Virtualization and Future Technologies

Project 2016-02 Standards Drafting Team:
The Case for Change

April 2019

Table of Contents

Executive Summary

What the interconnected power grid does for Bulk Electric System (BES) reliability, virtualization does for the computing infrastructure supporting vital control systems. Individual utilities interconnected their power systems to form a power grid to share spare capacity for meeting demand peaks and surviving contingencies such as generating unit and transmission line outages. Virtualization connects processors, networks, and storage into 'computing grids' that allow our vital systems and applications to meet peak demands and survive outages of individual components.

This is accomplished by abstracting servers, networks, and storage into virtual or logical resources that can be independent of specific underlying hardware such as individual processors, circuits, and disks. A control system and its underlying operating system become a virtual machine and can move to any available hardware. This greatly increases reliability and resiliency of our control systems that support BES reliability.

Virtualization technologies also allow enhanced cyber security controls and the ability to move access controls from the edge of our networks to much deeper inside of them. This is analogous to having generation close to load centers to reduce the susceptibility to outages. These newer security controls allow us to provide tighter security by moving access controls from an outer perimeter closer to the actual code performing reliability tasks.

Virtualization and advanced technology are challenging the way we characterize the Critical Infrastructure Protection (CIP) standards' objectives and how we develop technical requirements. Use of virtualization and advanced technology can provide benefits for implementing both operational and security enhancements to a system. The goal is to require technology-enforced controls as alternatives to the current prescriptive requirements like those requiring a physically structured architecture, without forcing the use of the new technology. The existing standards with their prescriptive language limit the ability to take full advantage of the new technologies. The Project 2016-02 Modifications to CIP Standards Standard Drafting Team (SDT) is drafting new requirements to support virtualization capabilities. This leaves Responsible Entities with the option to maintain a non-virtualized environment and use backward compatibility to preserve current CIP investments and security postures.

This white paper represents the views of the SDT and presents a sampling of virtualization concepts. It explains, in the SDT's opinion, how the use of these concepts is inhibited by some of the CIP standards' definitions and requirements. It also introduces the SDT's ideas on how to address these issues while maintaining compatibility with current state. The goal is to allow for the use of these technologies and newer security controls by removing prescribed "how's" and replacing them with clear "what's" that would allow both current state and these enhanced features.

Introduction

Note: The Project 2016-02 SDT developed this white paper to explain the need to change the cyber security CIP Reliability Standards. This white paper has not been approved or endorsed by NERC and is solely the views of the Project 2016-02 SDT.

In the history of the CIP standards, industry has seen many versions and changes. Some have been straightforward and almost self-explanatory. Others, complex and time-consuming to implement. Still others have been foundational, like the ‘do-over’ change from Versions 1 – 3 to Version 5. Not surprisingly, the prospect of another set of fairly major changes, this time involving virtualization sparks a great deal of industry concern. It raises valid questions about the timing and the drivers of those changes, and whether or not they are truly necessary. The white paper is designed to provide detailed answers to those questions from the viewpoint of the SDT. This introductory section provides a short, high level introduction to the white paper discussion on what brought about the SDT’s focus on future revisions to the CIP standards to address virtualization.

Recognizing the continuing growth in technology innovation, many entities in the Electricity Sector have implemented virtualization as part of their CIP programs. Many of these same entities, however, have implemented this new technology without taking full advantage of virtualization’s advanced capabilities. There are a number of reasons for this from the constraints of the current CIP architecture to the ongoing ambiguity around how new virtualization technology applies to CIP compliance. Some of those who are implementing virtualization are experiencing a great deal of uncertainty and difficulty around developing implementation strategies that will support compliance and achieve greater reliability and security.

These issues began coming to light following the NERC “Virtualization Summit” in 2015. Ultimately, the Version 5 Transition Advisory Group (V5TAG) heard about the industry concerns and determined that the issues around the standards and virtualization would be best addressed by a drafting team. The Project 2016-02 SDT was assigned the task to address the technological innovation in virtualization within the CIP standards.

The SDT’s purpose of incorporating the virtualization concept into the CIP standards is not to merely augment the current standards. The SDT’s intent is to better position the CIP standards to be applicable to any future technological innovation. Leveraging the abstraction that virtualization provides will allow the industry to more readily adopt new technology and increase security posture. This paper presents the SDT’s case for change to the NERC CIP standards that is needed to allow for the innovative security techniques and new concepts brought about by virtualization.

As virtualization has progressed, many of these types of issues cannot be addressed with Implementation Guidance. Documenting a possible way to implement a requirement is of great value, but Implementation Guidance cannot, for example, change current requirements so they do not prescribe a perimeter-based model, or allow remediation VLANs. It also cannot add new requirements that are needed for issues like management plane isolation or handling shared infrastructure. These new concepts and techniques require changes to the standards to make them viable and to clarify how they should be secured.

Chapter 1: Virtualization Benefits

The most basic concepts of CIP are essentially unchanged from the Urgent Action 1200 standard in 2003. The primary focus of those standards was the “critical cyber asset”; an “electronic device” such as a server, workstation, or relay as a physical object. It had an operating system, always on and performing its function, and communicating with other components over routable protocols. It was protected by traditional firewalls at the network edge looking at source and destination protocol addresses and ports as the only mechanism by which to make network access control decisions.

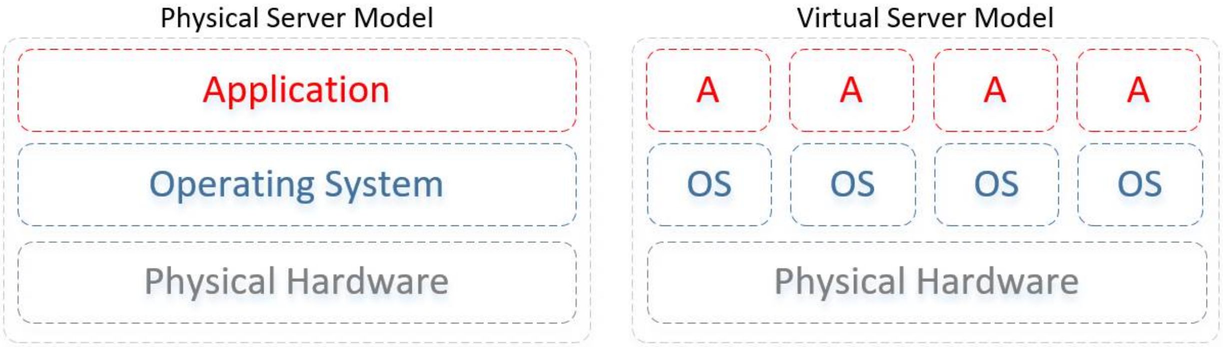

Figure 1

Today, virtualization has changed this scenario. With virtualization, physical devices are no longer the primary units of organization. An entire control system infrastructure can be virtualized (such as “software defined data centers”) and only exist as logical constructs. An EMS database server may never exist as a discrete physical object. With containers (operating system emulation), there may not be a concrete tie between application logic and an operating system. Virtual machines can be created and destroyed dynamically and are neither always on, nor tied to specific hardware. Workloads may mirror their information for reliability purposes across great distances without using routable protocols. With micro-segmentation, network access control lists are now much more granular than IP addresses at a perimeter. They are enhanced by policy-based control templates enforcing access at a “user to workload” level throughout the system infrastructure. Electronic access control may no longer be based solely on routable protocol addresses or found only at an ESP boundary.

Benefits of Virtualization

Virtualization technologies bring many reliability benefits to BES Cyber Systems (BCS).

- Increased uptime, very fast recovery capability and flexible architecture that can instantly adapt to changing workloads.
- Virtualization allows for racks of CPUs, memory and disks to be tied together with high speed mesh networks and viewed simply as raw computing resources.
- If the workload on a particular virtual server is nearing capacity, the infrastructure orchestration system can create and configure an additional server on the fly, bring it online to help with the peak workload, and then destroy it when it is no longer needed.

- If a physical machine runs out of resources, the workload can be moved to another physical machine dynamically based upon relative load. When a virtual server or workstation is not in use, it is similar to a physical server that is powered off.
- This flexible and dynamic architecture also allows improved security controls such as those provided by micro-segmentation.
- Users can be granted access to specific workloads that can be placed dynamically throughout the infrastructure with managed access to provided services.

More secure, more reliable, cost effective technology solutions

New and innovative security capabilities continue to evolve to address current threats, particularly in virtualized environments. These capabilities can increase the security of hosted BES Cyber Systems. For example, *distributed firewalls* based on software defined access policies can enforce access controls at a much deeper level within the infrastructure to help prevent an attacker's lateral movement through the network. *Privileged introspection* allows security service such as anti-malware, to operate in a tamper-proof way outside of the instances they protect. *Zero trust models* allow communication to be protected end to end between individual processes across cyber assets without having to trust that devices in the path (such as all firewalls and switches) are configured correctly to protect the data. We'll consider each of these and other examples in more detail to show the benefits and the case for change in the CIP standards.

Chapter 2: Virtualization challenges in the NERC CIP Standards

Virtualization brings benefits to reliability and resiliency of our BES Cyber Systems. It also brings challenges with determining how some of the newer concepts fit within the framework of the NERC CIP standards. Some of these newer techniques and concepts are:

- Identification of Virtual Cyber Assets
- Distributed Firewalls vs. Perimeter models
- Zero Trust models
- Virtualized Firewall Interfaces
- Virtual Storage challenges
- Management Plane Isolation
- Privileged Introspection
- Remediation VLANs

Each of these topics will be covered in detail in the following sections.

Identification of Virtual Cyber Assets

One of the issues requiring change within the CIP standards is the need to clarify the treatment of virtual machines (VMs) under today's definitions. The foundational term Cyber Asset is defined as "Programmable electronic devices, including the hardware, software, and data in those devices." This definition, since it includes the hardware, does not fit with virtualized environments. A literal use of this foundational definition would mean a server hardware platform may be a BES Cyber System, the hypervisor software could be the "operating system", and all the virtual machines running on the hardware are applications or simply data. The CIP-007 requirements, most of which are aimed at an operating system instance, would not be applied to these virtual machines since they aren't hardware. This is referred to in the Figure 2 as the "Physical System w/Software Model."

To require that VM operating systems be classified correctly, the tie between a virtual machine and the hardware needs to be clarified. The standards and their associated definitions need to clearly support the "VM as a CA Model" (Virtual Machine as a Cyber Asset) as pictured in Figure 2. Here it is clear that the physical server hardware and the hypervisor are in scope and each VM is subject to requirements such as CIP-007.

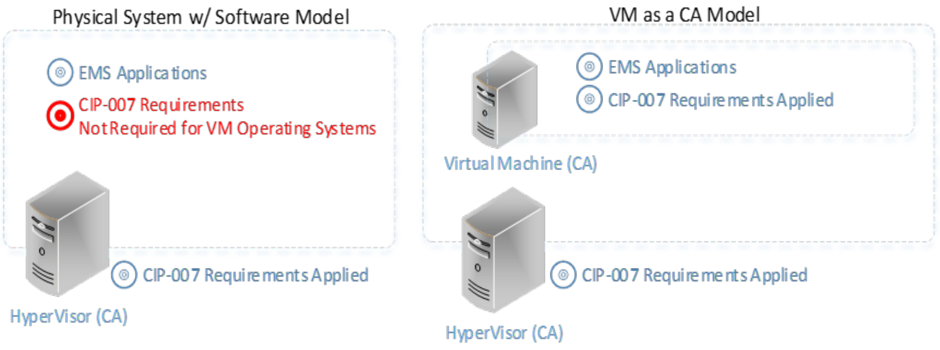
Physical System w/ Software Model
EMS Applications
CIP-007 Requirements
Not Required for VM Operating Systems
CIP-007 Requirements Applied
HyperVisor (CA)
VM as a CA Model
EMS Applications
CIP-007 Requirements Applied
Virtual Machine (CA)
CIP-007 Requirements Applied
HyperVisor (CA)

Figure 2

The Cyber Asset definition assumes a 1:1 relationship between a BES Cyber Asset and its hardware. Virtualization, however, breaks that concept because in some cases hundreds of virtual BES Cyber Assets can share a pool of hardware resources (compute, network, and storage). The outdated concepts need to be changed to recognize the one-to-many relationships that are now possible between a BES Cyber Asset, an instance of an operating system supporting applications, and the underlying hardware so it is clear how each is to be treated. There are three different types of asset classes involved in today's BES Cyber Systems:

- Self-contained devices that are composed of dedicated hardware, some form of operating system or firmware, and the application code. This is the traditional definition of a Cyber Asset and applies to things such as digital relays, RTUs, physical operator workstations and dedicated physical servers.
- Virtual "cyber assets" composed of an operating system and applications or containers minus any dedicated hardware. These types of assets are logical or virtual constructs by nature and exist only in memory or files. They can, however, appear from a network perspective the same as any other host.
- Shared Infrastructure consists of the hardware resource pools (compute, network, storage) and is shared by virtual cyber assets and can host numerous virtual cyber assets, networks, or storage locations.

The current CIP standards only recognize the first class and changes are needed to properly address the other two. This will allow BES Cyber Systems, which can be composed of all three asset classes, to be properly identified and classified and the proper requirements applied to each one.

Distributed Firewalls vs. Perimeter Models

Distributed firewalls based on software defined access policies can enforce access controls deeply within the infrastructure to help prevent an attacker's lateral movement through the network. The CIP standards (CIP-005) today require a perimeter-based model as shown in the left side of Figure 3. It has an Electronic Security Perimeter (ESP) as a logical border around a group of cyber assets, and defined interfaces (on a firewall for example) as an Electronic Access Point (EAP). This perimeter model is a prescriptive topology, and for many scenarios is still a valid way to perform network security. What may be prescribed, however, for a small network of similar cyber assets may not be ideal for a large network of virtualized BES Cyber Assets.

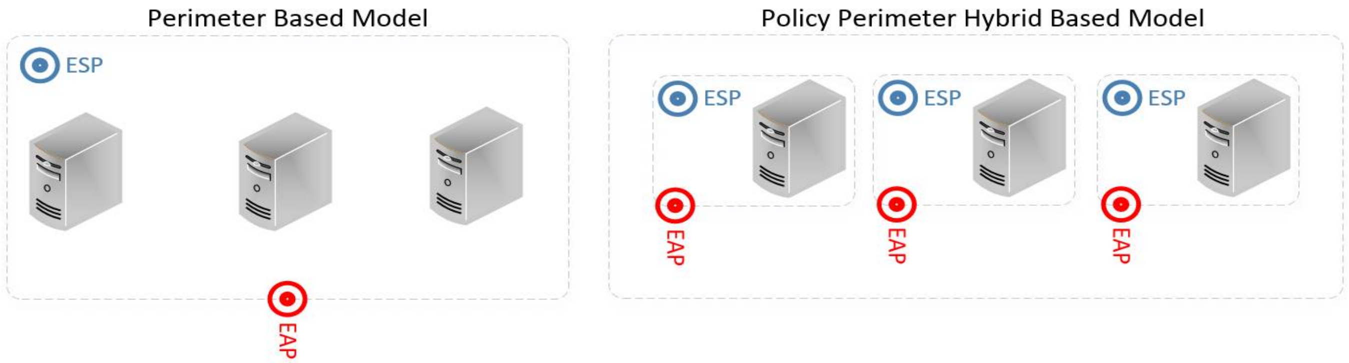

Figure 3

The issue is that network access is permitted or denied at the outer EAP "gate." Even if access is only needed to a single cyber asset inside the perimeter and a strict rule is instituted at the EAP "gate" to allow someone access to that single cyber asset, that control is implemented only at the gate. Once inside the perimeter, there are no network level access controls that would prevent hopping ("pivoting") from that cyber asset to all other cyber assets inside the perimeter. Attackers today depend upon this; if they can get a foothold on a single cyber asset inside a perimeter they can work over time to move laterally within the environment from there.

Virtualized environments, especially those of cloud providers that support tenants from various customers, had to be designed to implement new, more granular security techniques to not only mitigate this threat but to also keep customer's workloads totally isolated from one another. These new techniques are philosophically simple – every cyber asset is inside its own "perimeter" and the EAP "gates" are specific to each virtual cyber asset. This is shown on the right side of Figure 3 in the "Policy Perimeter Hybrid Based Model." In this model, network access is controlled at an individual cyber asset level and access to a single cyber asset gives you no additional access to any other cyber asset. An attacker's ability to move laterally or pivot in this environment is greatly reduced. These techniques are also configured at a *network access policy* level and the infrastructure dynamically implements it at various levels throughout the entire infrastructure. Even as these cyber assets move dynamically within the infrastructure, these access controls move with them. All of this, however, precludes the ability of administrators to provide a "list" of potentially hundreds or thousands of ESPs, EAPs, or the discrete rules per EAP. This is because it is all dynamically generated by the infrastructure to control access in accordance with a higher-level access policy.

These same techniques can be used to mitigate the same risks within virtualized CIP environments. The problem is that CIP requires the perimeter model with EAP "gates" as the only prescribed way to implement network level security. The SDT is planning to change the CIP standards and definitions in a way that does not *preclude* the perimeter based model, but also does not t *prescribe* that model as the only way. For example, the language in CIP-005 R1.2 "All External Routable Connectivity must be through an identified Electronic Access Point (EAP)." This creates an issue since historically this is a physical interface on a physical device. More current, virtualized architectures distribute the policy enforcement tasks to multiple devices and each of them validates that the network communications passing through it are permitted. The current standards do not align with these architectures because there is no single interface identified that is responsible for policy enforcement.

Identifying non-routable EAPs and the OSI Model

Related to the ESP issue, today's CIP Standards are limited to prescriptive topologies even as potentially more secure, reliable and cost-effective solutions are available. ESP's are defined today by the access provided at the network layer (OSI Layer 3) and are therefore limited to making access decisions based on routable protocol addresses. This method does not support security evaluation or compliance via security solutions at any of the other layers of the OSI model. This presents two issues:

- There may be network access into and out of the perimeter at different layers other than by a purely routable protocol that allow for things like high speed replication of data. This will be discussed further in the "Multi-Site Data Center Extensions (Super ESP)" section below.
- Virtualized environments have much more context and can enforce network access control at more granular levels in much better ways. These include by user, process, or certificate and are not limited to only a source/destination IP address of a routable protocol.

In many environments today, the perimeter model is sufficient and routable protocol addresses are all an entity has to make access control decisions. However, the standards should not prescribe this as the only way and should allow entities to use these more granular controls that may operate at other layers.

Zero Trust Models

The zero trust network model is a new and different way of thinking about network security. In most cases it can be implemented within the security model that an entity currently has deployed.

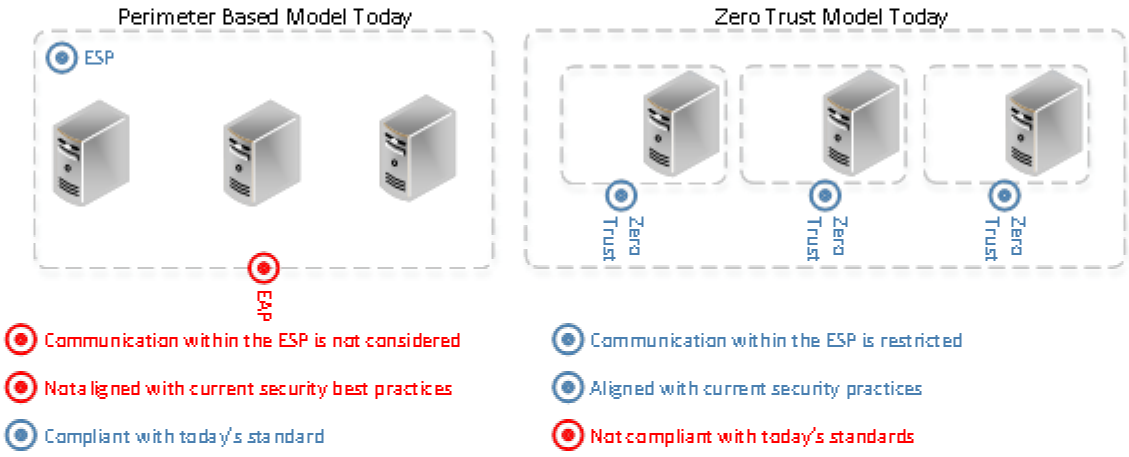
Perimeter Based Model Today
ESP
EAP
Zero Trust Model Today
Zero Trust
Zero Trust
Zero Trust
Communication within the ESP is not considered
Nota ligned with current security best practices
Compliant with today's standard
Communication within the ESP is restricted
Aligned with current security practices
Not compliant with today's standards

Figure 4

Most entities have a perimeter model deployed, which is what the current ESP model requires as in Figure 4. Traditionally, this consists of a firewall (or firewalls), with a single path into the network (EAP). This model protects the BES Cyber System indirectly, by protecting the network topology. This access point is protected by rules and policies that tell the firewall what kind of traffic to pass and what kind of traffic to stop. This model can be loosely compared to a gated community, where the fence is the firewall, and the gate and security guard are the access control.

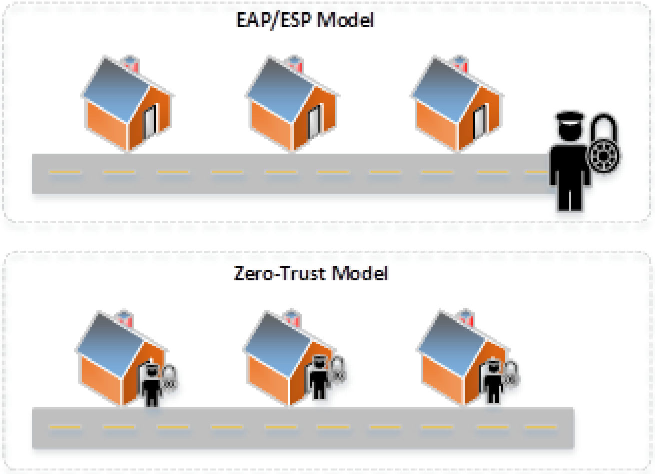

Figure 5

The overarching problem with this model is that if attackers should get past the firewalls and through the gate, they would have free rein throughout the network and could pose a serious threat to other BES Cyber Systems within the ESP.

A zero trust model brings the protection directly to the cyber asset or even to the process level. Using network policies, whitelisting, certificates, and other technologies, the ability of an intruder to move around is severely restricted. A zero trust model usually starts with simple items like restricting network traffic to the device except that which is necessary for it to perform its job function. It also restricts what logins can access the device. Finally, it introduces more complex ideas such as digitally signed executables, restrictions on what time a process can be started, and end to end encrypted traffic, even within the local network. This model protects the BES Cyber System directly by protecting the workloads, the devices and the topology. Unfortunately, this model is very difficult to describe using the current language constructs of the CIP requirements.

Referring to the gated community analogy in Figure 5, this would put a fence in every yard, a lock on every door and an alarm system in every house. This effectively isolates the intruder to the public areas of the neighborhood.

Virtualized Firewall Interfaces ('Firewall on a Stick')

As the use of network segmentation grows and firewalls are used to control traffic flows between the network segments, firewalls need higher densities of interfaces. To obtain these higher densities, products today rely on fewer physical ports and instead use virtualized network interfaces to emulate many physical interfaces. Some modern firewalls are service modules (a "firewall on a stick") installed in a switch or router where virtual interfaces are the only option to route traffic to and from the firewall as shown in Figure 6.

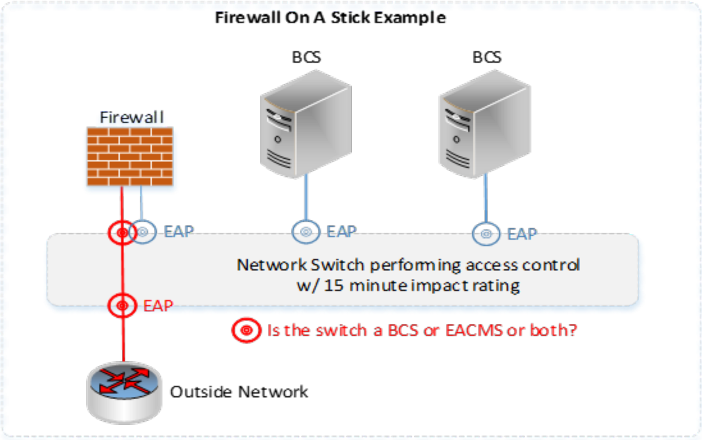
Firewall On A Stick Example
BCS
BCS
Firewall
EAP
EAP
EAP
Network Switch performing access control
w/ 15 minute impact rating
EAP
Is the switch a BCS or EACMS or both?
Outside Network

Figure 6

The CIP standards are ambiguous on how to properly identify these assets. If this were simply a network switch allowing two BES Cyber Systems to communicate with each other with a 15 minute impact if the switch is unavailable, then the switch itself is part of a BES Cyber System. In this example, the entity wanted to increase security of the communications between these two systems. To accomplish this, it put each BCS on its own network segment so it can filter all traffic between these two systems through a firewall. A "firewall on a stick" is installed in the switch and the ports on the switch are divided into virtual networks so the traffic between different segments is directed to the firewall and controlled. Since the firewall either has no physical interface or may use its few physical interfaces for redundancy, the outside network must connect to the same switch and that port is configured to the "outside" interface of the firewall. All communications to and from the outside network and between the two BCS are all controlled via the firewall. The entity has increased the security of these systems, but it is now undecided on how to identify the cyber assets in this scenario. As one example, are all the network switch ports now EAPs on an EACMS, or is this switch still just a part of the BCS? As access control is pushed deeper into networks as in this example, the current paradigms are ambiguous. The answer may lie in identifying new asset classes and developing new objective requirements that do not include prescriptive network topology assumptions. Trunks and sub-interfaces are common network technologies to aggregate data streams and better utilize limited physical interfaces[1].

[1] Trunking is the use of a single link to aggregate multiple data streams. This technology is widely used in networking and telecommunications. Sub-interfaces allow for multiple LAN segments to share a common physical interface.

Virtual Storage Challenges

As we move to virtual storage, there are issues with some of today's CIP concepts. For example, the CIP-011-2, requirement R2 restricts unauthorized access to BES Cyber System Information on systems meant for reuse outside of the CIP environment or systems meant for disposal. This causes a challenge for virtualized infrastructure. Today, when disposing of a physical medium or high impact BES Cyber Asset, access to BCSI can be restricted by sanitizing or destroying the media associated with the physical asset. If storage media is provided from a Storage Area Network (SAN) or Network Attached Storage (NAS) array, there may be no easy way to identify the specific drives where the data might have been stored. Additionally, there is no significant benefit to pulling drives from an array to sanitize or destroy them if there is another way to manage this media.

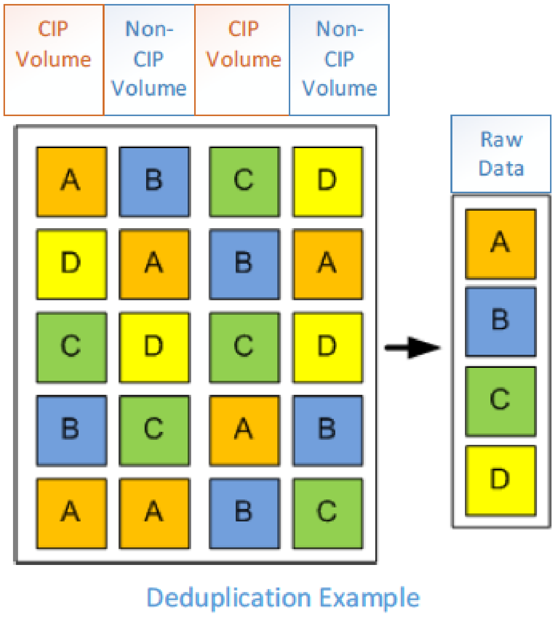

Figure 7

To extend this challenge further, many SAN and NAS environments can de-duplicate and move data to various storage tiers within the array to optimize access times and reduce latency. This process analyzes the data in small enough chunks to create single instances of data that are used within multiple volumes that may be assigned to any number of devices. The chunks with the highest access rates are then placed on fast storage media because the data is constantly referenced. This effectively causes data re-use if storage media is shared between CIP and Non-CIP assets even when the information contained within these de-duplicated chunks is not readily retrievable individually.

That was one example. The CIP standards were built on a paradigm of disk storage being physical disks dedicated to a Cyber Asset. Updates to the standards are needed to reflect the realities of virtual storage technologies that are in use today.

Management Plane Isolation

Because virtualized servers, networks, switches, firewalls, and storage are logical constructs, controlling access and communications to the management plane of these systems is imperative. Access to the management plane (interface/console/etc.) allows a user to create, modify, or delete these objects or entire infrastructures from one place, or move objects from one zone or network to another. Administrative level or "management plane" access to the hypervisors is therefore absolutely critical to the security and reliability of the hosted systems. These types of access must be brought into the scope of CIP standards if hosting BES Cyber Systems and will require changes to the CIP standards.

Another challenge of the perimeter model is it can drive less secure topologies in order to provide clear adherence to the perimeter model. In Figure 8 below on the right is the desired separation between the management plane of a virtualized environment and the production or "data plane". The management port of the BCA is connected to a switch on the management network. The production network switch is connected to the network interface on the BCA. This creates two different network paths to the BCA: one for normal production traffic and a separate network path with separate access controls for administrative or management plane access. It is unclear, however, if the BCA has now created two ESP's with two EAPs, and whether or not it is now an EACMS. It is not functioning as an EACMS, but from the topology it appears as one. The prescriptive CIP model may create a temptation to place the management interface into the production ESP as in the picture on the left. This is very clear from a CIP standard adherence perspective but is the less secure choice.

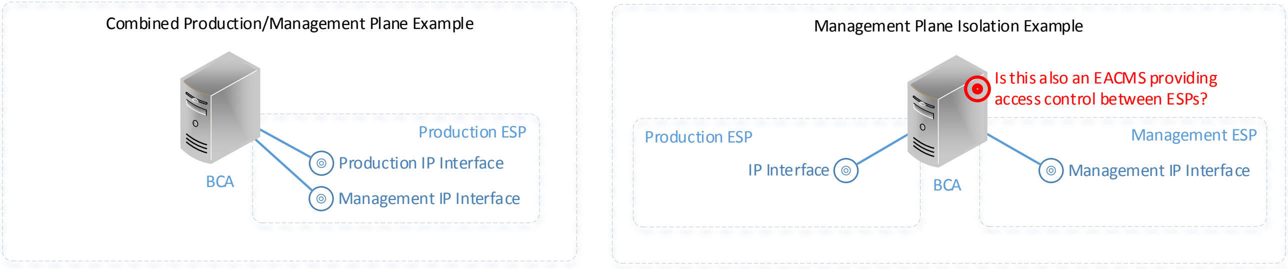
Combined Production/Management Plane Example
Production ESP
Production IP Interface
BCA
Management IP Interface
Management Plane Isolation Example
Is this also an EACMS providing access control between ESPs?
Production ESP
IP Interface
BCA
Management ESP
Management IP Interface

Figure 8

Privileged Introspection

Privileged introspection allows security services, such as malware detection, to be performed outside of the operating system and applications it is protecting in a tamper-proof way. Previously, malware that could gain elevated privileges on a cyber asset could then disable security services such as anti-virus solutions running on the same OS. Virtualized environments, however, can allow privileged introspection where anti-virus or application whitelisting services run as part of the hypervisor; outside of the operating system they are protecting so malware with elevated privileges inside the OS cannot reach them. Another benefit of privileged introspection is it allows integration with distributed firewalls because it can inspect network traffic entering or leaving virtual machines and make decisions with vastly more context. For example, these firewalls can control access not simply at a network address level (IP source/destination) but down to the level of a process with a particular certificate.

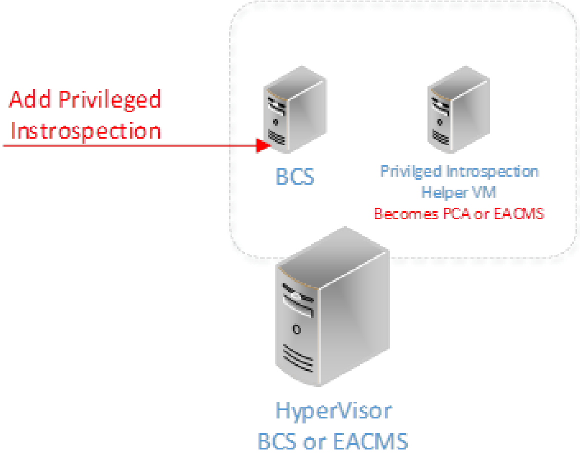
Add Privileged Instrospection
BCS
Privilged Introspection Helper VM
Becomes PCA or EACMS
HyperVisor
BCS or EACMS

Figure 9

Even though the current CIP-007 R3 is already written at a security objective level and can work with privileged introspection, there can still be CIP standards issues brought into play. If this functionality uses "helper VMs", it is not clear how to identify and classify these under the current CIP standards. Since the status of malware detection is now outside the VM's, it requires grouping the VM's protected by privileged introspection into a particular BES Cyber System.

Remediation VLANs

As another example of the need for change in the CIP standards brought about by new virtualization technologies, CIP-010 R3.3 states, “Prior to adding a new applicable Cyber Asset to a production environment, perform an active vulnerability assessment of the new Cyber Asset...” This is a valid requirement that cyber assets be in a CIP compliant state before being added to a production CIP environment. The issue is that this is written with the idea that an entity’s new cyber asset is a separate physical asset it can configure in a compliant state prior to it being placed in a “production environment” and connected to a production network.

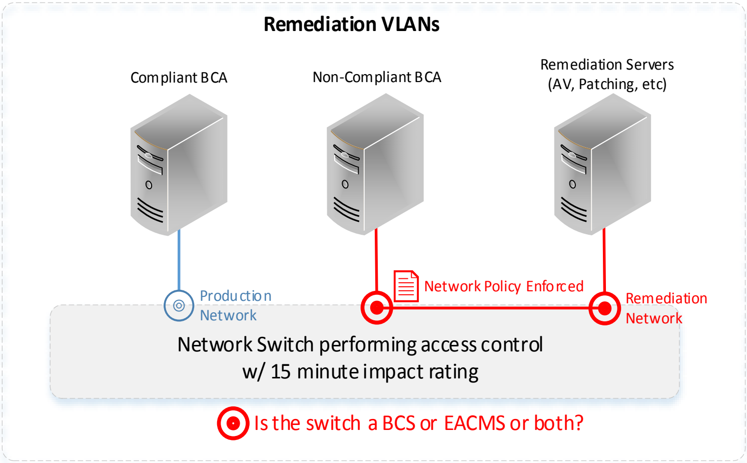
Remediation VLANs
Compliant BCA
Non-Compliant BCA
Remediation Servers (AV, Patching, etc)
Production Network
Network Policy Enforced
Remediation Network
Network Switch performing access control w/ 15 minute impact rating
Is the switch a BCS or EACMS or both?

Figure 10

Virtualized environments, however, allow for a concept known as remediation VLANs. When a new cyber asset, such as a server or operator workstation is instantiated from an image, it is already “in” a production environment. However, it is brought up in a remediation VLAN where its network visibility is restricted to those services (patching, anti-virus updates, etc.) that it needs to assess vulnerabilities, update itself, and become compliant. Only after it is fully compliant with policy, can it be moved to a production VLAN where it has the full network visibility. Should this virtual machine become non-compliant with the network policy, it is moved back to the remediation VLAN to be made compliant again. This is an advanced security control, beyond what CIP-010 envisions, where real time checks are performed to ensure a cyber asset is compliant with security policy. The way the CIP standards are currently written, however, could preclude such advanced controls since this Remediation VLAN must be either part of an ESP or a separate ESP in order to host BCAs per CIP-005-5 R1. Requirements such as those in CIP-005 and CIP-010 need to be modified to allow for these more advanced features and controls.

Multi-Site Data Center Extensions (Super ESP)

There are technology solutions now that support very high levels of resiliency, so much so that virtual machines can be seamlessly moved from physical infrastructure at one data center across great distances to physical infrastructure at another data center. For example, control center functions could move easily and seamlessly from a primary control center to a backup control center across town or across the state. For this to work, an entity would use tunneling protocols that make the two data center networks appear as one local network (“layer 2 adjacency”). You’ve “stretched” a LAN across a WAN. The issue with CIP is that it assumes a LAN at each site with a defined ESP and access control at a point on that ESP at each site, with all the WAN communications equipment in between these ESPs exempted from scope. In

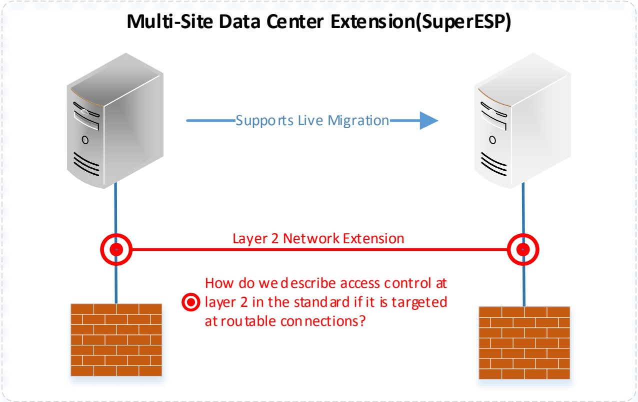
Multi-Site Data Center Extension(SuperESP)
Supports Live Migration
Layer 2 Network Extension
How do we describe access control at layer 2 in the standard if it is targeted at routable connections?

Figure 11

this situation, however, we have a single “local network” stretched between different sites creating at one level a “Super ESP”. Complicating the issue is that underneath this logical network extension, the data may actually ride in encrypted tunnels created at a routable protocol layer. At that layer, firewalls on an ESP can’t really provide any security, but that is prescriptively where the CIP standards say it must be done.

Chapter 3: Future Concept Example

Previously the proposed solution to address these issues with virtualization was a full embrace of the cyber system concept (retiring device level terms like Cyber Asset and BES Cyber Asset) and writing the technical requirements (e.g. CIP-005, CIP-007, CIP-010) at a more objective level. However, through industry feedback several important issues came to light:

- Eliminating the BES Cyber Asset definition and moving to the BES Cyber System as the sole foundational object with no further granularity would cause a complete overhaul and "do-over" of entities' CIP-002 processes with insufficient benefit for the effort required.
- Requiring grouping systems by function exclusively conflicts with a number of different ways the industry has used the systems grouping flexibility; for example to group all Cyber Assets of a like operating system for patching purposes.
- Stating security objectives in requirements must be at a level where they are clearly measurable.

With these issues in mind, the future concept is to leave the foundational definitions largely unmodified. A "Cyber Asset" would remain as-is and be inclusive of the hardware as it is today. However, including the hardware for virtual cyber assets does not work, even though it needs to be protected the same as a discrete physical cyber asset. To resolve this the SDT will create a separate definition to capture virtual cyber assets so that they can be added to the applicability of the appropriate requirements. The hardware that provides compute, network, and storage resources would become its own new term to capture shared infrastructure and allow requirements to be written for risks unique to that environment. BES Cyber Systems would remain as-is and allow for all the various ways entities have used that concept to date.

The direction remains to establish objective level requirements but these need to be written at the appropriate level. For example, if CIP-005 R1 were written as "Mitigate the risk of unauthorized network access", it is certainly an objective but is too high level to be measurable. No one would be able to state when, or to what degree this objective had been accomplished. If CIP-005 R1 were written as "Only allow known valid layer 3 IP addresses into or out of the network and implement this only at a Cyber Asset interface located only at the network edge", it would be a prescribed topology and a prescriptive "how." This should be avoided as it precludes much of the newer technologies presented in this paper. The latter is essentially the CIP-005 R1 of today and is one of the primary requirements that needs to change to remove this prescriptive topology and "how's".

An example of an alternative objective requirement that clearly describes a measurable "what" but avoids prescriptive "how's" for CIP-005 R1 could conceptually be:

> *Deny all access to and from the networks on which high and medium impact BES Cyber Systems and their associated PCAs are connected and only allow network communication that has documented access permissions including the reason for granting access.*

This is a much clearer objective to accomplish but it avoids telling the entity how or where they must implement it or how their network must be architected. As for backward compatibility, an entity could still define an ESP with designated EAPs and provide firewall configurations as they do today. However, an entity using the advanced virtualization techniques from earlier in this paper, or future techniques to control network access that have not yet been contemplated can still show they meet this same objective.

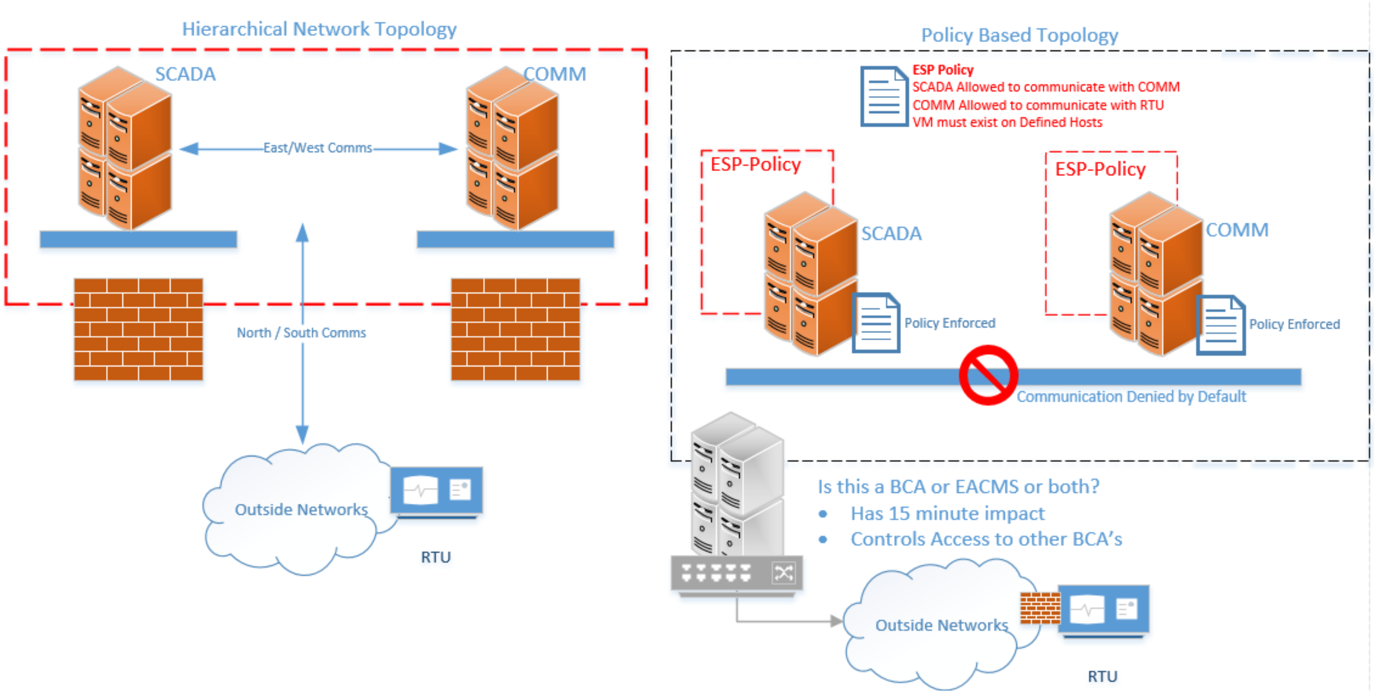
Hierarchical Network Topology
SCADA
COMM
East/West Comms
North / South Comms
Outside Networks
RTU
Policy Based Topology
ESP Policy
SCADA Allowed to communicate with COMM
COMM Allowed to communicate with RTU
VM must exist on Defined Hosts
ESP-Policy
SCADA
Policy Enforced
ESP-Policy
COMM
Policy Enforced
Communication Denied by Default
Is this a BCA or EACMS or both?
Has 15 minute impact
Controls Access to other BCA's
Outside Networks
RTU

Figure 12

As shown in Figure 12, the current state ESP/EAP model on the left continues to be a valid architecture and topology with which to meet the objective. It is prescribed under the current requirements, but is not the ONLY way to accomplish the objective. On the right is a policy-based model in a fully virtualized system which can also accomplish the same objective once it is no longer prohibited by the requirement language. It has no distinct EAP; no distinct 'point' or 'interface' where the access controls are implemented. This is because the access policy is implemented throughout the infrastructure and the controls move with the virtual workloads. Access control is applied at a deeper level, closer to the functions that require protection and the policy allows no communication to occur that isn't explicitly allowed. The objective stated in the requirement can be met in either scenario.

For the asset classification, nothing changes for the architecture on the left. For the virtualized architecture on the right, the concept is to clarify that the systems in orange are virtual cyber assets, but without including the hardware. They exist as a virtual entity and need much of the same protection as a physical cyber asset. The systems depicted in grey are the Shared Infrastructure, a new asset class that hosts virtual cyber assets. Having these additional asset classes would allow the CIP standards to require protection appropriate to each class, whether a discrete physical system, a virtual cyber asset, or shared infrastructure hosting many virtual BES Cyber Systems.

Chapter 4: Conclusion

This paper has stepped through the benefits of virtualization in terms of increased reliability and resiliency for BES Cyber Systems. It also discusses how virtualization provides newer techniques for securing access to those systems, such as distributed firewalls and zero trust models. It has highlighted some of the issues created by several of these concepts and the required topologies contained within today's CIP requirements. Many of these issues revolve around the proper classification of cyber assets in virtualized environments and having asset classes whereby requirements can be properly scoped. Several other issues revolve around the prescriptive "how" of today's CIP-005 ESP/EAP model, the CIP-010 data requirements in a virtual storage world, or how CIP-011's requirements need minor changes to allow for technologies such as remediation VLAN's.

The CIP standards require some changes to:

- Address risks unique to virtualized environments such as the sharing of hardware resources and management plane access.
- Provide clarity around how to identify and categorize the various types of cyber assets in a virtualized infrastructure and scope requirements appropriately.
- Allow entities to fully implement newer security techniques in these environments that can provide higher levels of access control that are easier to manage, and don't require purchasing extra hardware to show compliance with prescribed topologies.

Many of these issues can be addressed as shown in the CIP-005 R1 example concept, where the prescriptive 'how' has been removed and replaced with a clearly stated security objective – one that fully allows for current state but also allows these newer techniques. Developing technology-agnostic security objectives can not only solve the issues presented by today's virtualization technologies, but also help address future issues brought about by technologies that aren't even contemplated today.