Report of the Federal PKI Certificate Policy Working Group Cloud PKI Tiger Team

Version 1.0

October 18, 2024

Abstract

This paper documents the conclusions of the PKI in the Cloud Tiger Team of the Federal PKI CPWG

Table of Contents

1. Executive Summary

In October 2022, the Federal Public Key Infrastructure (FPKI) Policy Authority (PA) tasked the Certificate Policy Working Group (CPWG) to establish a Tiger Team to investigate the possibility of implementing a Federal PKI (FPKI) compliant framework using commercially available cloud services, (e.g. Amazon Web Service, Microsoft Azure, etc.) without compromising the security requirements defined in the Federal PKI Certificate Policies.

The Tiger Team spent over a year on the following activities:

- Identifying current policy challenges to cloud adoption and documenting the purpose of each requirement in light of the risks they address,
- Engaging prospective vendors and reviewing their capabilities for cloud PKI implementations,
- Conducting a PKI component risk analysis and
- Defining important terms and critical cloud security controls to address identified policy challenges.

The Tiger Team has concluded that it is not currently possible to define a single cloud design architecture that would be compliant with Federal PKI policies, for the following reasons:

- Variety of cloud technologies available for use by providers makes it impractical to attempt to define a single set of recommendations that would cover every potential cloud PKI design or architectural selection
- Due to rapid cloud technology evolution, any recommendation for current cloud technologies could rapidly become outdated
- Numerous deployment and delivery options for cloud technologies have the potential to significantly impact the overall security posture of the solution.

The Tiger Team has determined that the best recommendation to provide at this time is a guidance document that will inform future efforts by FPKI members to evaluate cloud deployments. This guidance document incorporates:

- A list of current FPKI policies which prevent deployment of cloud hosted PKI elements
- A list of challenges introduced by cloud vs. on-premise PKI deployments
- Identification of the components of a complete CA deployment, identifying the threats for each component and the associated risk
- A list of considerations to assist FPKI members to analyze a cloud hosted implementation of PKI components.

The Tiger Team is releasing this document to facilitate Federal PKI implementers design and propose FPKI-compliant cloud based solutions and submit them for review and approval by the Federal PKI Policy Authority. Critical characteristics of any technical proposed solution will need to address all cloud equivalent controls identified in Appendix A, in addition to addressing all baseline controls included in FPKI policies and the FPKI SP 800-53 overlay.

2. Introduction

2.1. Origin of the Cloud Tiger Team

The PA created a member poll to prioritize topics for investigation by the CPWG. This poll was conducted from August to September 2022 and the highest priority identified was investigation of FPKI components in the cloud.

A Cloud Tiger Team subsequently kicked off in January 2023 with 45 initial participants identified from CPWG stakeholder organizations.

2.2. Cloud Tiger Team Approach

The Cloud Tiger Team approached this topic in three phases:

- Initial fact-finding
- Definition of a Risk Assessment Model
- Development of guidance and final recommendations

2.2.1. Initial Fact Finding

During the initial fact finding phase, the Cloud Tiger Team identified FPKI policy challenges to adoption of cloud technologies for PKI implementations, and performed a basic survey and security assessment of cloud vendors nominated by Cloud Tiger Team members.

2.2.1.0. Identify Challenges

In preparation for the Cloud Tiger Team kickoff, the FPKIPA Support Team extracted the perceived limitations to cloud implementations from FPKI certificate policy, specifically [COMMON]. These limitations generally included personnel controls, procedural technical security controls, and audit requirements from [COMMON] Sections 5, 6, and 8. Some examples of these policy limitations include Physical Security, Trusted Roles, Multi-Person Controls, Network Security Controls, System Development controls, and Audit scope, and the full list of policy limitations is available upon request to fpki@gsa.gov.

Once the Tiger Team was established and the scope and approach were determined, these policy limitations were validated and expanded to document risk categories and potential threats each policy requirement was meant to address in order to facilitate any subsequent analysis of FPKI implementations in cloud-based architectures.

2.2.1.1. Vendor Security Assessment

Once the policy limitations for cloud implementation were defined and risks based on those controls were identified, the Cloud Tiger Team developed a vendor security assessment questionnaire. This questionnaire presented several inquiries regarding the controls implemented by vendors to address the policy limitations. The Cloud Tiger Team provided the questionnaire to cloud vendors as recommended by the members.

The Cloud Tiger Team received responses from vendors offering the following services:

- HSM as a Service
- Managed PKI Service
- Infrastructure as a Service

The information provided by the vendors clarified a number of details regarding operational security practices, but none of the solutions appeared to comply with all of the documented requirements organically (aka “out of the box”). This was expected, but the information provided by the vendors informed our ongoing discussions on this topic.

2.2.2. PKI Component Risk Assessment Model

While the Tiger Team coordinated responses from vendors they also developed a high level PKI component list with associated threats. The component list broke out all architectural sub-components of a fully compliant FPKI solution and mapped each sub-component to the policy challenges identified in the previous phase.

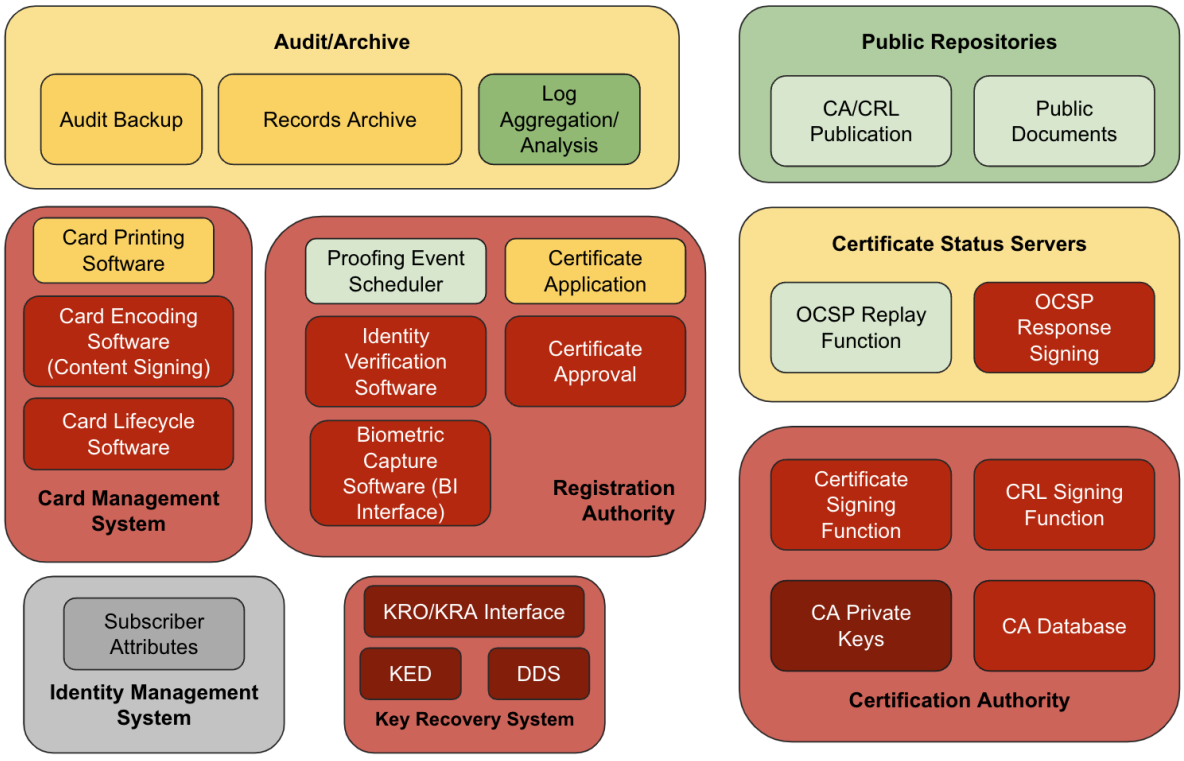

PKI Component Risk Assessment Model

The Cloud Tiger Team performed a cursory threat assessment of the subcomponents of the PKI to identify any “low-risk” components that could be prioritized for migration to the cloud. This model primarily defined the ratings based on the potential impact of a compromise on a PKI sub-component.

Provided the impact diagram produced by the Tiger Team, organizations are able to independently assess the likelihood of a threat materializing in a specific cloud environment, versus the existing on-premise solution. With the impact understood, organizations can more easily develop a comprehensive risk assessment for a proposed cloud implementation of a PKI.

3. Cloud Tiger Team Outcomes

Following the development of the risk model, the Cloud Tiger Team determined that the most effective approach to preserve our security posture while enabling future iterations of cloud hosted architectures, was to develop a guidance document for implementers who wish to propose an architecture against the FPKI policy requirements.

3.1. Orthogonal Policy Analysis

Each policy limitation identified from the beginning of the Cloud Tiger Team was subjected to a decision tree analysis which included the following questions:

- Will the removal of the policy requirement facilitate cloud adoption?
- Can [or has] the requirement be reasonably implemented in a cloud environment?
- Does complete removal of the policy requirement introduce a known risk vector?
- Is the requirement relevant to all system software layers in a virtualized instance such as a cloud hosted VM?
- Can the requirement be rewritten to support cloud implementations without impacting component confidentiality, integrity or availability?

Based on this analysis, the Cloud Tiger Team identified the following potential outcomes for each requirement under consideration:

- If removing the requirement presented no threat to the overall security of the community, it would be removed.
- If the Tiger Team could identify a modification of the requirement to support cloud implementation without compromising the security of the existing PKI implementations, a change proposal would be drafted for submission to the FPKI community.
- For all other requirements, the existing language would be maintained, and the Tiger Team would develop guidance to support further analysis of potential cloud implementations.

Although the Cloud Tiger Team did consider potential modification for four of the identified policy challenges, ultimately no recommendations were made regarding policy changes in order to prevent the introduction of new potential risks. As a result, the following section on Cloud Equivalent Controls Guidance was developed based on all of the policy statements that were subjected to the orthogonal policy analysis.

3.2. Cloud Equivalent Controls Guidance

The purpose of the cloud equivalent controls is to provide a framework for developing alternative PKI architectures, such as a cloud hosted PKI. Additionally, it documents the control objectives that must be considered before requesting a modification of established FPKI policy to accommodate a new deployment model.

Each table in the Appendix A addresses one of the Tiger Team identified policies limiting cloud adoption. Each table provides:

- the exact text from policy,
- definition of the control objective underlying the policy requirement,
- documentation of implicit assumptions present in the policy language, based on the current, on-prem architecture
- a list of threats that the policy is meant to mitigate, and
- guidance for implementers identifying the policy concerns that any alternative architecture, such as a cloud implementation, would need to address in order to propose a change to the existing policy requirements.

Prospective implementers of FPKI components in the cloud should use this guidance when designing and documenting security characteristics of their implementations to demonstrate compliance with the intent underlying FPKI policy requirements.

Appendix A - Cloud Equivalent Controls Documentation

The following sections have been extracted from [COMMON] for the purpose of conveying policy barriers to cloud adoption. The [FBCA CP] has equivalent requirements though the section references may differ slightly between the two policies.

Control Objective Table

As an aid to the reader, this document provides the following table which lists the control objectives for the cloud equivalent controls and provides a link to the section in this appendix that addresses the identified objective.

Control Objective	Relevant Section
Ensure that audit logs are not modified prior to review, and that the logs appropriately reflect the operations of the CA or component in question for a given period of time	5.4.4
Ensure that auditors provide unbiased feedback on CA operations	8.3
Ensure that audits verify the proper operation of all relevant elements of the PKI infrastructure	8.4
Ensure that the responsible management entity has effective control over the scope and methodology of the PKI audit of the PKI infrastructure	8
Ensures no one party can act unilaterally for key generation or unauthorized key usage	5.2.4
Ensures no one party can use a private key backup or restore a key from backup in order to perform unauthorized key usage, such as certificate or CRL signature.	6.2.4
Prevent destruction of key material through HSM destruction or zeroization maliciously or accidentally.	5.1.2.1
Prevent misuse of CA private key by a single malicious insider threat.	6.2.2
Prevent side-channel attacks, or interdiction/seeding of hardware intended to perform sensitive cryptographic functions.	6.6.1
Prevent unauthorized access to PKI components	6.7
Prevent unauthorized physical access to private key storage or private key materials which can result in the compromise of the CA signing function	5.1.2.1
Prevent unauthorized use of keys even in the event the HSM is physically compromised	6.4.2
Prevents unauthorized and unaudited access of relevant components	6.5.1
Prevents unauthorized execution of functions not assigned to specific roles (enforces permissions)	6.5.1
Prevents unintended discovery of PKI components, and reduces cyber attack surface area	6.7
Trusted Role background checks reduce the risk of untrustworthy persons having privileged access to sensitive resources.	5.3.2

5.1.2.1 Physical Access for CA Equipment

Current Policy Language
At a minimum, the physical access controls for CA equipment, as well as remote workstations used to administer the CAs, must: ● Ensure that no unauthorized access to the hardware is permitted. ● Ensure that all removable media and paper containing sensitive plain-text information is stored in secure containers. ● Be manually or electronically monitored for unauthorized intrusion at all times. ● Ensure an access log is maintained and inspected periodically. ● Require two-person physical access control to both the cryptographic module and computer systems.
Identified Control Objectives
● Prevent unauthorized physical access to private key storage or private key materials which can result in the compromise of the CA signing function ● Prevent destruction of key material through HSM destruction or zeroization maliciously or accidentally.
Current Policy Perspective
This policy assumes that all individuals physically interacting with the CA have full access, could gain full access, or could cause a denial of service. ● The generic access restrictions limits access by unauthorized individuals ● The two person control requirement limits the ability of any single authorized individual to abuse their logical access rights and is meant to limit the capabilities of potential insider threats.
Threats
● An unauthorized individual or a single authorized administrator leverages physical access to access data such as private key material through unauthorized channels, for example through the escalation of their privileges. ● Physical destruction of Hardware Storage Modules with CA Key Material. ● Destruction of Key Material by activation of tamper protection mechanisms.
Guidance for Equivalent Controls Implementation
● Demonstrate that the risk of access to CA functions and CA key material is reduced to an acceptable level ○ Access to the physical hardware on which virtualized functions are running does not grant logical access to virtual systems ○ The risk of destruction of the CA, or of HSMs housing key material is mitigated through appropriate physical or logical controls.

5.2.4 Number of Persons Required per Task

Current Policy Language
Where multiparty control is required, at least one of the participants must be an Administrator. All participants must serve in a trusted role as defined in Section 5.2.1.
Identified Control Objectives
● Ensures no one party can act unilaterally for key generation or unauthorized key usage
Current Policy Perspective
This policy assumes that all individuals interacting with the CA have full administrative access, or could gain full access. ● The generic access restrictions limits access by unauthorized individuals ● The two person control requirement limits the ability of any single authorized individual to abuse their logical access rights and is meant to limit the capabilities of potential insider threats.
Threats
● An unauthorized individual or a single authorized administrator leverages logical access to access data through unauthorized channels, for example through the escalation of their privileges. ● Key compromise, denial of service, or issuance of fraudulent certificates for malicious purposes
Guidance for Equivalent Controls Implementation
● Demonstrate that, for a subset of administrative functions, the risk of privilege escalation is reduced to an acceptable level. ○ This could be achieved via internal segregation of system functions or other mechanisms that block privilege escalation.

5.3.2 Background Check Procedures

Current Policy Language
Trusted Roles must receive a favorable adjudication after undergoing a background investigation covering the following areas: ● Employment; ● Education; ● Place of residence; ● Law Enforcement; and ● References. The period of investigation must cover at least the last five years for each area, excepting the residence check which must cover at least the last three years. Regardless of the date of award, the highest educational degree must be verified. Adjudication of the background investigation must be performed by a competent adjudication authority using a process consistent with [Executive Order 12968], or equivalent.
Identified Control Objectives
● Trusted Role background checks reduce the risk of untrustworthy persons having privileged access to sensitive resources.
Current Policy Perspective
This policy assumes that all individuals interacting with the CA have full administrative access, or could gain full access. ● The generic access restrictions limits access by unauthorized individuals ● The two person control requirement limits the ability of any single authorized individual to abuse their logical access rights. ● This control is meant to limit the probability of a potential insider threat being appointed to a trusted role, it is based off of established personnel security practices and assumes that past behavior is the only indicator for future behavior
Threats
● An unauthorized individual or a single authorized administrator leverages logical access to access data through unauthorized channels, for example through the escalation of their privileges. ● Unauthorized actions from insider threat to include loss of confidentiality, integrity, or availability of critical data/PKI services, or the misissuance of certificates
Guidance for Equivalent Controls Implementation
● Demonstrate that, for a subset of administrative functions, the risk of privilege escalation is reduced to an acceptable level. ○ This could be achieved via internal segregation of system functions or other mechanisms that block privilege escalation. ● For the identified subset of acceptable risk functions, identify how individuals with access to those functions are vetted for trustworthiness. ○ This process need not conform to the full set of Trusted Role background check requirements.

5.4.4 Protection of Audit Log

Current Policy Language
System configuration and operational procedures must be implemented together to ensure that only authorized individuals may move or archive audit records and that audit records are not modified before review. Collection of the audit records from the CA system must be performed by, witnessed by or under the control of trusted roles who are different from the individuals who, in combination, command the CA signature key.
Identified Control Objectives
● Ensure that audit logs are not modified prior to review, and that the logs appropriately reflect the operations of the CA or component in question for a given period of time
Current Policy Perspective
● Under current deployments, in on-premise configurations, audit logs for all relevant components, and all layers of the system are owned or exclusively controlled by the entity operating the CA. ● In cloud configurations, under a shared responsibility model, the entity operating the CA infrastructure may not have visibility to audit logs or events occurring on systems that are under the responsibility of the service provider.
Threats
● Attacks against the CA infrastructure are undetected due to gaps in the audit records
Guidance for Equivalent Controls Implementation
● For effective monitoring of the CA infrastructure, monitoring must include all relevant infrastructure components. ● Document which elements of shared [cloud] infrastructure are possible threat vectors for CA specific assets, and demonstrate that CA assets are truly segregated from the excluded elements ● Document and demonstrate that logs from shared infrastructure elements are protected from tampering before being reviewed. ● Document and demonstrate that controls are in place to prevent individuals with control over the signing key from managing the process of audit log collection and transfer ● Document and demonstrate that the individuals who have permissions that would allow them to modify or delete audit records are sufficiently vetted for trustworthiness.

6.2.2 Private Key (n out of m) Multi-Person Control

Current Policy Language
A single person must not be permitted to activate or access any cryptographic module that contains the complete CA private signing key. CA signature keys may be backed up only under two-person control. Access to CA signing keys backed up for disaster recovery must be under at least two-person control. The names of the parties used for two-person control must be maintained on a list that must be made available for inspection during compliance audits.
Identified Control Objectives
● Prevent misuse of CA private key by a single malicious insider threat.
Current Policy Perspective
● In on-premise configurations, the CA cryptographic modules are dedicated to the organization managing the CA, and the initialization and configuration of the module is audited to ensure that multi-party control is in effect for the entire time in which CA key material is present in the module. ● For current providers, the list of individuals who may interact with CA cryptographic modules is small, and a list of them is maintained at all times. ● Auditors are able to verify that no unauthorized personnel have accessed the CA cryptographic modules
Threats
● A single malicious insider can access the private key in order to make unauthorized copies or otherwise make unauthorized use of the key material (e.g., fraudulent certificate signature).
Guidance for Equivalent Controls Implementation
● Document and demonstrate that the cryptographic module hosting a complete private CA key is under multi-party control, and that a complete list of all individuals who may physically or logically access the module and activate the private key at any point in time can be produced. ○ Note that different parties may be responsible for physical vs logical access. For example, a hosting provider may have physical access, and a customer may have logical access. ● Document private key activation factors and any physical controls or processes that are in place to enforce multi party control (see 6.4.2) ● Document and demonstrate that all backups containing complete copies of the key are under the same or greater levels of control.

6.2.4 Private Key Backup

Current Policy Language
All backups of the CA, CSS and PIV Content Signing private signature keys must be accounted for and protected under the same multi-person control as the original signature key. At least one copy of the CA private signature key must be stored off-site. For all other keys, backup, when permitted, must provide security controls consistent with the protection provided by the original cryptographic module. Backed up private signature key(s) must not be exported or stored in plaintext form outside the cryptographic module.
Identified Control Objectives
● Ensures no one party can use a private key backup or restore a key from backup in order to perform unauthorized key usage, such as certificate or CRL signature.
Current Policy Perspective
● In the absence of technical information about the protection mechanisms implemented by key storage solutions for private key backup, we must rely on the same level of protection applied to active key material in an HSM. Part of this involves multi-party control for backups of key material. ● FIPS 140-3 level 2 and 3 validated HSMs provide the highest commercially available level of protection for private key material, and for this reason they are the foundation of technical trust within a CA. Once private key material leaves the protected boundary of the cryptographic module, however, the opportunities for attack increase significantly.
Threats
● Impersonation of a CA, CSS or PIV Content Signer via use of key material obtained from a backup to sign objects (e.g., fraudulently signed PKI objects such as subscriber certificates).
Guidance for Equivalent Controls Implementation
● If muti-party control of all backups of CA private key material cannot be assured throughout the entire lifecycle of the backups and in any location they might be kept, the implementer must describe how unauthorized restoration of the private key is prevented. ● Describe and demonstrate that all key material backed up has controls of equal or greater strength than those enforced by the HSM itself. ● Demonstrate that restoration of key material to the original HSM or a different HSM may only be done under controls as stringent as those required for key generation. ● Consider all keys and all scenarios where a key may exist in an unprotected state. The following list provides examples but is not exhaustive: ○ Keys in application memory that may be found in core dumps or through other mechanisms ○ Keys used to cryptographically protect signing keys while on disk or in backups

6.4.2 Activation Data Protection

Current Policy Language
Data used to unlock private keys must be protected from disclosure by a combination of cryptographic and physical access control mechanisms. Activation data must be: ● recorded and secured at the level of assurance associated with the activation of the cryptographic module, and must not be stored with the cryptographic module.
Identified Control Objectives
● Prevent unauthorized use of keys even in the event the HSM is physically compromised
Current Policy Perspective
● Private key controls rely in part on protection of the credentials that administrators use to authenticate to the HSM for activation. In a private data center context, this can be accomplished by ensuring that the activation material is stored at the same security level, but in a different location. ○ “Credentials” in this context refers to all elements related to authentication, including any physical objects such as authentication tokens, or logical elements such as PINs used to unlock tokens. ● In different contexts, if that assurance cannot be provided, then we must ensure that private keys are protected.
Threats
● An attacker uses private key activation data to perform unauthorized/fraudulent signing with the private key.
Guidance for Equivalent Controls Implementation
● Document and demonstrate how activation material is protected from unauthorized access or use, including any physical security controls or processes in place such as protection of administrator card sets ● Document and demonstrate how internal controls minimize unauthorized access to or use of activation data.

6.5.1 Specific Computer Security Technical Requirements

Current Policy Language
For CAs, KEDs, and DDSs operating under this policy, the computer security functions listed below are required. These functions may be provided by the operating system, or through a combination of operating system, software, and physical safeguards. The CA and its ancillary parts must include the following functionality (these functions are applicable to all system software layers where applicable): ● authenticate the identity of users before permitting access to the system, data, or ● applications; ● manage privileges of users to limit users to their assigned roles; ● enforce domain integrity boundaries for security critical processes; ● require a trusted path for identification of all users; Methods used to administer the CAs, KEDs, or DDSs must not bypass applicable two-person controls. In addition, the computer security functions listed below are required: ● authenticate the identity of users before permitting access to the system or applications; ● manage privileges of users to limit users to their assigned roles; ● enforce domain integrity boundaries for security critical processes; ● prohibit object reuse or require separation for random access memory;
Identified Control Objectives
● Prevents unauthorized and unaudited access of relevant components ● Prevents unauthorized execution of functions not assigned to specific roles (enforces permissions)
Current Policy Perspective
● In a traditional, data center and physical server centric deployment of a CA, the operating system must enforce separation between processes and between users to protect the software performing the signing function from being leveraged by other processes. Virtualization and cloud technologies introduce much more complexity, since it introduces additional targets for attack such as the hypervisor or the infrastructure itself. ● Traditional CA, KED and DDSs deployments rely on some well known operating system functions to ensure that prevent malicious processes from capturing user credentials (e.g. Trusted Path). In a remote environment, such as a cloud, the risk of user credential capture and replay must be addressed in a different manner.
Threats
● Attacks on the CA signing function from other processes that can access shared resources such as CPU, memory or block storage. ● Use of management functions to achieve privileged access to managed resources, this can include privileged virtual consoles, hypervisor functions or other means of accessing or managing virtualized or cloud resources
Guidance for Equivalent Controls Implementation

- Describe and demonstrate that processes performing signing functions on behalf of a CA, KED, or DDS are protected from all processes running in the environment, this can include other processes on the system, or may include processes running on a hypervisor, or a cloud infrastructure. For each process that could potentially access shared resources used by the signing function, demonstrate how that process is restricted from performing improper signing or accessing sensitive data such as keys or credentials.
- Include all parties involved in the provisioning and managing of services, including third party resellers of cloud services, or individuals supporting underlying functions such as data center personnel or networking administrators.
- Describe authenticators used to access the PKI components as a trusted role, especially if not PIV/CAC enabled

6.6.1 System Development Controls

Current Policy Language
Hardware and software used to administer or operate the CA must be procured in a fashion to reduce the likelihood that any particular component was tampered with (e.g., by ensuring the vendor cannot identify the PKI component that will be installed on a particular device). The CA hardware and software, including all system software layers, must be dedicated to operating and supporting the CA (i.e., the systems and services dedicated to the issuance and management of certificates). There must be no other applications, hardware devices, network connections, or component software installed that are not parts of the CA operation, administration, monitoring and security compliance of the system. CA hardware and system software layers may support multiple CAs and their supporting systems, provided all systems have comparable security controls and are dedicated to the support of the CA in compliance with this CP
Identified Control Objectives
● Prevent side-channel attacks, or interdiction/seeding of hardware intended to perform sensitive cryptographic functions.
Current Policy Perspective
● Current policy assumes that access to physical CA resources at any time can lead to unauthorized access to the CA and its functions. Even before the machine is delivered and built, additional hardware or firmware components may be installed which may be used by an attacker to access sensitive data or perform unauthorized functions. ● Virtualized or cloud environments add additional complexity by introducing “virtualized” hardware. A virtual machine appears to a customer as a single physical server, but is in fact composed of software interfaces managed by a hypervisor or other layer.
Threats
● An individual with privileged access to hardware or virtual hardware, and knowledge that the hardware will be used for sensitive transactions, may make modifications to the device which will enable them to perform unauthorized actions resulting in access to or use of the private signing key. ● Malicious software installed on shared resources may be able to exfiltrate data
Guidance for Equivalent Controls Implementation
● Describe and demonstrate how physical hardware is protected during manufacturing and transport, and how the ultimate customer and use case for the hardware is hidden prior to the hardware being placed in service. ● Describe and demonstrate how elements of a virtual machine, including software and configuration, are protected from tampering by unauthorized personnel, or by authorized personnel attempting unauthorized actions. ● Describe and demonstrate appropriate intrusion detection and prevention solutions that have been put in place to limit impacts of malware such as data exfiltration

6.7 Network Security Controls

Current Policy Language
Protection of CA and KRS equipment must be provided against known network attacks. All unused network ports and services must be turned off. Any network software present on the CA and KRS equipment must be necessary to the functioning of the CA application. Any boundary control devices used to protect the network on which PKI equipment is hosted must deny all but the necessary services to the PKI equipment. Repositories, CSSs, KRA/KRO, and remote workstations used to administer the CAs must employ appropriate network security controls. Networking equipment must turn off unused network ports and services. Any network software present must be necessary to the function of the equipment.
Identified Control Objectives
● Prevents unintended discovery of PKI components, and reduces cyber attack surface area ● Prevent unauthorized access to PKI components
Current Policy Perspective
● In a traditional, data center and physical server centric deployment of a CA, the determination of unnecessary services is solely up to the CA operator. Virtualization and cloud technologies introduce much more complexity, since other stakeholders may be operating services on the same infrastructure components. ● In multi-tenant shared infrastructure, remote workstations will access infrastructure components for CA and non-CA systems, at both high and low risk levels.
Threats
● Attacks on the CA signing function from other systems that can access shared resources such as CPU, memory or block storage, or shared software such as hypervisors. ● Use of management functions to achieve unauthorized privileged access to managed resources, this can include privileged virtual consoles, hypervisor functions or other means of accessing or managing virtualized or cloud resources
Guidance for Equivalent Controls Implementation
● Describe and demonstrate that processes performing signing functions on behalf of a CA, KED, or DDS are protected from all processes running in the environment, this can include other processes on the system, or may include processes running on a hypervisor, or a cloud infrastructure. For each process that could potentially access shared resources used by the signing function, demonstrate how that process is restricted from performing improper signing or accessing sensitive data such as keys or credentials. ● Describe and demonstrate that CA resources can be isolated from other resources not dedicated to CA operations ● Include all parties involved in the provisioning and managing of services, including third party resellers of cloud services, or individuals supporting underlying functions such as data center personnel or networking administrators.

- Describe authenticators used to access the PKI components as a trusted role, especially if not PIV/CAC enabled
- Additionally refer to CA/B forum guidance for network security controls in order to address industry best practices: https://cabforum.org/working-groups/netsec/documents/

8 Compliance Audit and Other Assessments

Current Policy Language
The organization's PMA must be responsible for ensuring annual audits are conducted for all PKI functions regardless of how or by whom the PKI components are managed and operated
Identified Control Objectives
• Ensure that the responsible management entity has effective control over the scope and methodology of the PKI audit of the PKI infrastructure
Current Policy Perspective
• For transitive trust to function across the community, each participant must ensure that the proper functions are conducted and audited, regardless of whether the participant performs the function themselves. The PMA is the business entity accepting overall responsibility on behalf of the participant, so it is essential that they be ultimately responsible for ensuring that audits are carried out. • In scenarios where infrastructure is shared, the PMA may rely on common audits if they address the entire scope, but must perform their own audits if necessary to address any gaps in scope.
Threats
• An attacker takes advantage of a lapse in control effectiveness that was not detected due to a lack of verification of control effectiveness.
Guidance for Equivalent Controls Implementation
• Any participant who wishes to leverage a shared responsibility model to address audit requirements, must identify the specific parties performing the audit of each of the items within the required audit scope, and must provide that information, along with evidence of the completed audits, as part of their usual annual review package.

8.3 Assessor’s Relationship To Assessed Entity

Current Policy Language
The compliance auditor either must be a private firm that is independent from the entities (CA and RAs) being audited, or it must be sufficiently organizationally separated from those entities to provide an unbiased, independent evaluation.
Identified Control Objectives
● Ensure that auditors provide unbiased feedback on CA operations
Current Policy Perspective
● Under current deployments, the PMA may send their own audit team to validate controls under a shared infrastructure. This ensures that gaps in the service providers audit coverage may be addressed in the overall audit. ● For some cloud service providers, third party audits by customers are not permitted, meaning that the independent auditor will not be able to verify the effectiveness of some of the controls in scope of the audit.
Threats
● An attacker takes advantage of a lapse in control effectiveness that was not detected due to a lack of verification of control effectiveness.
Guidance for Equivalent Controls Implementation
● Any participant who wishes to leverage a shared responsibility model to address audit requirements, must identify the audit scope for all audits performed by service providers, and must be able to address any gaps in audit coverage. The participant must provide that information, along with evidence of the completed audits, as part of their usual annual review package. All audits must be performed by auditors who meet the independence requirement of the appropriate policy.

8.4 Topics Covered By Assessment

Current Policy Language
All aspects of the CA/RA operation must be subject to compliance audit inspections.
Identified Control Objectives
• Ensure that audits verify the proper operation of all relevant elements of the PKI infrastructure
Current Policy Perspective
• Under current deployments, the PMA may send their own audit team to validate controls under a shared infrastructure. This ensures that gaps in the service providers audit coverage may be addressed in the overall audit. • For some cloud service providers, third party audits by customers are not permitted, meaning that the independent auditor will not be able to verify the effectiveness of some of the controls in scope of the audit.
Threats
• An attacker takes advantage of a lapse in control effectiveness that was not detected due to a lack of verification of control effectiveness.
Guidance for Equivalent Controls Implementation
• Any participant who wishes to leverage a shared responsibility model to address audit requirements, must identify the audit scope for all audits performed by service providers, and must be able to address any gaps in audit coverage. The participant must provide that information, along with evidence of the completed audits, as part of their usual annual review package.   ○ FedRAMP certified systems must review the Security Controls Overlay of NIST Special Publication 800-53 Revision 5 Security Controls for FPKI Systems for guidance on additional controls beyond those specified in 800-53.