NERC

NORTH AMERICAN ELECTRIC RELIABILITY CORPORATION

Response to Comments

Project 2016-02 Modifications to CIP Standards Draft 2 Posting ending September 1, 2021

Additional information is available on the project page. If you have questions, contact Senior Standards Developer, Jordan Mallory (404-446-2589).

Background Information

Project 2016-02 (1) addresses the Federal Energy Regulatory Commission (Commission) directives contained in Order No. 822 and (2) considers the Version 5 Transition Advisory Group (V5TAG) issues identified in the CIP V5 Issues for Standard Drafting Team Consideration (V5TAG Transfer Document).

The V5TAG, which consisted of representatives from NERC, Regional Entities, and industry stakeholders, was formed to issue guidance regarding possible methods to achieve compliance with the CIP Version 5 standards and to support industry's implementation activities. During the V5TAG's activities, it identified certain issues with the CIP Reliability Standards that would be better addressed by a standard drafting team (SDT) for the CIP Reliability Standards. The V5TAG developed the CIP Version 5 Transition Advisory Group Issues for Consideration document to formally recommend that the SDT address these issues and consider modifications to the standard language during the standards development process. Among other issues, the V5TAG stated "The CIP Version 5 standards comments.

There were 93 sets of responses, including comments from approximately 218 different people from approximately 137 companies representing 10 of the Industry Segments as shown in the table on the following pages.

Question 1.

Are the two options for identification of SCI within CIP-002 clear and is it understood that when SCI is included in the CIP Systems that it is treated like the CIP System, it is a part of for CIP Requirement Applicability?

Q1 Comment Themes

- **The options are mostly clear and acceptable with the exception of the phrase "independent SCI supporting"**
- **The term 'SCI' is still unclear and ambiguous**
- **Does not agree that the two options for identification of SCI are clear. The distinction between SCI included in a BES Cyber System (BCS) and SCI operating independently adds an unnecessary level of complexity to the standards**

SDT Response: Thank you for your comments. The SDT agrees with this theme that much more clarity around SCI is needed and has taken a different and simplified approach for the next draft. The two options remain for entities to either:

1) dedicate a virtualized infrastructure to hosting Virtual Cyber Assets (VCA) of the same CIP impact rating or 'associated with' the same impact rating (i.e., "all-in") or,
2) share the virtualized infrastructure hosting Virtual Cyber Assets (VCA) of different CIP impact ratings or non-CIP impact rated (previously known as "SCI Identified independently").

The terminology of this latter case ("identified independently") will no longer be used. In order to simplify these options, the term SCI will be redefined to ONLY refer to the second option.

For the "all-in" scenario, the virtualized infrastructure hardware components are now simply Cyber Assets that are a part of the systems they are hosting, such as BCAs in a BCS, or as Cyber Assets in an virtual EACMS, PACS, or PCA. In this scenario, every VCA hosted on the Cyber Assets, and the Cyber Assets themselves, are of the same impact rating. *There is no SCI in this option*. The SDT believes this is the way most entities with existing virtual systems have identified and categorized their virtual infrastructures to date and this will allow that to remain unchanged.

The SDT in response to this comment theme has modified the SCI definition so that it ONLY applies to the second scenario, where the "S" in "SCI" means the underlying hardware *is shared among VCAs of differing impact levels* (i.e., 'mixed mode'). This simplification allows for distinct scoping of new additional requirements for the proper isolation required as well as restricting access to the management plane where the isolation is configured.

To recap:

- If the entity implements BCS and any associated EACMS or PACS or PCAs as VCAs on an underlying hardware infrastructure, and all is treated as or 'associated with' the same (highest) impact rating, this is "all-in" and the underlying hardware is identified as Cyber Assets that may also be BCAs in a BCS or part of the EACMS or PACS they host. There is no SCI in this scenario. No CIP-002 recategorization is required.
- If the entity hosts VCAs of varying impact levels and wishes to use the security controls to keep each impact level environment isolated from each other, then the underlying hardware is identified as SCI and additional requirements for such isolation are required for this environment.

Question 2.

The Applicable Systems column may include "SCI identified independently…" Is this clear or is additional clarification (such as "SCI identified as supporting, but not part of…") needed?

SDT Response:

Thank you for your comments. Most entities answered this question, which is closely related to Q1 with shared themes from Q1. Please see Q1 responses for the SDT response to those themes.

Question 3.

The SDT modified the ERC definition to reference "outside the asset containing". This is to allow scoping based on connectivity of the logging systems as required by CIP-007 Requirement R4 as well as the scoping of requirement parts in CIP-004 and CIP-006 based on risk. Do you agree with the proposed change? If not, please provide the basis for your disagreement and an alternate proposal.

Q3 Comment Theme:

- **Use of "asset" for scoping is unclear**

SDT Response: Thank you for your comments. The SDT agrees with comments that the "asset containing" is not defined precisely enough for use in this foundational scoping term of ERC. While the intent of ERC is to restrict the scope of certain requirements for assets/sites that do not have the requisite type of communications, this has traditionally been accomplished by using the ESP as the demarcation point of ERC, not the "asset containing". The SDT will return ERC to being ESP-based. In addition, the ERC definition essentially mirrors the currently approved definition with one change of replacing concepts of "inside or outside" an ESP to "through its ESP" to better position it for Zero Trust architectures that are not strictly network perimeter based.

Q3 Comment Theme:

- **Expands the existing definition by changing from "BES Cyber System" to "CIP System" and by changing from "to access" to "to communicate".**

SDT Response: Thank you for your comments. The SDT agrees and has reinstated the BES Cyber System as the subject of ERC definition and reverted to the currently approved "access" language. The ERC definition now essentially mirrors the currently approved definition with one change of replacing the concepts of "inside or outside" an ESP to "through its ESP" to better position it for Zero Trust architectures that are not strictly network perimeter based.

Question 4.

The SDT proposes that the modified ESP definition can be used for both traditional firewall based networks, as well as future networks such as zero trust. Do you agree with the proposed change? If not, please provide the basis for your disagreement and an alternate proposal.

Q4 Comment Theme:

- **EACMS appears to be required**

SDT Response: Thank you for your comments. The SDT agrees with the commentors. The "enforced by an EACMS" phrase, while meant to help define the extent of the ESP especially in Zero Trust architectures, did require an EACMS in order for an ESP to exist. As commenters rightly pointed out, this negated being able to define ESPs for isolated networks. The SDT addressed this concern by reinstating the current ESP definition and added a second component , "; or a logical boundary defined by one or more EAPs" to address zero trust.

Q4 Comment Theme:

- **ESP does not include routable scoping**

SDT Response: Thank you for your comments. The SDT agrees and worked routable protocol scoping back into the definition.

Q4 Comment Theme:

- **PCA may be redundant in ESP definition**
- **ESP no longer includes boundary**

SDT Response:Thank you for your comments. The SDT agrees with the various comments in this theme, which revolved around being able to define an ESP around isolated networks that do not have ERC but still need an ESP for defining PCAs, knowing the extent of TCA connections, etc. While the SDT's draft 2 proposed definition was more objective in order to incorporate rather different security models (traditional 'castle/moat' perimeter vs. Zero Trust), the SDT agrees the definition needs more specificity in order to not lose its other boundary-like functions that it fulfills even in isolated networks that do not have external connectivity. The SDT also agrees that the PCA definition adequately defines those Cyber Assets, therefore the ESP definition doesn't need to specifiy it includes or groups PCAs.

Q4 Comment Theme:

- **The word "policy" is not well understood**

SDT Response: Thank you for your comments. The SDT agrees that 'policy' needs more context in order to be well understood whether it is a business or administrative policy document or a technical configuration of a firewall policy. The SDT addressed these comments by removing the term 'policy' where it was intended to reflect technical controls, (e.g. software designed networks, instruction detection policies, or firewall rules/policies, etc.) and replaced it with 'configuration or settings'.

Q4 Comment Theme:

- **The use of 'CIP Systems' in the definition is too broad**

SDT Response: Thank you for your comments. The SDT agrees as all types of 'CIP Systems' are not required to be within an ESP. While not a requirement in CIP-005, it is an implication of the definition. The SDT eliminated this new definition and its reference from the ESP definition.

Question 5.

The SDT modified the IRA definition based on industry comments. Do you agree with the proposed change? If not, please provide the basis for your disagreement and an alternate proposal.

Q5 Comment Theme:

- **The protocol conversion bullet could expand scope beyond the intent (Control Center operator to field device as IRA). The second bullet point is unclear and ambiguous.**

SDT Response: Thank you for your comments. The SDT asserts that the 2nd bullet was correct but agrees it was complex and has made modifications to simplify it. The intent is to cover the situation where an interactive user is using a routable protocol (i.e., an IP address) and is able to access a BES Cyber System and operate or configure it, often called terms such as “engineering access”, even if the BES Cyber System itself uses only non-routable communications, such as a serial cable to a console port. In this case, the CIP-005 R2 controls should apply to the routable portion of the communications, and this bullet in the IRA definition is designed to capture that scenario, but with simplified language in our next draft.

Q5 Comment Theme:

- **“real-time” needs deletion/clarification/capitalization.**

SDT Response:

Thank you for your comments. The SDT agrees. The intent was to further define that ‘interactive’ access was in ‘real-time’, not access a user has programmed for a future time such as a scheduled batch job that runs when the user is no longer logged in. While this is a risk, its is not the subject of CIP-005 R2’s security controls, which are designed for a real-time, interactive user. Batch jobs can’t execute through a “pane of glass” interface of a typical Intermediate System, or enter a user’s current code from a physical token, or receive a text with a one-time-password from the user’s phone. Therefore, IRA is not all-inclusive but is designed around this high risk area of interactive, ‘real-time’ interaction with a BES Cyber System and isolating the remote user’s machine (through an Intermediate System), strongly authenticating the user (through multi-factor authentication), and protecting the session’s traffic over less trusted networks (encryption from the Intermediate System to the client). However, as commenters pointed out, the term ‘real-time’ was confusing or unnecessary, so the SDT deleted it from the definition.

Q5 Comment Theme:

- **Add “IRA does not include system-to-system process communications” back to the definition.**

SDT Response: Thank you for your comments. The SDT believes this comment was made related to the above comment concerning the term ‘real-time’ in the definition. There was a sense of if the term ‘real-time’ stayed in, this phrase needed to be added back to state ‘real-time’ communications did not include system-to-system communications (e.g., a system polling an RTU). The SDT removed the term ‘real-time’ to respond to these comments and asserts that the first phrases of the definition precludes system-to-system communications without an explicit negative phrase in the definition.

Question 6.

The SDT modified the Management Interface definition based on industry comments. Do you agree with the proposed change? If not, please provide the basis for your disagreement and an alternate proposal.

Q6 Comment Theme:

- **The second bullet concerning “lights out” needs clarification/is unclear/is vendor specific.**

SDT Response: Thank you for your comments. The SDT does not assert that "lights out management" is a vendor-specific term, but due to the number of comments the SDT changed this second bullet back to the language that had been proposed in the SDT's draft 1 definition of 'Management Module' as suggested by many in the comments.

Q6 Comment Theme:

- **Ambiguity about the intent of the 'user interface, logical interface, or dedicated physical port**

SDT Response: Thank you for your comments. The SDT agrees and replaced this phrasing with "An administrative interface on a Shared Cyber Infrastructure or Electronic Access Control or Monitoring System that…" which is clearer as to the intent of a Management Interface rather than a description of the several forms it may take.

Question 7.

As discussed in the CIP Definitions and Exemptions Technical Rationale (TR), the SDT believes that the use of configurations or policy in the modified ESP definition can reduce the burden of documenting ESPs in a zero trust environment. Do you agree with the proposed change? If not, please provide the basis for your disagreement and an alternate proposal.

Q7 Comment Theme:

- **Agreement/Disagreement with reduced documentation burden**

SDT Response: Thank you for your comments. Most of the comment themes for Q7 were either agreement or disagreement with the reduction in documentation burden based on the overall ESP themes reported in Q4. See Q4 responses to those specific themes. The SDT would like to take this opportunity to clarify the question and its position, which is in agreement with many points commenters made.

The goal of CIP-005 R1 has, from the beginning of CIP-005, been to establish a "policy or configuration" of who/what can communicate with a BCS or its associated PCA's inside a protected network (ESP) and ensure that each element of that policy/configuration allowing communication is necessary; that it implements the security principle of least privilege. This is accomplished in today's more prescriptive perimeter-based architecture by having a documented list of ESPs, which leads to a documented list of EAPs, which then leads to documentation of the policy/configuration within each EAP which is the basis for determining whether the ESP is compliant with the intent of the requirement. So today, ESPs and EAPs are mechanisms for getting to the actual policy/configuration of what is allowed to communicate across the boundary. That intent does not change; the SDT's approach states the eventual policy/configuration as the direct object of the requirement rather than indirectly through documented lists of ESPs and EAPs.

However, the SDT is attempting to change the standards so that other models such as Zero Trust (ZT) are not prohibited or discouraged by overly prescriptive language. As several commenters pointed out, in ZT the number of ESP/EAPs will increase as ESPs "shrink" from protecting a network segment down to protecting a single resource or service, so in that sense going directly to the policy/configuration reduces the burden of the intermediate step of documenting ESPs as the number of those may increase exponentially in full ZT implementations. As other commenters pointed out, ZT with its more granular access control will potentially greatly increase the number of policy/configuration statements of "what can talk to what", so it does not reduce documentation in that regard. The SDT agrees.

Question 8.

The SDT added new and revised several defined terms to incorporate virtualization and future technologies within the CIP Standards. Do you agree with the proposed changes to the NERC Glossary terms? If not, please provide the basis for your disagreement and an alternate proposal.

NOTE: Many comments for Q8 referred back to issues with the specific questions on specific definitions in Q1-6. Please see responses to those issues in the appropriate question above.

Q8 Comment Theme:

- **Intermediate System: The statement "The Intermediate System must not be located inside the Electronic Security Perimeter" should be retained**

SDT Response: Thank you for your comments. The SDT disagrees and asserts that the "where" of the Intermediate System is a requirement that is embedded within the current definition. If an entity places the Intermediate System inside an ESP such that "public" encrypted traffic is allowed in from yet-to-be authenticated users, that should be a violation of a requirement, not simply not meeting a definition. The SDT removed this phrase from the definition and authored a requirement in CIP-005 R2.6 to address this situation.

Q8 Comment Theme:

- **CIP Systems: Remove "TCA" from "CIP Systems". TCAs are included in the CIP System definition. However, CIP System is used in the definition of ESP.**

SDT Response: Thank you for your comments. The SDT, in response to comments, removed the nested use of "CIP Systems" in other definitions. Upon further analysis, it was then only used in a non-essential clarification in CIP-010 R1.2 that has been removed. The SDT has deleted the proposed term "CIP Systems".

Q8 Comment Theme:

- **Cyber System: Revert back to Cyber Assets or go with generic "cyber system" term, leaving it undefined**
- **Cyber Asset: Consider not excluding SCI, keeping the original**

SDT Response: Thank you for your comments. After mindful consideration of comments, the SDT disagrees that Cyber Asset is now one form an object, with applicable CIP requirements, can take; an inclusive hardware-based electronic device with all the software and data in the device. Virtualization fundamentally breaks the 1:1 tie between hardware resources and the varied and dynamic VCA that may happen to execute on it (be "in the device"). The SDT intent is that the objects of CIP requirements can now take one of three forms: Cyber Assets, Virtual Cyber Assets, or SCI. This term ("Cyber System") was created as an attempt to reduce admistrative burden on Registered Entities by establishing this as shorthand for those three forms so all three do not have to be repeated to update every occurrence of 'Cyber Asset' within all CIP documentation. In addition, in Draft 1, the SDT attempted using the generic words 'cyber system' as an undefined term in the exemptions and stakeholder comments were numerous that this term must be a defined glossary term.

The SDT analyzed the comments regarding removal of the exclusion for SCI from the Cyber Asset definition, and after consideration has chosen to retain this exclusion in the Cyber Asset definition for sound cause. If the underlying hardware resource pool of SCI is a "Cyber Asset" then that definition consumes ALL software (VCA) on the SCI. The all-inclusive nature of the Cyber Asset definition is a core issue with SCI scenarios, and the exclusion is necessary to prevent this condition.

Q8 Comment Theme:

- **Protected Cyber Asset (PCA): Additional clarifications on the qualifier of "excluding Virtual Cyber Assets that are being actively remediated prior to introduction to the ESP."**

SDT Response: Thank you for your comments. The SDT's intent is this is part of a set of changes to ensure the CIP requirements and definitions do not prohibit or discourage automated means of vulnerability assessment and remediation in virtualized environments (a.k.a "Remediation VLANs"). In these instances, a VCA already resides "in a production environment" but whenever it is instantiated, it is placed on an isolated network where it is assessed against the security policy and remediated as needed to meet the expected security hygiene level before being placed on a production network, all in an automated fashion. However, although during this assessment/remediation process it is isolated from a network perspective, it may for a very brief period of time meet the 2nd part of the proposed PCA definition in that it may be executing on the same hypervisor as a BCS. The SDT asserts that since it is isolated from a network perspective while it is being remediated in an automated fashion, it should not require separate and distinct hypervisor infrastructure for this brief period of time. This is the reason for this exclusionary qualifier in the PCA definition.

Q8 Comment Theme:

- **Overall: Some definitions include requirements on CPU and memory (isolation/shared) and are not future proof**

SDT Response: Thank you for your comments. The SDT intent in these definitions and related requirements is to mitigate the risk unique to virtualization of hardware based 'side channel' attacks from other VCAs executing on the same hardware (CPU/Memory). Attacks such as Rowhammer where one process can potentially change code or data in physically adjacent memory belonging to another process, or Spectre/Meltdown where a hardware CPU vulnerability could allow access to other processes are the risk, as a couple examples. The SDT has considered higher level objective language such as "mitigate the risks of side-channel attacks" as suggested by other commenters but finds that to be overly broad. This sharing of physical CPU and memory is a virtualization specific risk that is being mitigated and is specific enough that entities can show affinity rule configurations as evidence of compliance when hosted on SCI. Using such terms may not be ultimately future-proof, but the SDT is having to balance the security objectives with enough specificity that entities know how to comply and how to provide evidence of compliance.

Q8 Comment Theme:

- **PCA: The CPU/memory affinity requirements are problematic in that it essentially places a requirement inside a definition. The implicit requirement is that a BES Cyber Asset cannot share CPU/memory with a non-BCS of the same impact rating or a Protected Cyber Asset.**

SDT Response: Thank you for your comments. After mindful consideration of comments, the SDT disagrees that this is a requirement within a definition. The SDT intent is for it to be each entity's choice whether to 1) employ affinity rules or, 2) separate VCA's of differing impact levels onto separate hypervisors. However, should an entity choose to not employ methods to keep VCAs of different impact levels from executing on the same CPU/memory subjecting those VCAs to hardware-based side-channel attacks, then those VCAs that are not BCAs are properly defined as an associated PCA; creating the need for this language within the PCA definition. If the entity chooses to employ methods to separate VCAs onto different hypervisors or server blades so they do not execute on the same CPU/memory simultaneously, then the risk is mitigated and they no longer meet the definition of a PCA from that perspective. There is no requirement to employ such methods, just the consequence of becoming a PCA if the entity chooses not to.

Q8 Comment Theme:

- **Overall: Since the Glossary modifications are the foundation to all Standard changes, NERC should seek approval of the new terms prior to any changes being introduced in the Standards to reduce potential misunderstanding or misinterpretation**

SDT Response: Thank you for your comments. The SDT suggests that posting definitions alone would not give stakeholders the necessary context of how a definition is used within the standards and would hamper stakeholder acceptance. The SDT asserts that posting these together with the applicable Standard changes is a necessary part for gaining stakeholder acceptance, especially for the very context-sensitive technical terms in CIP, is dependent upon having the context of its use within requirements.

Question 9.

The SDT revised CIP-002 based on industry comments. Do you agree with the proposed changes to the CIP-002 Reliability Standard? If not, please provide the basis for your disagreement and an alternate proposal.

NOTE: Several commenters reported they cannot support CIP-002 changes due to concerns with the definition of SCI expressed in Q1. Please see responses to Q1 above.

Q9 Comment Theme:

- **CIP-002 R1 requirements should be restored and all other SCI language in CIP-002 Attachment 1 should be removed**

SDT Response: Thank you for your comments. After the latested modification of the SCI definition (See Q1 above), the SDT agrees and has reverted the CIP-002 requirements to currently approved language. As our scope was virtualization, we only covered the identification of SCI. However, the identification of "all other" types of Cyber Systems other than the BES Cyber Systems (EACMS, PACS, PCAs, SCI, etc.) is a larger issue that should be considered holistically.

Q9 Comment Theme:

- **Concern that in the Exemptions both data communication links and communication networks should be exempt as in currently approved language.**

SDT Response: Thank you for your comments. The SDT's intent was that shortening to the phrase "communication links" included both, but in response to comments the SDT has reincorporating the full language back into the exemptions.

Q9 Comment Theme:

- **Request clarification of capitalization of "High Impact," "Medium Impact" and "Low Impact."**

SDT Response:

Thank you for your comments. The SDT agrees and the capitalization of these terms was reviewed and corrected.

Question 10.

The SDT revised CIP-005 based on industry comments. Do you agree with the proposed changes to the NERC Glossary terms? If not, please provide the basis for your disagreement and an alternate proposal.

NOTE: Many comment themes in CIP-005 are based on issues with individual definitions (ESP, EAP, IRA, SCI, etc.). For responses to those themes, please see the appropriate question above.

Q10 Comment Theme:

- **Request clarification between Part 1.2 and Part 1.3. Part 1.2 reads "Permit only needed and controlled communications to and from Management Interfaces, and deny all other communications" which seems to include Part 1.3 which reads "Deny network communications from Applicable Systems of Part 1.1 to the Management Interface, per system capability."**
- **In addition electronic access controls should be in-place on the Management Interface to only allow the appropriate administrators access to it. The prescriptive and broad ban of traffic from the data plane seems excessive. The separation occurs in cloud-based hosting because the hosting service is the administrator and the tenant is not. For on-premise that is not the case and the administrator and tenant are the same entity.**

SDT Response:

Thank you for your comments. The SDT asserts that having the VCAs, on SCI, not have access to the management plane is a best practice, but agrees with the issues brought up by the commenters. Since a person using IRA to a Management Interface will have the controls of CIP-005 R2 applied, and the overarching R1.2 (now 1.3 due to the reinstatement of 1.1 and the renumbering) permits only needed communications to and from the Management Interfaces with justification, the SDT agrees, and deleted the Draft 2 Part 1.3 language.

Q10 Comment Theme:

- **The host-based firewall exclusion in R1 is not appropriate for objective-based model/needs clarification/may be needed for Zero Trust implementation/should be in Implementation Guidance.**

SDT Response: Thank you for your comments. The SDT agrees. The intent of the exclusion was to avoid the scenario of a BCS being plugged into the Internet with only the internal OS firewall turned on. A compromise of the host through any OS vulnerability was then also a compromise of the firewall itself. It met the higher level objective, but with an insufficient separation of the security control from the object protected. The SDT has deleted the exclusion language and reinstated R1.1 in order to address the issue.

Q10 Comment Theme:

- **Requirement 1.1 includes a new insertion "between intelligent electronic devices", what is an intelligent electronic device?**

SDT Response:

Thank you for your comments. The SDT's intent was to bring forward the exclusion from the low impact scenario where ESPs and firewalls are excluded in very time sensitive applications such as relays communicating with each other for asset protection functions. The term "intelligent electronic devices" is a term of art, but remains undefined and the SDT agrees is an overly generic term. The SDT is proposing to use the NERC Glossary defined term of "Protection System" in its place, and moved the exclusion to Part 1.2 after the reinstatment of Part 1.1.

Q10 Comment Theme:

- **Hall of mirrors is created in R2.1 where an Intermediate System is needed for an Intermediate System**

SDT Response:

Thank you for your comments. The SDT agrees. This was a document error that was only in the clean but not the redline versions of the standard. This has been corrected.

Q10 Comment Theme:

- **R1.6 on detecting malicious communications needs consideration of the "where" in light of Zero Trust and changes to the ESP.**

SDT Response:

Thank you for your comments. The SDT has changed the requirement in response so that it refers to the communications that are entering or leaving an ESP and changed the Applicable Systems column to describe which ESPs – it applies to high impact BCS and medium impact BCS at Control Centers. This should make the requirement focus on specifically what communications need to be detected, but avoids implication of where the technology must reside to do so.

Question 11.

The SDT revised CIP-007 based on industry comments. Do you agree with the proposed changes to the NERC Glossary terms? If not, please provide the basis for your disagreement and an alternate proposal.

Q11 Comment Theme:

- **Requirement 1.2 need clarification on what "non-programmable communications components located inside both a PSP and ESP" includes.**

SDT Response: Thank you for your comments. The SDT asserts that this is a redline error and is from currently approved CIP-007-6 and is not a change made by this SDT. We apologize that it showed as a change in the redline. It covers, for example, "non-programmable" communication components such as a unmanaged network switch that may not meet the definition of "Cyber Asset" but are still the object of this existing requirement.

Q11 Comment Theme:

- **Request clarification on CIP-007 Part 1.1. Why is this Requirement so technology specific? We refer to "Internet Protocol ports" instead of "ports." There is inconsistency between CIP-007 Parts 1.1 and 1.2 where 1.1 uses "Internet Protocol ports" while 1.2 uses "physical input/output ports" and "network connectivity." We expected consistency.**
- **Under Internet Protocol (IP), the transport layers (e.g. TCP, UDP) create and maintain the ports for the underlying services.**

SDT Response: Thank you for your comments. The SDT agrees that logical ports are at the TCP/UDP level and not IP (per RFC 791). Based on the feedback, the SDT has taken a simplified approach to this requirement by stating it as a clearer objective of system hardening – disable or prevent unneeded routable protocol network accessibility on each Applicable System. This intent is that in the virtualization underlay, the hardware resource pools being managed by hypervisors communicating with each other, the entity can manage unneeded network accessibility by disabling underlay services that are not needed. This also allows entities to achieve the objective at either a port or service level; some unneeded accessibility may be disabled/prevented at a TCP/UDP port level, others by disabling a service in the OS. The objective is to reduce the network attack surface of the Applicable System and no longer prescribes that this be accomplished at either a port or a service level.

The intent of the different 1.1 and 1.2 requirement parts remains in that 1.1 is concerned with operating system (OS) level ports and services, while 1.2 is concerned with physical hardware ports, for example, RJ45 ethernet, USB ports, and DB9 serial ports.

Q11 Comment Theme:

- **Requirement 1.3 is greatly limiting. One objective of virtualization is to optimize the usage of computer resources (CPU, memory, etc.) and this limits the possible gain of using virtualization.**

SDT Response:

Thank you for your comments. The SDT is allowing entity choice on the treatment of computer resources by either considering the computing resources and all the hosted VCAs as the same impact level (high watermarked) through the PCA definition (everything that can execute simultaneously on the same CPU/memory becomes an 'associated PCA'). If the entity chooses this path, the computing resources are not SCI and this requirement no longer applies. This requirement only applies if the entity chooses to retain differing impact levels of hosted VCAs and isolates them from each other by configuring affinity rules to keep VCAs of differing impact levels from executing on the same underlying hardware resources.

Q11 Comment Theme:

- **Clarity on Requirement 2 and patching cadence for dormant VCAs.**

SDT Response:

Thank you for your comments. The SDT has added more qualifications to the definition of VCA to clarify that a dormant, non-instantiated VM image is not a VCA, and thus not an Applicable System of R2. Other clarifications have been made in CIP-010 to clearly allow for automated vulnerability assessment and remediation of a VCA upon instantiation in a remediation VLAN.

Q11 Comment Theme:

- **Request explicit language in CIP-007 R2 that these patching Requirements do not include patching in the cloud.**

SDT Response: Thank you for your comments. The SDT does not have authorization in its SAR to address the aspects of off-premise cloud computing, of which this is one of numerous issues that would need to be addressed.

Question 12.

The SDT revised CIP-010 based on industry comments. Do you agree with the proposed changes to the NERC Glossary terms? If not, please provide the basis for your disagreement and an alternate proposal.

Q12 Comment Theme:

- **Overall: Changes are not needed/aren't related to virtualization**

SDT Response:

Thank you for your comments. The SDT asserts changes are needed to CIP-010 for virtualization issues. The SDT is making CIP-010 R1 a requirement with a clear change management security objective to alleviate several virtualization-related issues such as how to treat parent/child images (such as in VDI), application containers, clarifying the use of automated vulnerability assessment capability for VMs with remediation VLANs, and inclusion of authorization of change to the security controls affecting VM affinity and network isolation between VMs of differing impact levels.

Previously, CIP-010 R1 was concerned with five prescriptive attributes of a system, several of which caused issues with the virtualization concepts noted above. Rather than continuing to expand a list of prescriptive attributes in R1.1 with a backwards-looking requirement to update a baseline within 30 calendar days after a change is made and in addition determining affected CIP-005 and CIP-007 security controls, the core of the requirement is now the identification and testing of those CIP-005 and CIP-007 security controls and authorization of changes. It now focuses on the security objective of authorizing changes and allows entities to define those changes, taking into account the virtualization issues, that can affect the CIP-005 and CIP-007 security controls.

As an example of a virtualization issue, the SDT considered VDI scenarios where dozens or hundreds of VCAs may be dynamically instantiated from parent images. The SDT's intent is that the goal of R1 is not to maintain perfect baseline configuration of a hundred individual dynamically created VCAs but to focus on the much higher risk of authorizing and thus controlling change to the parent image from which a hundred dynamic images will be created. A forward looking requirement to authorize a change is of more value to reliability than a requirement to update baseline documentation in a certain timeframe after making a change. Another example is dormant VMs that may rarely instantiate. With automated vulnerability assessment and remediation of patches, AV updates, etc. at instantiation, the goal of R1 is not to have the entity somehow track when this automated process happens to each image file and then prove the update of baseline documentation within 30 calendar days. These are two examples of virtualization enabled scenarios where the SDT asserts CIP-010 needs to change to accommodate these advances in technology in a much more dynamic and automated environment. The SDT will incorporate example measures that include baseline configurations as one tool an entity may choose for this requirement.

Q12 Comment Theme:

- **R1/ Part 1.1.1: Add baselines and parent/child (VDI) examples to the measures as clarifying examples**
- **What level of OS change must be authorized?**
- **Why is lack of OS relevant to patching?**
- **The Measure mentions patch 'implementation' vs the requirement for change authorization**

SDT Response: Thank you for your comments. In response to these and numerous other comments on CIP-010 R1 and to allow for scenarios brought about by virtualization, the SDT has chosen to go with an objective level requirement rather than continue to maintain a prescriptive list of system attributes subject to change management processes. This requirement now requires entities to define the types of changes that will be in the change mgt process that includes changes that may impact CIP-005 and CIP-007 controls. Then, for those changes, the requirement focuses on authorizing those changes and the verification that changes do not adversely impact those security controls. This was the core of the previous version's requirement as well, but is now more future-proof and objective oriented.

Q12 Comment Theme:

- **R1/Part 1.1.2: Clarify "application containers"**

SDT Response: Thank you for your comments. The SDT in our draft 1 posting proposed a new glossary term for "Self-Contained Application (SCA)". Most stakeholder comments said that term is unnecessary and to use the IT industry "application container" term instead of proposing a NERC specific definition. The SDT intent for including this phrase is not to create something new but to simply help clarify that "application containers" as understood in the IT industry are to be treated like other application software and not as some unique form of VCA.

Q12 Comment Theme:

- **R1/Part 1.1.4: Ports and services/clarity on dynamic port ranges/tracking ports in the virtualization underlay rather than services/Clarify the authorization for change is to the OS/SW/Config that affects open ports, not every time a port opens/closes.**

SDT Response: Thank you for your comments. The SDT agrees with many of these comments and based on these and others has, in conjunction with changes to CIP-007 R1, made both of these requirements more objective-oriented. This allows these requirements to incorporate virtualization-specific scenarios and modifies the language in such a way that as technology continues to change, these requirements will not require continual update due to maintaining prescriptive lists of system attributes and baseline config documentation of such attributes – including ports and services. As the network accessibility of a system and reducing the network 'attack surface' remains a CIP-007 R1 security objective, then changes to Applicable Systems that can affect this remain the subject of CIP-010 R1 as well.

Q12 Comment Theme:

- **R3/Part 3.3: the requirement to build outside of the production environment and perform the vulnerability scan has been removed**
- **Clarify timeline between the VA and production**
- **Clarify "prior to becoming an Applicable System"**

SDT Response: Thank you for your comments. The SDT asserts the objective to perform a vulnerability assessment of a new Applicable System before it can have an impact in production remains, but has been reworded to allow for virtualization scenarios and automated means of implementing this requirement. The currently approved language states "Prior to adding a new applicable Cyber Asset to a production environment…" which has worked well with the traditional Cyber Asset such as Cyber Assets that arrive on a loading dock. However, it does not work well with a VCA that is created in an isolated environment WITHIN a virtualized infrastructure. It does not arrive from 'outside', it is created in the environment. The SDT changed the language for clarification. The changes should also clarify that automated tools that implement this functionality, such as remediation VLANs, are a good and allowable practice even though it occurs in a production "environment", but not with production network connectivity and the ability to impact other production VCAs.

The SDT's intent is that "prior to becoming an Applicable System" means that prior to the object in question fulfilling the role of one of the Applicable Systems and their definitions. For example, before it can have the 15 minute impact of the BCA definition, or perform the electronic access control function of an ESP or BCS. At the time the object fulfills those roles and thus meets the corresponding definition, that's when it "becomes an Applicable System" of this requirement, and prior to that point in time the vulnerability assessment must have been performed.

The SDT notes that the maximum allowable timeframe between the vulnerability assessment and the production use of the system has not been prescribed to date and is not a virtualization related issue. The SDT is changing the language to clearly allow for automated tools that will perform this task faster and every time a VCA is booted, not just its initial introduction to the environment.

Q12 Comment Theme:

- **R4: TCA scope changes from High/Med and PCAs to "everything but lows" and includes all non-CIP and thus the scoping is entirely in the definitions of TCA/RM**

SDT Response: Thank you for your comments. The SDT agrees and has reinstated updated applicability statements within the CIP-010 R4 requirement language and removed the applicability statements from the previous draft that were in Attachment 1.

Question 13.

The SDT revised CIP-003, CIP-004, CIP-006, CIP-008, CIP-009, CIP-011, and CIP-013 (conforming changes) based on industry comments. Do you agree with the proposed changes to these Reliability Standards? If not, please provide the basis for your disagreement and an alternate proposal.

Q13 Comment Theme:

- **CIP-006 R1.10 was moved to CIP-005 R1.4 but the associated PCAs aren't included**

SDT Response: Thank you for your comments. The SDT agrees and has included associated PCA's in the scope of the requirement.

Q13 Comment Theme:

- **CIP-009 should not include SCI in applicability. Recovery of the functionality may occur outside of recovering underlying SCI.**

SDT Response: Thank you for your comments. The SDT agrees that the SCI itself does not always require its own recovery plan. The objective of this requirement is not to recover SCI simply as SCI, but to recover the functionality of the systems it supports that remain in scope. Entities could have recovery plans that recover the functionality in other ways if underlying SCI is the issue. However, entities should be aware that if their plan to recover an Applicable System is dependent upon recovery of the SCI, the SCI should be a part of the documented plan. With the exception of Part 1.5, the SDT removed SCI from the Applicable Systems where it had been included before in Draft 2.

Question 14.

Please provide any additional comments for the SAR drafting team to consider, if desired.

Q14 Comment Theme:

- **Associated Data Center (as required by the latest ERT v5 spreadsheet) is not listed in the draft standard**

SDT Response: Thank you for your comments. The 'associated Data Center' is an included component of the currently approved Control Center definition in the NERC glossary. It is not the subject of any changes by this SDT.

Q14 Comment Theme:

- **Planned and Unplanned changes language should be added to the Implementation Plan, similar to the language included in previous CIP implementation plans.**

SDT Response: Thank you for your comments. The SDT agrees and has copied those sections into the proposed Implementation Plan so those provisions remain with this version.

Q14 Comment Theme:

- **CIP-013 is part of this update, is there a reason why CIP-014 is not part of this update?**

SDT Response: Thank you for your comments. CIP-014 is a physical security standard and is not in the suite of cyber security standards, unlike CIP-006 that is specifically related to the physical security of BES Cyber Systems. While CIP-014 has a few similarities with CIP-002, it is not affected by virtualization technology. CIP-013 on the other hand, is in the suite of Cyber Security standards and is affected by virtualization – such as the procurement of SCI.