NERC
NORTH AMERICAN ELECTRIC RELIABILITY CORPORATION

Virtualization and Future Technologies: Case for Change Webinar Q & A

Project 2016-02 Modifications to CIP Standards

The Project 2016-02 Modifications to CIP Standards Drafting Team (SDT) held a Virtualization: Case for Change Webinar on April 24, 2019. The following items are questions and answers from the webinar. The SDT encourages industry to review this document prior to responding to the Case for Change Informal Comment period.

Q1. The White Paper states "Here it is clear the physical server hardware and the hypervisor are in scope and each VM is subject to requirements such as CIP-007" (page 3). Is it the SDT's position that this is true for hypervisors ONLY containing EACMS?

A1. No, the Cyber Asset definition is used for all asset classes including EACMS, BCA, PCA, etc. and it is inclusive of hardware. What the SDT is pointing out is that:

- If the Cyber Asset definition is inclusive of hardware and;
- If the entity classifies their VM as a Cyber Asset; then the hardware is included in that definition.

Q2. With the Zero Trust Model, now it looks like it would allow the co-located systems without a separate EAP (gated) around the ESP, and allow VM's to be an EAP/ESP individually and would not be requiring physical segmentation from NON-Compliance systems.

A2. Zero trust is a methodology for implementing access control and integrity. In principle the services on top in this model would not even trust their own underlying architecture. We will talk later in the presentation about the SDT's strategy for handling mixed-mode or mixed trust systems through the use of new requirements targeted at new definitions.

Q3. Since a Cyber System can be one CA, doesn't the model on the right just provide defense in depth, like physical security?

A3. An ESP within an ESP? The SDT does not believe the standard should prescribe a particular technology solution to these issues. In addition, within the standard the entity's network access control is only measured at the network edge. We believe the entity does not get credit for the defense in depth strategy, so it does not encourage them to use it.

Q4. As you discussed the BCSI issues, how will this team work with the separate SDT working on BCSI?

A4. Both CIP SDTs will keep open collaboration regarding BCSI through the development of each project.

Q5. The Case for change paper speaks to Distributed Firewalls as a more secure option in a virtualized involvement. Why can't this approach be effectively used in today's environment either behind an identified EAP or have the large ESP broken down into multiple ESPs?

A5. The SDT does not believe the standard should prescribe a particular technology solution to these issues. Within the standard the entity is only measured at the edge and we believe the entity does not get credit for the defense in depth strategy so it does not encourage them to use it.

Q6. To clarify: if a hypervisor hosts EACMS and non-EACMS, is the hypervisor an EACMS? Are the other non-EACMS VMs high watermarked? We've been told EACMS do not high watermark like BCAs.

A6. The SDT believes that how to handle the described scenario is unclear. This is the primary reason that it was presented in the challenges section of this presentation. EACMS relies on the CA definition and the CA definition includes hardware.

Q7. In regards to VLANs on a switch and port security. If I am using a trunk port connecting to a firewall, would this make the switch an EACMS? When any type of port security is used, i.e. mac security or 802.1x ... would this also make the switch an EACMS?

A7. The SDT believes that this is also an area that needs clarity. Through our analysis the switch itself performing the access control could fall into the EACMS category. However, if it has a 15 minute impact it may also fall into the classification of a BCA.

Q8. Why is the SDT working on still addressing hardware instead of addressing the Virtual environment risk (using the highest risk application principle)? SDT is not the first to draft virtual security, PCI council and NIST already have a risk based approach.

A8. The SDT believes that it is not simply addressing the hardware but that the existing definitions are limiting our ability to describe virtual environments. We have been looking to other frameworks as guides to effectively describe these new technologies while maximizing our ability to be backward compatible with existing programs.

Q9. BCSI explanation was very confusing. If SAN is part of the environment, why should not the storage area be the focus?

A9. The SDT is addressing the storage architecture as a core area of virtualization as you will see on many of our drawings. While changes may be needed, we are not directly addressing the handling of BCSI at this time because there is a separate team evaluating the topic.

Q10. For Entities that are moving toward virtualized environments, will the SDT provide some guidance for those moving to it prior to the finalization of the revised standard? Alternatively, will there be a pilot implementation group similar to the V5 TAG pilot?

A10. The SDT plans to provide as much clarification as possible through guidance, outreach, and white papers. Many entities are already using virtualization in their environments and we will try to work with NERC to help provide a smooth transition.

Q11. What is SCI?

A11. Shared Cyber Infrastructure. It is a new definition to target requirements at the shared infrastructure on which virtual objects (machines, networks, and storage) execute (As depicted in grey in the drawings).

Q12. Does the proposed approach for the "shared infrastructure" categorization create a dual standard for compliance between data center (compliance at SCI) and substation environments (compliance at BCA level with ports and services and passwords, etc.)?

A12. The SDT plans to use the new term for Shared Cyber Infrastructure to target applicability within the standards.

Q13. From slide 27, how does the approach to 'deny all access to and from networks, with "networks" being the operative term, permit the use of a zero trust model. Doesn't this still prescribe the gated community model and preclude 'house'-level protections?

A13. The language on the screen was draft and we realize that it may not be the correct target for this requirement. We did want to share this approach with you for your consideration as we continue to draft better requirements.

Q14. How do you address the problem of VLAN hopping/crossover? Do you make a distinction between port based and logic based VLANs?

A14. To date, the SDT has been using the concept of logical isolation that can be created by various means including those you mentioned. The entity would be responsible to prove that their logical isolation is sufficient to meet the requirements.

Q15. Can you clarify the segregation requirements of shared storage? I understand the need for separate hosts for CIP and non-CIP assets - does that mean we need separate storage arrays or will logical access controls that segregate LUNs to hosts suffice?

A15. As we have been describing it so far, we would need to draft a new requirement targeted at the shared infrastructure to create logical separation. Those requirements are not yet complete.

Q16. Are there any lessons learned that will be delivered soon about those challenges?

A16. The SDT has been working on a model we call Pinecone Power, a fictitious entity that will be the foundation for guidance and future lessons learned whitepapers.