

NERC

NORTH AMERICAN ELECTRIC
RELIABILITY CORPORATION

Cyber Security — System Security Management

Technical Rationale and Justification for
Reliability Standard CIP-007-7

April 2024

RELIABILITY | RESILIENCE | SECURITY



**3353 Peachtree Road NE
Suite 600, North Tower
Atlanta, GA 30326
404-446-2560 | www.nerc.com**

Table of Contents

Technical Rationale for Reliability Standard CIP-007-7.....	3
Introduction.....	3
New and Modified Terms and Applicability	3
General Considerations	4
Requirement R1.....	4
Requirement R2.....	6
Requirement R3.....	7
Requirement R4.....	7
Requirement R5.....	8
Former Background Section from Reliability Standard CIP-007-6	9
Background.....	9
Technical Rationale for Reliability Standard CIP-007-6.....	11
Guidelines and Technical Basis.....	11

Technical Rationale for Reliability Standard CIP-007-7

Introduction

This document explains the technical rationale and justification for the proposed Reliability Standard CIP-007-7. It provides stakeholders and the ERO Enterprise with an understanding of the technology and technical requirements in the Reliability Standard. This Technical Rationale and Justification for CIP-007-7 is not a Reliability Standard and should not be considered mandatory and enforceable.

Updates to this document now include the Project 2016-02 – Modifications to CIP Standards Drafting Team’s (SDT’s) intent in drafting changes to the requirements.

Background

The V5TAG, which consists of representatives from NERC, Regional Entities, and industry stakeholders, was formed to issue guidance regarding possible methods to achieve compliance with the CIP V5 standards and to support industry’s implementation activities. During the course of the V5TAG’s activities, the V5TAG identified certain issues with the CIP Reliability Standards that were more appropriately addressed by a standard drafting team (SDT). The V5TAG developed the V5TAG Transfer Document to explain the issues and recommend that they be considered in future development activity. As Project 2016-02 was formed to address the directives in FERC Order 822 issued on January 21, 2016, that team also received addressing the V5TAG issues as part of its Standard Authorization Request (SAR).

One of the areas of issue was virtualization. The V5TAG Transfer document said, “The CIP Version 5 standards do not specifically address virtualization. However, because of the increasing use of virtualization in industrial control system environments, questions around treatment of virtualization within the CIP Standards are due for consideration. The SDT should consider revisions to CIP-005 and the definitions of Cyber Asset and Electronic Access Point that make clear the permitted architecture and address the security risks of network, server, and storage virtualization technologies.”

New and Modified Terms and Applicability

This standard uses new or modified terms and contains new or modified exemptions in Section 4 Applicability. The rationale for this global content can be found in “CIP Definitions and Exemptions Technical Rationale.pdf” for reference when reading the technical rationale that follows.

General Considerations

Requirement R1

General Considerations for Requirement R1

Shared Cyber Infrastructure (SCI) is mutually exclusive from BES Cyber System (BCS) by definition. To clarify treatment of virtualization concepts in CIP-007-7, the SDT added “SCI supporting an Applicable System in this Part”. This approach keeps the SCI applicability parallel to each existing variant of Applicable Systems in the same Requirement Part (i.e., Medium impact BCS vs. Medium impact BCS with External Routable Connectivity (ERC) vs. Medium impact BCS at Control Centers etc.)

Change Rationale Requirement R1:

The SDT has changed the name of the R1 table to “System Hardening” to clearly reflect the security objective of the entries in the table, which is to reduce a systems’ attack surface. The parts of this requirement do this by:

- limiting network access to only needed routable protocol network accessibility (i.e., logical ports or services),
- limiting access to physical I/O ports (i.e., USB, network, console ports) to only what is necessary, and
- limiting the shared hardware (CPU and memory resources) attack surface between certain groups of VCAs on SCI.

As R1 is broader than only ports and services, the table name was changed to reflect this. The SDT notes that this is merely a label on the Requirement Parts table, linking it to the main Requirement 1 language above it, and does not change or imply any other Requirements or Requirement Parts.

The SDT chose to include “SCI supporting an Applicable System in this Part” language in the Applicable Systems column of the requirement to ensure that controls which are applicable to the hardware portion of a physical Cyber Asset would remain applicable to the hardware supporting the Virtual Cyber Assets (VCA) used in applicable BCS, EACMS, PACS or PCA.

Change Rationale Requirement R1 Part 1.1: Requirement R1 Part 1.1 requires “disable or prevent unneeded routable protocol network accessibility on each Applicable System, per system capability”. The SDT updated the Requirement Part language to state a security objective concerning “*routable protocol network accessibility*” as opposed to “ports and services”. As this is a new phrase, the intent of this phrase with some examples and rationale for this change is as follows.

Routable protocol network accessibility - The objective of this phrase in the Requirement Part is to reduce the attack surface on an applicable system by preventing any unnecessary accessibility to the system over a network using routable protocols. “Accessibility” as used in R1.1 is at the routable protocol network level and does not include physical access, logon to the physical console, code on TCA/RM, etc. Port numbers (TCP/UDP) are at times the best way to track this accessibility, at other times documenting enabled services is better. For example, reducing network accessibility in the physical underlay of SCI between hypervisors or on fabric-based networks may be best performed at a *services* level; turning off or disabling virtualization services that are not needed, rather than documenting the often proprietary and dynamic port numbers which may be of little value. However, in the overlay where an entity may be hosting a database server VCA, it may be easier to show that network accessibility on that VCA is limited to SQL server and remote admin enabled port numbers. In Zero Trust Architectures (ZTA), it may be neither ports or services, but instead a “user to tagged workload” level access control policy where accessibility is described and protected at a more granular, yet higher level than a port #, enforced at a Policy Enforcement Point (PEP) on the applicable system. The SDT has moved to this objective language to avoid prescribing only one way to perform and document network accessibility in all these various scenarios and implementations. In addition, it is limited to routable protocol network accessibility such that non routable network communications (e.g., SAN over Fiber Channel) do not fall within scope of the Requirement Part. The objective is to know the ways a system

can be accessed from the network via routable protocols and have no unnecessary attack surface from that perspective.

In this Requirement Part, the SDT used the verbs “Disable or prevent”. In some cases, the entity may be able to disable a service or remove/uninstall software that is providing unneeded network accessibility to the applicable system. In other cases, the entity may not be able to disable a service, but can prevent access to it in another layer, such as the underlying OS with a host-based firewall, a PEP, or other means of filtering traffic. In instances where the entity can do neither (e.g., a firmware-based ‘black box’ device with limited configuration capabilities), the SDT chose to add ‘per system capability’ to make the requirement conditional on the ability of the applicable system to meet it, if the entity can show that it is incapable.

The SDT also added the clarifier “on each Applicable System” to indicate that the intent of this requirement is for an entity to perform the configuration actions *on* each Applicable System, hardening the system from its routable protocol network peers rather than having a single method such as an EAP network firewall rule that would disable such accessibility for a group of Applicable Systems on a network. In other words, merely filtering a port/service on a firewall at an ESP network boundary (CIP-005 R1 controls) does not meet the intent of CIP-007 R1.

Change Rationale Requirement R1 Part 1.3:

Requirement R1 Part 1.3 is a new requirement intended to apply security controls in Shared Cyber Infrastructure scenarios to mitigate the risk of hardware-based vulnerabilities between groups of Applicable Systems of differing impact categorizations. This requirement addresses these vulnerabilities at the CPU and memory level (e.g., Downfall, Inception, Spectre, Meltdown, Rowhammer) that could allow for compromise across processes simultaneously executing on the same shared hardware. Virtualization technology typically has “affinity” controls that can be used to tag VCAs by their impact category and then configure affinity rules such that hypervisors do not instantiate or move VCAs of differing trust levels to the same CPU or memory resources at the same time. Note the requirement excludes shared storage resources, although included in the definition of SCI. The SDT intent is it is allowable to have logical disks attached to systems of differing impact categorization stored in the same SAN for example, and since these storage solutions have internal CPU and memory, which is not the intended object of this requirement.

Note that the SCI definition covers *only* those portions of entity’s virtualization infrastructure that support systems of differing impact levels. As VCAs that share the same CPU or memory with any part of the highest rated BCS are “Associated PCAs”, this requirement is only needed to cover instances where different impact categorizations are supported and the entity wishes to keep other VCAs from becoming “associated PCAs”, thus maintaining the separation by impact category. This may occur when part of the same SCI supports different BCS and the entity does not wish to high watermark all the BCS together, but instead employ affinity rules to prevent VCAs that are high or medium impact from being instantiated or moved to CPU and memory resources where other VCAs are executing.

For example, this requirement would apply where the same SCI supports both a high impact BCS on one part of a virtualization cluster and a low impact BCS on a different part of the virtualization cluster. The proper use of host affinity controls would not require the low impact BCS to be considered an “associated PCA” of the high as long as:

1. the low impact BCS is not in the ESP of the high impact BCS, and
2. controls are in place to prevent the SCI from executing these BCS on the same CPU and memory resources.

As an example, if an entity has a substation intelligent electronic device (IED) management system that is used as an EACMS for medium and low impact BCS in a single VCA, it is an associated EACMS of the medium impact BCS and protected as such. It can therefore execute on hypervisors with other medium or high impact BCS. If it is implemented as two separate VCAs, one for medium impact IEDs and another for low impact IEDs, then the low impact one cannot execute on the same hypervisor as the medium impact one, unless it is categorized as an “associated PCA” of the medium BCS as well. If the entity wishes to have it remain as low impact those two VCAs sharing the same CPU or

memory resources would not meet this requirement. In situations where high or medium impact VCAs are allowed to execute on the same CPU or memory resources with VCA's of other impact levels and those other VCA's are inadvertent PCAs, this requirement is a single place for that violation rather than every requirement that has PCA in the scope.

Requirement R2

Change Rationale Requirement R2:

The SDT chose to include the "SCI supporting an Applicable System in this Part" language in the Applicable Systems column of the requirement in order to ensure that controls which are applicable to the hardware portion of a physical Cyber Asset would remain applicable to the hardware supporting the VCA(s) used in applicable BCS, EACMS, PACS or PCA.

The SDT made conforming changes where necessary with Requirement R2 to remove reliance on the Cyber Asset term, choosing to reference the Applicable Systems instead.

Additionally, the SDT chose to insert the word "cyber" as a clarifier to security patches for consistency with the term used in R2.1.

Backward Compatibility

CIP-007-7 Requirement R2 retains backward compatibility for entities that do not utilize SCI.

The SDT intends that entities take full advantage of their virtualization infrastructure in order to ease the overhead associated with patch management. While many of the entity's existing processes will remain the same, (such as those associated with tracking the source of cyber security patches, evaluation of the applicability of available patches and mitigation plans for those applicable patches that cannot be installed) new or modified processes around the installation of patches can be used (for example, parent images, remediation VLANs, etc.).

Parent Images

One of the interesting nuances of virtualization is the concept of parent/child relationships.

Some VCAs utilize a "parent image" type methodology where a specific VCA (operating system and software) is merely a child instance of a parent image. In such cases, the application of cyber security patches to the parent image applies those changes to any children instantiated from that image. As there may be many dynamic child instances, the administrative overhead associated with patching of those child instances can be greatly reduced. The intent is that entities would track and document the R2 cyber security patching requirements at the parent image level rather than attempting to track and document patching for each individual VCA instantiated from the parent image.

Dormant Instances

There are times when a VCA may not be in use and, as a virtual instance, may be shut down to free its resources since it can be easily instantiated when needed. Thus, VCAs are unlike physical Cyber Assets that are up and running continuously and can receive applicable patches, software updates, configuration updates, etc.

Leveraging the built-in virtualization features allows idle resources to be reassigned to tasks at hand without incurring additional overhead of tracking which dormant virtual instances require patching. Dormant virtual instances are just files with the saved state of the VCA. Dormant instances are not VCAs themselves nor Applicable Systems until they become active instances again.

Where a dormant virtual instance is also the child instance of a parent image, the application of security patches to the parent image will also make the same changes automatically when needed to any out of date dormant child virtual instances when they are restarted.

In most cases, a dormant virtual child instance is made active again on a remediation type network where any missing security patches (compared to the parent image) are automatically applied before it is placed back into active service.

Saved Images

Similar to dormant child images, saved child images or snapshots may be also used as functional backups for CIP-009 purposes in order to provide fast restoration and better availability. If an active VCA becomes corrupted, a saved image can be made active where the application of any missing security patches (compared to the parent image) can be automatically applied, through remediation, before being placed back into active service. However, the SDT's intent is that VCA images that are used for backup purposes would not be subject to patching, as the backup may be needed to recover from a failed patch or if later discovered that a patch has negatively affected functionality.

Requirement R3

Change Rationale Requirement R3:

The SDT chose to include "SCI supporting an Applicable System in this Part" language in the Applicable Systems column of the requirement in order to ensure that controls which are applicable to the hardware portion of a physical Cyber Asset would remain applicable to the hardware supporting the VCA used in applicable BCS, EACMS, PACS or PCA.

Requirement R4

Change Rationale Requirement R4:

The SDT chose to include "SCI supporting an Applicable System in this Part" language in the Applicable Systems column of the requirement in order to ensure that controls which are applicable to the hardware portion of a physical Cyber Asset would remain applicable to the hardware supporting the VCA used in applicable BCS, EACMS, PACS or PCA.

Change Rationale Requirement R4 Part 4.1:

The SDT determined the entire structure of the "at the BCS level (per BES Cyber System capability) or at the Cyber Asset level (per Cyber Asset capability)" language could be simplified to "per system capability" since the Applicable Systems column clarifies what systems are included.

Additionally, the SDT chose to insert the word "security" as a clarifier to which events are to be logged.

Change Rationale Requirement R4 Part 4.2:

The SDT applied conforming changes to the Requirement Part language to remove the reliance on the "Cyber Asset" term, as well as the ambiguity that would have existed if the term "BES Cyber System" were left in place while the Applicability Column included SCI.

Change Rationale Requirement R4 Part 4.3:

The SDT clarified the retaining of applicable "security" event logs, as that is what is identified in Part 4.1.

The SDT also chose to remove the reliance on a TFE from Part 4.3 in favor of the updated term "per system capability". The SDT contends that the term still requires an entity to document the limit to the system's capability with regards to the requirement language, while not incurring the additional documentation overhead of a TFE.

Requirement R5

Change Rationale Requirement R5:

The SDT chose to include “SCI supporting an Applicable System in this Part” language in the Applicable Systems column of the requirement in order to ensure that controls which are applicable to the hardware portion of a physical Cyber Asset would remain applicable to the hardware supporting the VCA used in applicable BCS, EACMS, PACS or PCAs.

Change Rationale Requirement R5 Parts 5.1, and 5.6 – 5.7:

The SDT chose to remove the reliance on TFE in favor of the updated term “per system capability”. The SDT contends that the term still requires an entity to document the limit to the system’s capability with regards to the requirement language, while not incurring the additional documentation overhead of a TFE.

Change Rationale Requirement R5 Part 5.4:

The SDT applied conforming changes to the Requirement Part language to remove the reliance on the “Cyber Asset” term, replacing it with a per “system” capability.

Change Rationale Requirement R5 Part 5.5:

The SDT applied conforming changes to the Requirement Part language to remove the reliance on the “Cyber Asset” term, replacing it with a reference to the Applicable Systems.

Former Background Section from Reliability Standard CIP-007-6

The section **6. Background** has been retired and removed from the Standard, and preserved by cutting and pasting as-is below.

Background

Standard CIP-007 exists as part of a suite of CIP Standards related to cyber security, which require the initial identification and categorization of BES Cyber Systems and require a minimum level of organizational, operational and procedural controls to mitigate risk to BES Cyber Systems.

Most requirements open with, “*Each Responsible Entity shall implement one or more documented [processes, plan, etc.] that include the applicable items in [Table Reference].*” The referenced table requires the applicable items in the procedures for the requirement’s common subject matter.

The term *documented processes* refers to a set of required instructions specific to the Responsible Entity and to achieve a specific outcome. This term does not imply any particular naming or approval structure beyond what is stated in the requirements. An entity should include as much as it believes necessary in its documented processes, but it must address the applicable requirements in the table.

The terms *program* and *plan* are sometimes used in place of *documented processes* where it makes sense and is commonly understood. For example, documented processes describing a response are typically referred to as *plans* (i.e., incident response plans and recovery plans). Likewise, a security plan can describe an approach involving multiple procedures to address a broad subject matter.

Similarly, the term *program* may refer to the organization’s overall implementation of its policies, plans, and procedures involving a subject matter. Examples in the standards include the personnel risk assessment program and the personnel training program. The full implementation of the CIP Cyber Security Standards could also be referred to as a program. However, the terms *program* and *plan* do not imply any additional requirements beyond what is stated in the standards.

Responsible Entities can implement common controls that meet requirements for multiple high and medium impact BES Cyber Systems. For example, a single training program could meet the requirements for training personnel across multiple BES Cyber Systems.

Measures for the initial requirement are simply the documented processes themselves. Measures in the table rows provide examples of evidence to show documentation and implementation of applicable items in the documented processes. These measures serve to provide guidance to entities in acceptable records of compliance and should not be viewed as an all-inclusive list.

Throughout the standards, unless otherwise stated, bulleted items in the requirements and measures are items that are linked with an “or,” and numbered items are items that are linked with an “and.”

Many references in the Applicability section use a threshold of 300 MW for UFLS and UVLS. This particular threshold of 300 MW for UVLS and UFLS was provided in Version 1 of the CIP Cyber Security Standards. The threshold remains at 300 MW since it is specifically addressing UVLS and UFLS, which are last ditch efforts to save the Bulk Electric System. A review of UFLS tolerances defined within regional reliability standards for UFLS program requirements to date indicates that the historical value of 300 MW represents an adequate and reasonable threshold value for allowable UFLS operational tolerances.

“Applicable Systems” Columns in Tables:

Each table has an “Applicable Systems” column to further define the scope of systems to which a specific requirement row applies. The CSO706 SDT adapted this concept from the National Institute of Standards and Technology (“NIST”) Risk Management Framework as a way of applying requirements more appropriately based on impact and connectivity characteristics. The following conventions are used in the “Applicability Systems” column as described.

- **High Impact BES Cyber Systems** – Applies to BES Cyber Systems categorized as high impact according to the CIP-002-5.1 identification and categorization processes.
- **Medium Impact BES Cyber Systems** – Applies to BES Cyber Systems categorized as medium impact according to the CIP-002-5.1 identification and categorization processes.
- **Medium Impact BES Cyber Systems at Control Centers** – Only applies to medium impact BES Cyber Systems located at a Control Center.
- **Medium Impact BES Cyber Systems with External Routable Connectivity** – Only applies to medium impact BES Cyber Systems with External Routable Connectivity. This also excludes Cyber Assets in the BES Cyber System that cannot be directly accessed through External Routable Connectivity.
- **Electronic Access Control or Monitoring Systems (EACMS)** – Applies to each Electronic Access Control or Monitoring System associated with a referenced high impact BES Cyber System or medium impact BES Cyber System. Examples may include, but are not limited to, firewalls, authentication servers, and log monitoring and alerting systems.
- **Physical Access Control Systems (PACS)** – Applies to each Physical Access Control System associated with a referenced high impact BES Cyber System or medium impact BES Cyber System.
- **Protected Cyber Assets (PCA)** – Applies to each Protected Cyber Asset associated with a referenced high impact BES Cyber System or medium impact BES Cyber System.

Technical Rationale for Reliability Standard CIP-007-6

This section contains a “cut and paste” of the former Guidelines and Technical Basis (GTB) as-is of from CIP-007-6 standard to preserve any historical references. No modifications have been made.

Guidelines and Technical Basis

Section 4 – Scope of Applicability of the CIP Cyber Security Standards

Section “4. Applicability” of the standards provides important information for Responsible Entities to determine the scope of the applicability of the CIP Cyber Security Requirements.

Section “4.1. Functional Entities” is a list of NERC functional entities to which the standard applies. If the entity is registered as one or more of the functional entities listed in Section 4.1, then the NERC CIP Cyber Security Standards apply. Note that there is a qualification in Section 4.1 that restricts the applicability in the case of Distribution Providers to only those that own certain types of systems and equipment listed in 4.2.

Section “4.2. Facilities” defines the scope of the Facilities, systems, and equipment owned by the Responsible Entity, as qualified in Section 4.1, that is subject to the requirements of the standard. As specified in the exemption section 4.2.3.5, this standard does not apply to Responsible Entities that do not have High Impact or Medium Impact BES Cyber Systems under CIP-002-5.1’s categorization. In addition to the set of BES Facilities, Control Centers, and other systems and equipment, the list includes the set of systems and equipment owned by Distribution Providers. While the NERC Glossary term “Facilities” already includes the BES characteristic, the additional use of the term BES here is meant to reinforce the scope of applicability of these Facilities where it is used, especially in this applicability scoping section. This in effect sets the scope of Facilities, systems, and equipment that is subject to the standards.

Requirement R1:

Requirement R1 exists to reduce the attack surface of Cyber Assets by requiring entities to disable known unnecessary ports. The SDT intends for the entity to know what network accessible (“listening”) ports and associated services are accessible on their assets and systems, whether they are needed for that Cyber Asset’s function, and disable or restrict access to all other ports.

1.1. This requirement is most often accomplished by disabling the corresponding service or program that is listening on the port or configuration settings within the Cyber Asset. It can also be accomplished through using host-based firewalls, TCP_Wrappers, or other means on the Cyber Asset to restrict access. Note that the requirement is applicable at the Cyber Asset level. The Cyber Assets are those which comprise the applicable BES Cyber Systems and their associated Cyber Assets. This control is another layer in the defense against network-based attacks, therefore the SDT intends that the control be on the device itself, or positioned inline in a non-bypassable manner. Blocking ports at the ESP border does not substitute for this device level requirement. If a device has no provision for disabling or restricting logical ports on the device (example - purpose built devices that run from firmware with no port configuration available) then those ports that are open are deemed ‘needed.’

1.2. Examples of physical I/O ports include network, serial and USB ports external to the device casing. BES Cyber Systems should exist within a Physical Security Perimeter in which case the physical I/O ports have protection from unauthorized access, but it may still be possible for accidental use such as connecting a modem, connecting a network cable that bridges networks, or inserting a USB drive. Ports used for ‘console commands’ primarily means serial ports on Cyber Assets that provide an administrative interface.

The protection of these ports can be accomplished in several ways including, but not limited to:


- Disabling all unneeded physical ports within the Cyber Asset’s configuration
- Prominent signage, tamper tape, or other means of conveying that the ports should not be used without proper authorization
- Physical port obstruction through removable locks

The network ports included in the scope of this requirement part are not limited to those on the BES Cyber System itself. The scope of physical network ports includes those ports that may exist on nonprogrammable devices such as unmanaged switches, hubs, or patch panels.

This is a ‘defense in depth’ type control and it is acknowledged that there are other layers of control (the PSP for one) that prevent unauthorized personnel from gaining physical access to these ports. Even with physical access, it has been pointed out there are other ways to circumvent the control. This control, with its inclusion of means such as signage, is not meant to be a preventative control against intruders. Signage is indeed a directive control, not a preventative one. However, with a defense-in-depth posture, different layers and types of controls are required throughout the standard with this providing another layer for depth in Control Center environments. Once physical access has been achieved through the other preventative and detective measures by authorized personnel, a directive control that outlines proper behavior as a last line of defense is appropriate in these highest risk areas. In essence, signage would be used to remind authorized users to “think before you plug anything into one of these systems” which is the intent. This control is not designed primarily for intruders, but for example the authorized employee who intends to plug his possibly infected smartphone into an operator console USB port to charge the battery.

The Applicable Systems column was updated on CIP-007-6 Requirement 1, Part 1.2 to include “Nonprogrammable communication components located inside both a PSP and an ESP.” This should be interpreted to apply to only those nonprogrammable communication components that are inside both an ESP and a PSP in combination, not those components that are in only one perimeter as can be illustrated in the following diagram:

Location of Nonprogrammable Communication Components



Requirement R2:

The SDT’s intent of Requirement R2 is to require entities to know, track, and mitigate the known software vulnerabilities associated with their BES Cyber Assets. It is not strictly an “install every security patch” requirement; the main intention is to “be aware of in a timely manner and manage all known vulnerabilities” requirement.

Patch management is required for BES Cyber Systems that are accessible remotely as well as standalone systems. Standalone systems are vulnerable to intentional or unintentional introduction of malicious code. A sound defense-in-depth security strategy employs additional measures such as physical security, malware prevention software, and software patch management to reduce the introduction of malicious code or the exploit of known vulnerabilities.

One or multiple processes could be utilized. An overall assessment process may exist in a top tier document with lower tier documents establishing the more detailed process followed for individual systems. Lower tier documents could be used to cover BES Cyber System nuances that may occur at the system level.

2.1. The Responsible Entity is to have a patch management program that covers tracking, evaluating, and installing cyber security patches. The requirement applies to patches only, which are fixes released to handle a specific vulnerability in a hardware or software product. The requirement covers only patches that involve cyber security fixes and does not cover patches that are purely functionality related with no cyber security impact. Tracking involves processes for notification of the availability of new cyber security patches for the Cyber Assets. Documenting the patch source in the tracking portion of the process is required to determine when the assessment timeframe clock starts. This requirement handles the situation where security patches can come from an original source (such as an operating system vendor), but must be approved or certified by another source (such as a control system vendor) before they can be assessed and applied in order to not jeopardize the availability or integrity of the control system. The source can take many forms. The National Vulnerability Database, Operating System vendors, or Control System vendors could all be sources to monitor for release of security related patches, hotfixes, and/or updates. A patch source is not required for Cyber Assets that have no updateable software or firmware (there is no user accessible way to update the internal software or firmware executing on the Cyber Asset), or those Cyber Assets that have no existing source of patches such as vendors

that no longer exist. The identification of these sources is intended to be performed once unless software is changed or added to the Cyber Asset's baseline.

2.2. Responsible Entities are to perform an assessment of security related patches within 35 days of release from their monitored source. An assessment should consist of determination of the applicability of each patch to the entity's specific environment and systems. Applicability determination is based primarily on whether the patch applies to a specific software or hardware component that the entity does have installed in an applicable Cyber Asset. A patch that applies to a service or component that is not installed in the entity's environment is not applicable. If the patch is determined to be non-applicable, that is documented with the reasons why and the entity is compliant. If the patch is applicable, the assessment can include a determination of the risk involved, how the vulnerability can be remediated, the urgency and timeframe of the remediation, and the steps the entity has previously taken or will take. Considerable care must be taken in applying security related patches, hotfixes, and/or updates or applying compensating measures to BES Cyber System or BES Cyber Assets that are no longer supported by vendors. It is possible security patches, hotfixes, and updates may reduce the reliability of the system, and entities should take this into account when determining the type of mitigation to apply. The Responsible Entities can use the information provided in the Department of Homeland Security "Quarterly Report on Cyber Vulnerabilities of Potential Risk to Control Systems" as a source. The DHS document "Recommended Practice for Patch Management of Control Systems" provides guidance on an evaluative process. It uses severity levels determined using the Common Vulnerability Scoring System Version 2. Determination that a security related patch, hotfix, and/or update poses too great a risk to install on a system or is not applicable due to the system configuration should not require a TFE.

When documenting the remediation plan measures it may not be necessary to document them on a one to one basis. The remediation plan measures may be cumulative. A measure to address a software vulnerability may involve disabling a particular service. That same service may be exploited through other software vulnerabilities. Therefore disabling the single service has addressed multiple patched vulnerabilities.

2.3. The requirement handles the situations where it is more of a reliability risk to patch a running system than the vulnerability presents. In all cases, the entity either installs the patch or documents (either through the creation of a new or update of an existing mitigation plan) what they are going to do to mitigate the vulnerability and when they are going to do so. There are times when it is in the best interest of reliability to not install a patch, and the entity can document what they have done to mitigate the vulnerability. For those security related patches that are determined to be applicable, the Responsible Entity must within 35 days either install the patch, create a dated mitigation plan which will outline the actions to be taken or those that have already been taken by the Responsible Entity to mitigate the vulnerabilities addressed by the security patch, or revise an existing mitigation plan. Timeframes do not have to be designated as a particular calendar day but can have event designations such as "at next scheduled outage of at least two days duration." "Mitigation plans" in the standard refers to internal documents and are not to be confused with plans that are submitted to Regional Entities in response to violations.

2.4. The entity has been notified of, has assessed, and has developed a plan to remediate the known risk and that plan must be implemented. Remediation plans that only include steps that have been previously taken are considered implemented upon completion of the documentation. Remediation plans that have steps to be taken to remediate the vulnerability must be implemented by the timeframe the entity documented in their plan. There is no maximum timeframe in this requirement as patching and other system changes carries its own risk to the availability and integrity of the systems and may require waiting until a planned outage. In periods of high demand or threatening weather, changes to systems may be curtailed or denied due to the risk to reliability.

Requirement R3:

3.1. Due to the wide range of equipment comprising the BES Cyber Systems and the wide variety of vulnerability and capability of that equipment to malware as well as the constantly evolving threat and resultant tools and controls, it is not practical within the standard to prescribe how malware is to be addressed on each Cyber Asset. Rather, the Responsible Entity determines on a BES Cyber System basis which Cyber Assets have susceptibility to malware intrusions and documents their plans and processes for addressing those risks and provides evidence that they follow those plans and processes. There are numerous options available including traditional antivirus solutions for common operating systems, white-listing solutions, network isolation techniques, Intrusion Detection/Prevention (IDS/IPS) solutions, etc. If an entity has numerous BES Cyber Systems or Cyber Assets that are of identical architecture, they may provide one process that describes how all the like Cyber Assets are covered. If a specific Cyber Asset has no updateable software and its executing code cannot be altered, then that Cyber Asset is considered to have its own internal method of deterring malicious code.

3.2. When malicious code is detected on a Cyber Asset within the applicability of this requirement, the threat posed by that code must be mitigated. In situations where traditional antivirus products are used, they may be configured to automatically remove or quarantine the malicious code. In white-listing situations, the white-listing tool itself can mitigate the threat as it will not allow the code to execute, however steps should still be taken to remove the malicious code from the Cyber Asset. In some instances, it may be in the best interest of reliability to not immediately remove or quarantine the malicious code, such as when availability of the system may be jeopardized by removal while operating and a rebuild of the system needs to be scheduled. In that case, monitoring may be increased and steps taken to insure the malicious code cannot communicate with other systems. In some instances the entity may be working with law enforcement or other governmental entities to closely monitor the code and track the perpetrator(s). For these reasons, there is no maximum timeframe or method prescribed for the removal of the malicious code, but the requirement is to mitigate the threat posed by the now identified malicious code.

Entities should also have awareness of malware protection requirements for Transient Cyber Assets and Removable Media (“transient devices”) in CIP-010-2. The protections required here in CIP-007-6, Requirement R3 complement, but do not meet, the additional obligations for transient devices.

3.3. In instances where malware detection technologies depend on signatures or patterns of known attacks, the effectiveness of these tools against evolving threats is tied to the ability to keep these signatures and patterns updated in a timely manner. The entity is to have a documented process that includes the testing and installation of signature or pattern updates. In a BES Cyber System, there may be some Cyber Assets that would benefit from the more timely installation of the updates where availability of that Cyber Asset would not jeopardize the availability of the BES Cyber System’s ability to perform its function. For example, some HMI workstations where portable media is utilized may benefit from having the very latest updates at all times with minimal testing. Other Cyber Assets should have any updates thoroughly tested before implementation where the result of a ‘false positive’ could harm the availability of the BES Cyber System. The testing should not negatively impact the reliability of the BES. The testing should be focused on the update itself and if it will have an adverse impact on the BES Cyber System. Testing in no way implies that the entity is testing to ensure that malware is indeed detected by introducing malware into the environment. It is strictly focused on ensuring that the update does not negatively impact the BES Cyber System before those updates are placed into production.

Requirement R4:

Refer to NIST 800-92 and 800-137 for additional guidance in security event monitoring.

4.1. In a complex computing environment and faced with dynamic threats and vulnerabilities, it is not practical within the standard to enumerate all security-related events necessary to support the activities for alerting and incident response. Rather, the Responsible Entity determines which computer generated events are necessary to log, provide alerts and monitor for their particular BES Cyber System environment.

Specific security events already required in Version 4 of the CIP Standards carry forward in this version. This includes access attempts at the Electronic Access Points, if any have been identified for a BES Cyber Systems. Examples of access attempts include: (i) blocked network access attempts, (ii) successful and unsuccessful remote user access attempts, (iii) blocked network access attempts from a remote VPN, and (iv) successful network access attempts or network flow information.

User access and activity events include those events generated by Cyber Assets within the Electronic Security Perimeter that have access control capability. These types of events include: (i) successful and unsuccessful authentication, (ii) account management, (iii) object access, and (iv) processes started and stopped.

It is not the intent of the SDT that if a device cannot log a particular event that a TFE must be generated. The SDT's intent is that if any of the items in the bulleted list (for example, user logouts) can be logged by the device then the entity must log that item. If the device does not have the capability of logging that event, the entity remains compliant.

4.2. Real-time alerting allows the cyber system to automatically communicate events of significance to designated responders. This involves configuration of a communication mechanism and log analysis rules. Alerts can be configured in the form of an email, text message, or system display and alarming. The log analysis rules can exist as part of the operating system, specific application or a centralized security event monitoring system. On one end, a real-time alert could consist of a set point on an RTU for a login failure, and on the other end, a security event monitoring system could provide multiple alerting communications options triggered on any number of complex log correlation rules.

The events triggering a real-time alert may change from day to day as system administrators and incident responders better understand the types of events that might be indications of a cyber-security incident. Configuration of alerts also must balance the need for responders to know an event occurred with the potential inundation of insignificant alerts. The following list includes examples of events a Responsible Entity should consider in configuring real-time alerts:

- Detected known or potential malware or malicious activity
- Failure of security event logging mechanisms
- Login failures for critical accounts
- Interactive login of system accounts
- Enabling of accounts
- Newly provisioned accounts
- System administration or change tasks by an unauthorized user
- Authentication attempts on certain accounts during non-business hours
- Unauthorized configuration changes
- Insertion of Removable Media in violation of a policy

4.3 Logs that are created under Part 4.1 are to be retained on the applicable Cyber Assets or BES Cyber Systems for at least 90 days. This is different than the evidence retention period called for in the CIP standards used to prove historical compliance. For such audit purposes, the entity should maintain evidence that shows that 90 days were kept historically. One example would be records of disposition of event logs beyond 90 days up to the evidence retention period.

4.4. Reviewing logs at least every 15 days (approximately every two weeks) can consist of analyzing a summarization or sampling of logged events. NIST SP800-92 provides a lot of guidance in periodic log analysis. If a centralized security event monitoring system is used, log analysis can be performed top-down starting with a review of trends from summary reports. The log review can also be an extension of the exercise in identifying those events needing real-time alerts by analyzing events that are not fully understood or could possibly inundate the real-time alerting.

Requirement R5:

Account types referenced in this guidance typically include:

- Shared user account: An account used by multiple users for normal business functions by employees or contractors. Usually on a device that does not support Individual User Accounts.
- Individual user account: An account used by a single user.
- Administrative account: An account with elevated privileges for performing administrative or other specialized functions. These can be individual or shared accounts.
- System account: Accounts used to run services on a system (web, DNS, mail etc.). No users have access to these accounts.
- Application account: A specific system account, with rights granted at the application level often used for access into a Database.
- Guest account: An individual user account not typically used for normal business functions by employees or contractors and not associated with a specific user. May or may not be shared by multiple users.
- Remote access account: An individual user account only used for obtaining Interactive Remote Access to the BES Cyber System.
- Generic account: A group account set up by the operating system or application to perform specific operations. This differs from a shared user account in that individual users do not receive authorization for access to this account type.

5.1 Reference the Requirement’s rationale.

5.2 Where possible, default and other generic accounts provided by a vendor should be removed, renamed, or disabled prior to production use of the Cyber Asset or BES Cyber System. If this is not possible, the passwords must be changed from the default provided by the vendor. Default and other generic accounts remaining enabled must be documented. For common configurations, this documentation can be performed at a BES Cyber System or more general level.

5.3 Entities may choose to identify individuals with access to shared accounts through the access authorization and provisioning process, in which case the individual authorization records suffice to meet this Requirement Part. Alternatively, entities may choose to maintain a separate listing for shared accounts. Either form of evidence achieves the end result of maintaining control of shared accounts.

5.4. Default passwords can be commonly published in vendor documentation that is readily available to all customers using that type of equipment and possibly published online.

The requirement option to have unique password addresses cases where the Cyber Asset generates or has assigned pseudo-random default passwords at the time of production or installation. In these cases, the default password does not have to change because the system or manufacturer created it specific to the Cyber Asset.

5.5. Interactive user access does not include read-only information access in which the configuration of the Cyber Asset cannot change (e.g. front panel displays, web-based reports, etc.). For devices that cannot technically or for operational reasons perform authentication, an entity may demonstrate all interactive user access paths, both remote and local, are configured for authentication. Physical security suffices for local access configuration if the physical security can record who is in the Physical Security Perimeter and at what time.

Technical or procedural enforcement of password parameters are required where passwords are the only credential used to authenticate individuals. Technical enforcement of the password parameters means a Cyber Asset verifies an individually selected password meets the required parameters before allowing the account to authenticate with the selected password. Technical enforcement should be used in most cases when the authenticating Cyber Asset supports enforcing password parameters. Likewise, procedural enforcement means

requiring the password parameters through procedures. Individuals choosing the passwords have the obligation of ensuring the password meets the required parameters.

Password complexity refers to the policy set by a Cyber Asset to require passwords to have one or more of the following types of characters: (1) lowercase alphabetic, (2) uppercase alphabetic, (3) numeric, and (4) non-alphanumeric or “special” characters (e.g. #, \$, @, &), in various combinations.

5.6 Technical or procedural enforcement of password change obligations are required where passwords are the only credential used to authenticate individuals. Technical enforcement of password change obligations means the Cyber Asset requires a password change after a specified timeframe prior to allowing access. In this case, the password is not required to change by the specified time as long as the Cyber Asset enforces the password change after the next successful authentication of the account. Procedural enforcement means manually changing passwords used for interactive user access after a specified timeframe.

5.7 Configuring an account lockout policy or alerting after a certain number of failed authentication attempts serves to prevent unauthorized access through an online password guessing attack. The threshold of failed authentication attempts should be set high enough to avoid false-positives from authorized users failing to authenticate. It should also be set low enough to account for online password attacks occurring over an extended period of time. This threshold may be tailored to the operating environment over time to avoid unnecessary account lockouts.

Entities should take caution when configuring account lockout to avoid locking out accounts necessary for the BES Cyber System to perform a BES reliability task. In such cases, entities should configure authentication failure alerting.

Rationale:

During development of this standard, text boxes were embedded within the standard to explain the rationale for various parts of the standard. Upon BOT approval, the text from the rationale text boxes was moved to this section.

Rationale for Requirement R1:

The requirement is intended to minimize the attack surface of BES Cyber Systems through disabling or limiting access to unnecessary network accessible logical ports and services and physical I/O ports.

In response to FERC Order No. 791, specifically FERC’s reference to NIST 800-53 rev. 3 security control PE-4 in paragraph 149, Part 1.2 has been expanded to include PCAs and nonprogrammable communications components. This increase in applicability expands the scope of devices that receive the protection afforded by the defense-in-depth control included in Requirement R1, Part 1.2.

The applicability is limited to those nonprogrammable communications components located both inside a PSP and an ESP in order to allow for a scenario in which a Responsible Entity may implement an extended ESP (with corresponding logical protections identified in CIP-006, Requirement R1, Part 1.10). In this scenario, nonprogrammable components of the communication network may exist out of the Responsible Entity’s control (i.e. as part of the telecommunication carrier’s network).

Rationale for Requirement R2:

Security patch management is a proactive way of monitoring and addressing known security vulnerabilities in software before those vulnerabilities can be exploited in a malicious manner to gain control of or render a BES Cyber Asset or BES Cyber System inoperable.

Rationale for Requirement R3:

Malicious code prevention has the purpose of limiting and detecting the addition of malicious code onto the applicable Cyber Assets of a BES Cyber System. Malicious code (viruses, worms, botnets, targeted code such as Stuxnet, etc.) may compromise the availability or integrity of the BES Cyber System.

Rationale for Requirement R4:

Security event monitoring has the purpose of detecting unauthorized access, reconnaissance and other malicious activity on BES Cyber Systems, and comprises of the activities involved with the collection, processing, alerting and retention of security-related computer logs. These logs can provide both (1) the detection of an incident and (2) useful evidence in the investigation of an incident. The retention of security-related logs is intended to support post-event data analysis.

Audit processing failures are not penalized in this requirement. Instead, the requirement specifies processes which must be in place to monitor for and notify personnel of audit processing failures.

Rationale for Requirement R5:

To help ensure that no authorized individual can gain electronic access to a BES Cyber System until the individual has been authenticated, i.e., until the individual's logon credentials have been validated. Requirement R5 also seeks to reduce the risk that static passwords, where used as authenticators, may be compromised.

Requirement Part 5.1 ensures the BES Cyber System or Cyber Asset authenticates individuals that can modify configuration information. This requirement addresses the configuration of authentication. The authorization of individuals is addressed elsewhere in the CIP Cyber Security Standards. Interactive user access does not include read-only information access in which the configuration of the Cyber Asset cannot change (e.g. front panel displays, web-based reports, etc.). For devices that cannot technically or for operational reasons perform authentication, an entity may demonstrate all interactive user access paths, both remote and local, are configured for authentication. Physical security suffices for local access configuration if the physical security can record who is in the Physical Security Perimeter and at what time.

Requirement Part 5.2 addresses default and other generic account types. Identifying the use of default or generic account types that could introduce vulnerabilities has the benefit ensuring entities understand the possible risk these accounts pose to the BES Cyber System. The Requirement Part avoids prescribing an action to address these accounts because the most effective solution is situation specific, and in some cases, removing or disabling the account could have reliability consequences.

Requirement Part 5.3 addresses identification of individuals with access to shared accounts. This Requirement Part has the objective of mitigating the risk of unauthorized access through shared accounts. This differs from other CIP Cyber Security Standards Requirements to authorize access. An entity can authorize access and still not know who has access to a shared account. Failure to identify individuals with access to shared accounts would make it difficult to revoke access when it is no longer needed. The term “authorized” is used in the requirement to make clear that individuals storing, losing, or inappropriately sharing a password is not a violation of this requirement.

Requirement 5.4 addresses default passwords. Changing default passwords closes an easily exploitable vulnerability in many systems and applications. Pseudo-randomly system generated passwords are not considered default passwords.

For password-based user authentication, using strong passwords and changing them periodically helps mitigate the risk of successful password cracking attacks and the risk of accidental password disclosure to unauthorized individuals. In these requirements, the drafting team considered multiple approaches to ensuring this requirement was both effective and flexible enough to allow Responsible Entities to make good security decisions. One of the approaches considered involved requiring minimum password entropy, but the calculation for true information entropy is more highly complex and makes several assumptions in the passwords users choose. Users can pick poor passwords well below the calculated minimum entropy.

The drafting team also chose to not require technical feasibility exceptions for devices that cannot meet the length and complexity requirements in password parameters. The objective of this requirement is to apply a measurable password policy to deter password cracking attempts, and replacing devices to achieve a specified password policy does not meet this objective. At the same time, this requirement has been strengthened to require account lockout or alerting for failed login attempts, which in many instances better meets the requirement objective.

The requirement to change passwords exists to address password cracking attempts if an encrypted password were somehow attained and also to refresh passwords which may have been accidentally disclosed over time. The requirement permits the entity to specify the periodicity of change to accomplish this objective. Specifically, the drafting team felt determining the appropriate periodicity based on a number of factors is more effective than specifying the period for every BES Cyber System in the Standard. In general, passwords for user authentication should be changed at least annually. The periodicity may increase in some cases. For example, application passwords that are long and pseudo-randomly generated could have a very long periodicity. Also, passwords used only as a weak form of application authentication, such as accessing the configuration of a relay may only need to be changed as part of regularly scheduled maintenance.

The Cyber Asset should automatically enforce the password policy for individual user accounts. However, for shared accounts in which no mechanism exists to enforce password policies, the Responsible Entity can enforce the password policy procedurally and through internal assessment and audit.

Requirement Part 5.7 assists in preventing online password attacks by limiting the number of guesses an attacker can make. This requirement allows either limiting the number of failed authentication attempts or alerting after a defined number of failed authentication attempts. Entities should take caution in choosing to limit the number of failed authentication attempts for all accounts because this would allow the possibility for a denial of service attack on the BES Cyber System.