption; also known as asymmetric cryptography. ( <a href="#">NIST SP 800-77</a> )                                                                                                                                                                                                                                                                                                                                |
| <b>Public Key Infrastructure (PKI)</b>        | The framework and services that provide for the generation, production, distribution, control, accounting, and destruction of public key certificates. Components include the personnel, policies, processes, server platforms, software, and workstations used for the purpose of administering certificates and public-private key pairs, including the ability to issue, maintain, recover, and revoke public key certificates. ( <a href="#">NIST SP 800-53 Rev. 4</a> ) |
| <b>Registration Authority</b>                 | An entity authorized by the certification authority system (CAS) to collect, verify, and submit information provided by potential Subscribers which is to be entered into public key certificates. The term RA refers to hardware, software, and individuals that collectively perform this function. ( <a href="#">CNSSI 4009-2015</a> )                                                                                                                                    |
| <b>Representational State Transfer (REST)</b> | A software architectural style that defines a common method for defining APIs for web services.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Risk Management Framework</b>              | The Risk Management Framework (RMF), presented in <a href="#">NIST SP 800-37</a> , provides a disciplined and structured process that integrates information security and risk management activities into the system development life cycle.                                                                                                                                                                                                                                 |
| <b>Rivest, Shamir, &amp; Adleman (RSA)</b>    | An algorithm approved in [FIPS 186] for digital signatures and in [SP 800-56B] for key establishment. ( <a href="#">NIST SP 800-57 Part 1 Rev. 4</a> )                                                                                                                                                                                                                                                                                                                       |
| <b>Root certificate</b>                       | A self-signed certificate, as defined by <a href="#">IETF RFC 5280</a> , issued by a root certificate authority. A root certificate is typically securely                                                                                                                                                                                                                                                                                                                    |

installed on systems, so they can verify end-entity certificates they receive.

**Root certificate authority**

In a hierarchical public key infrastructure (PKI), the CA whose public key serves as the most trusted datum (i.e., the beginning of trust paths) for a security domain. ([NIST SP 800-32](#))

**Subject Alternative Name**

A field in an X.509 certificate that identifies one or more fully qualified domain names, IP addresses, email addresses, URIs, or UPNs to be associated with the public key contained in a certificate.

**Simple Certificate Enrollment Protocol (SCEP)**

A protocol defined in an IETF [internet](#) draft specification that is used by numerous manufacturers of network equipment and software who are developing simplified means of handling certificates for large-scale implementation to everyday users, as well as referenced in other industry standards.

**Secure Hash Algorithm 256**

A hash algorithm that can be used to generate digests of messages. The digests are used to detect whether messages have been changed since the digests were generated. ([FIPS 180-4 \[March 2012\]](#))

**Secure Transport**

Transfer of information using a transport layer protocol that provides security between applications communicating over an IP network.

**Server**

A computer or device on a network that manages network resources. Examples include file servers (to store files), print servers (to manage one or more printers), network servers (to manage network traffic), and database servers (to process database queries). ([NIST SP 800-47](#))

**Service Provider**

A provider of basic services or value-added services for operation of a network; generally refers to public carriers and other commercial enterprises. ([NISTIR 4734](#))

**Simple Mail Transfer Protocol (SMTP)**

The primary protocol used to transfer electronic mail messages on the internet. ([NISTIR 7387](#))

**Special Publication**

A type of publication issued by NIST. Specifically, the Special Publication 800-series reports on the ITR's research, guidelines, and outreach efforts in computer security, and its collaborative activities

|                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                        | with industry, government, and academic organizations. The 1800 series reports the results of NCCoE demonstration projects.                                                                                                                                                                                                                                                                                                                                       |
| <b>System Administrator</b>            | Individual responsible for the installation and maintenance of an information system, providing effective information system utilization, adequate security parameters, and sound implementation of established Information Assurance policy and procedures. ( <a href="#">CNSSI 4009-2015</a> )                                                                                                                                                                  |
| <b>Team</b>                            | A number of persons associated together in work or activity. (Merriam Webster) As used in this publication, a team is a group of individuals assigned by an organization’s management the responsibility to carry out a defined function or set of defined functions. Designations for teams as used in this publication are simply descriptive. Different organizations may have different designations for teams that carry out the functions described herein. |
| <b>Transport Layer Security (TLS)</b>  | An authentication and security protocol widely implemented in browsers and web servers. TLS is defined by <a href="#">RFC 5246</a> and <a href="#">RFC 8446</a> .                                                                                                                                                                                                                                                                                                 |
| <b>Trust Protection Platform (TPP)</b> | The Venafi Machine Identity Protection platform used in the example implementation described in this practice guide.                                                                                                                                                                                                                                                                                                                                              |
| <b>User Principal Name</b>             | In Windows Active Directory, this is the name of a system user in email address format, i.e., a concatenation of username, the “@” symbol, and domain name.                                                                                                                                                                                                                                                                                                       |
| <b>Validation</b>                      | The process of determining that an object or process is acceptable according to a pre-defined set of tests and the results of those tests. ( <a href="#">NIST SP 800-152</a> )                                                                                                                                                                                                                                                                                    |
| <b>Web Browser</b>                     | A software program that allows a user to locate, access, and display web pages.                                                                                                                                                                                                                                                                                                                                                                                   |

## Appendix F      References

- [1] U.S. Department of Commerce, Security Requirements for Cryptographic Modules, Federal Information Processing Standards (FIPS) Publication 140-2, (including change notices as of 12-03-2002). <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.140-2.pdf>.
- [2] Joint Task Force Transformation Initiative, Security and Privacy Controls for Information Systems and Organizations, Draft NIST Special Publication (SP) 800-53 Revision 5, August 2017. <https://csrc.nist.gov/CSRC/media//Publications/sp/800-53/rev-5/draft/documents/sp800-53r5-draft.pdf>.
- [3] E. Barker, Recommendation for Key Management: Part 1: General, NIST Special Publication (SP) 800-57 Part 1, Revision 4, January 2016. <http://doi.org/10.6028/NIST.SP.800-57pt1r4>.
- [4] Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1, National Institute of Standards and Technology, April 16, 2018. See <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf>.
- [5] T. Dierks, E. Rescorla, The Transport Layer Security (TLS) Protocol Version 1.2, RFC 5246, Internet Engineering Task Force, August 2008. <https://www.ietf.org/rfc/rfc5246.txt>.

## Appendix G Supplemental Architecture Configurations

### G.1 Mail Server Configuration Files

The Postfix mail server and Dovecot mail client were both used to create an alert and administrative email server for all alerts received from the various TLS security components used in the TLS lab. The main.cf is the primary configuration file for Postfix and the dovecot.conf is used to configure the Dovecot mail user agent. Links to both files used in the TLS lab are provided below as a quick start to setting up the same mail server and client used in the TLS lab. The main.cf and dovecot.conf files are stored in the same repository as this Volume D document on the NCCoE web page.

- <https://www.nccoe.nist.gov/sites/default/files/library/supplemental-files/sp1800-16/main.cf>
- <https://www.nccoe.nist.gov/sites/default/files/library/supplemental-files/sp1800-16/dovecote.conf>