NERC
NORTH AMERICAN ELECTRIC
RELIABILITY CORPORATION

Technical Rationale and Justification for Reliability Standard CIP-005-7

November 2018

DRAFT – WORK IN PROGRESS

RELIABILITY | ACCOUNTABILITY

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

Table of Contents

Introduction

This document explains the technical rationale and justification for the proposed Reliability Standard CIP-005-7. The Standard Drafting Team (SDT) is proposing changes to CIP-005-6 as it updates the standards based on technology innovation and changes such as the increasing use of virtualization.

CIP-005 - From ESP's to "Logical Isolation" – The Case for Change

In previous versions of the CIP standards, CIP-005 required declarations of Electronic Security Perimeters (ESP) based on OSI layer 3 routable protocols. All BES Cyber Assets (BCA) that were connected to a network with routable protocols had to reside inside a declared ESP. Any External Routable Connectivity (ERC) to the BES Cyber Assets inside the ESP had to enter and exit via an Electronic Access Point (EAP). This limited the traffic entering or leaving the ESP to only that which was needed. It denied all other traffic by default.

This "castle and moat with drawbridge" protection, where the castle is the BES Cyber Systems (BCS), the moat is the ESP, and the drawbridge is the EAP, has been in place for many years. For a number of situations, it still offers sufficient protection. This protection configuration, however, drives certain network architecture decisions that may hinder the adoption of newer and in some cases more secure designs. Network access control is expanding beyond perimeter-based security at routable protocol address levels into other levels, with security present within the network fabric itself. The entire network is becoming "the firewall." As infrastructure components become smarter, entities can describe network policy at layers other than that of a routable protocol address, and the infrastructure itself can enforce these access policies at multiple levels.

The nature of virtualization is driving much of this as workloads are dynamically created, moved, and destroyed to meet changing demand. With virtualization, what resides in an ESP might be very dynamic, and rulesets in any EAPs will need to be dynamic as well. This is driving solutions more towards policy-based user access controls for workloads that the infrastructure then dynamically applies within the environment. EAPs themselves are also dynamic, and as network access controls become embedded into the infrastructure, we may go from a few designated EAPs to thousands of dynamic locations where access controls are enforced.

This is the reason the CIP standards are changing from "perimeters" to "zones" and why the ESP has been replaced with the Logical Isolation Zone (LIZ). A perimeter typically implies an outer boundary; something you can draw a rectangle around on a network diagram. As network security becomes more integrated into the network fabric, or is virtualized and dynamic, it is no longer a perimeter that you can draw a box around. This is driving us to employ the zone construct for the future. An ESP is one form of logical isolation zone making it backward compatible, but a logical isolation zone can be broader than the layer 3, routable protocol subnet.

As firewalls continue to evolve, IP addresses and port numbers, which are simply protocol data, will become less common methods of identifying and filtering communications. What needs to be secured are the processes behind those layer 3 network addresses and the actual communications they are processing. IP addresses and ports have traditionally been filtered at the network layer to ensure communications include only what is needed. As technology continues to evolve, more secure future architectures may not be based on routable protocol addressing schemes.

Logical Isolation – What Does It Mean?

ESPs deal only with layer 3 routable protocol addressing. Virtualization and its accompanying shared infrastructure has other characteristics such as shared computing hosts, shared storage, shared virtual networks and switches, all of which posing new security concerns. To adapt to these changes, CIP-005 includes a requirement that high and medium impact BES Cyber Systems be “logically isolated” from all other systems (regardless of protocol) to replace the routable protocol-based ESP requirement.

If a BCS is executing on a virtual host along with another related virtual machine that may not have a 15-minute impact, (e.g. a control system and its data historian), the entity must either consider all these workloads as part of a single system or declare the system with the 15-minute impact as a BCS and the historian as a separate system. If the latter is chosen, the entity will need to prove that the two systems are logically isolated and that every communication between the two is limited to only what is necessary.

Logical isolation is also relevant to other layers of a computing stack. Imagine in a virtualized data center where you are looking at a BES Cyber System architecture “from the side”. You can see the different layers of the system, from the storage on a Storage Area Network (SAN) through the virtual networks, switches, and firewalls at the virtual network layer, up to the virtual hosts that are executing the application workloads. If you rotate the view until you are looking down at the system from above, you should be able to see the touch points of all these layers to other systems. Logical isolation refers to this top-down view. An entity needs to be able to show that only necessary data flows are allowed through any of these layers that have an interface to another system. For example, at a storage layer, BCS systems should have their storage logically isolated from other systems that are not part of the BCS. At a networking layer, there should be no communications channels that allow the BCS to talk to any other system that is not controlled. At a virtual machine level, there should be logical isolation between VMs. As you “look down” through the computing stack, you should only see interface points that are controlled and locked down to a least-privilege position.

Logical isolation does not mean complete isolation. It means that only known, controlled communications can occur between a system and anything outside of its Logical Isolation Zone, and that all other communication is blocked. Serial communications such as RS-232 or RS-485 are logically isolated communications methods as well. These types of communications move data from the TX pin on one end of a cable to the RX pin on the other end of the cable. There is no addressing scheme or routing/firewall capability, so it meets the intent of logical isolation.

Management and Data Planes and CIP-005 R3

As the SDT considered virtualization, it identified a risk though not unique to virtualized infrastructures, is amplified by it. As virtualized servers, networks, switches, firewalls, and storage are logical constructs, controlling access and communications to the management plane of these systems becomes 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 and must be brought into the scope of CIP standards if hosting BES Cyber Systems.

Another example from a networking perspective is virtual LANs (VLANs) which are a method of logical isolation that can be used to meet CIP-005 R1. However, such logical isolation is critically dependent upon proper configuration of the isolation and management of access to those logical controls within the management plane of the switch or other network device. CIP-005 R3 brings that management plane of the switch into scope of the CIP standards, requiring that entities restrict access to the management functionality of these systems and isolate that access from the data plane (the normal, operational, non-administrative access).

Examples of this separation of the management plane and data plane can be accomplished in several different ways:

- A hypervisor can put its management interface IP address on a physical port or in a logical interface connected to a VLAN that is trunked to an upstream switch.
- Many hardware vendors include "integrated lights-out" or dedicated management interfaces that allow you to power on/off a cyber system.
- Physically isolated out-of-band network for management interfaces.
- Logically isolated out-of-band network (where logical constructs like VLAN, VXLAN, MPLS, VRF's (logical routers) are used to maintain logical isolation)

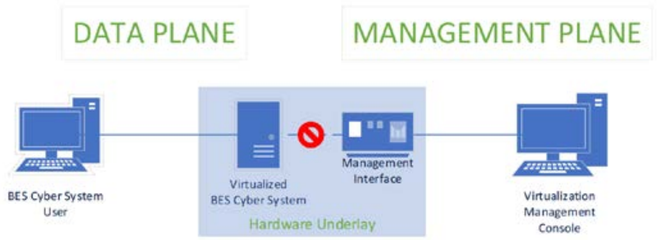

"Super-ESPs", CIP-005 R1.2, and the 4.2.3.3 Exemption

One of the issues with the ESP construct carried forward into the logical isolation model is the situation where entities have BES Cyber Systems that include components at separate locations. For example, if an entity has a need to replicate data at high speed between two databases in two different geographic locations to improve the resilience and reliability of BES Cyber Systems, the entity may have issues with the 4.2.3.2 exclusion in the standards that exempts the network and communications gear that is "between discrete ESPs." If this replication protocol is not a routable protocol, then no ESP can be created. The entity needs to be able to have a "Super-ESP" that can span more than one location.

The ESP model, along with the 4.2.3.2 exclusion within the CIP standards applicability does not lend itself to this construct. As technology evolves, there will be many more instances where a BES Cyber System may need to span locations. Another example could be the protection of a Transmission line consisting of digital devices at each end of the line with an IEC 61850 high speed network between them. This type of system, even with its components located in two geographically separate substations, can be shown to be logically isolated from other systems. The issue is that the 4.2.3.2 exclusion for limiting the scope of the involved communications gear between the sites does not apply. It is no longer "between" two ESP's or two logical isolation zones but is now completely "inside" a single BES Cyber System or a single logical isolation zone. Traditionally this situation could be addressed by installing an EAP at each site as an ESP boundary that would then allow the communications equipment between the sites to be subject to the exemption. There are issues however, such as:

- Installing a firewall at each end of an essentially point to point circuit with rules allowing devices at either end to talk to the other adds complexity, single points of failure, and provides little security benefit.
- Some of these scenarios are extremely time sensitive and the devices need to communicate at very high speed to perform their asset protection function. Going through layers of security could cause latency and reliability issues that could hamper the asset protection function of the system.

This 'Super ESP' construct has been addressed with a new exclusion and a new requirement part in CIP-005 R1. The new exclusion in 4.2.3.3 allows communication networks and data communication links used by an isolated BES Cyber System that spans more than one geographic location to be exempt from the standards. The new R1.2 in CIP-005 pairs with this exemption and requires that the data over this exempted communication be protected to preserve its integrity and confidentiality, with the exception in R1.2 for those time sensitive protection functions.

"SUPER-ESP" EXCLUSION

Entity has 2 BCS with "BCS 1" having components in two locations. A VPN has been established between the components of BCS 1 at each location to logically isolate that BCS across the locations. Further logical isolation occurs at other touch points to other networks or out of scope systems. The entity chooses an appropriate demarcation point at each location and the networks between those points are excluded unless CIP-012 is applicable.

(NOTE: Red Cloud icons only denote some form of logical isolation for the BCS, Blue Cloud icons denote some form of logical isolation for the non-BCS)

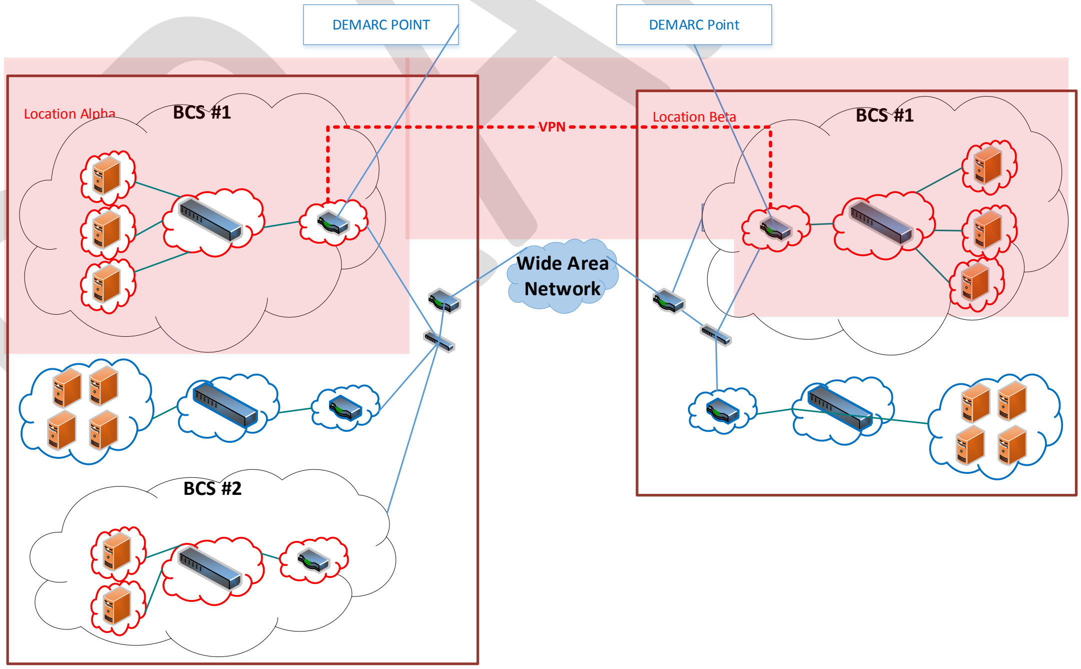

External Routable Connectivity and Interactive Remote Access

External Routable Connectivity (ERC) is used in the CIP standards for different purposes, including:

1. Establishing when EAPs are required
2. Limiting scope of ~38 requirement parts to those locations that have a high enough level of remote connectivity to support the requirement

The move to the more objective-based logical isolation concept shifts the requirement to an obligation that shows the system is logically isolated with controlled communications at any interface point or shared infrastructure. This can be accomplished without dictating any architecture or access control method. ERC is still needed as a scoping mechanism due to the vast scale of systems and their components within a geographically distributed BES. Many requirement parts should be scoped based on whether the system has ERC for the following reasons:

- The risk is increased for systems with ERC. The requirement should apply to those systems with an increased attack surface and risk due to their connectivity/accessibility.
- Locations that have legacy connectivity such as non-routable serial leased circuits should not have to increase their level of remote connectivity and attack surface to meet security requirements. For example, it would not be advisable to put in an IP network into a site to get SNMP traps out for alerts if a serial circuit with reduced attack surface is all that is needed for operations.

One issue with the ERC definition from the V5TAG transfer document has been that of BES Cyber Assets (BCA) that only speak non-routable protocols over a serial port. These BCAs are not in an ESP and therefore can be considered to not have ERC because it is defined in terms of an "associated ESP." These BCAs, however, can have Interactive Remote Access through an upstream serial-to-IP conversion. The SDT has kept ERC as-is with only conforming changes in order to not disrupt its scoping function as noted above. However, the IRA definition has been modified so that a device with only a serial, non-routable connection could now have IRA and be subject to CIP-005 R2. Appropriate controls (CIP-005 R2) are now required for these Interactive Remote Access sessions without regard to ERC.

The following diagrams show different scenarios and whether ERC and/or IRA exist in the situation.

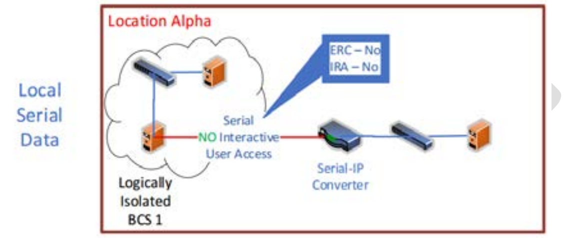

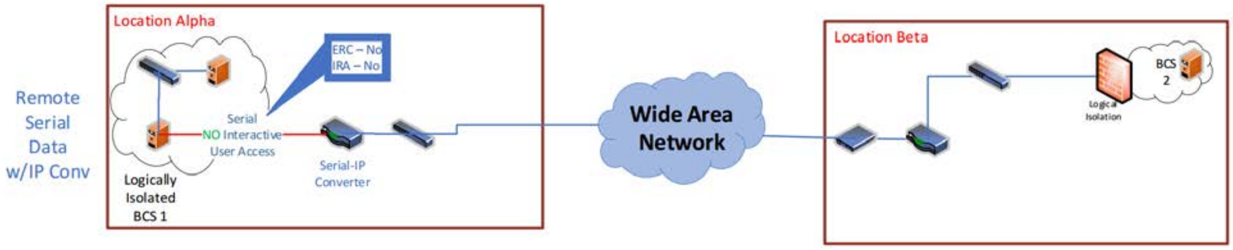

Remote Serial Access

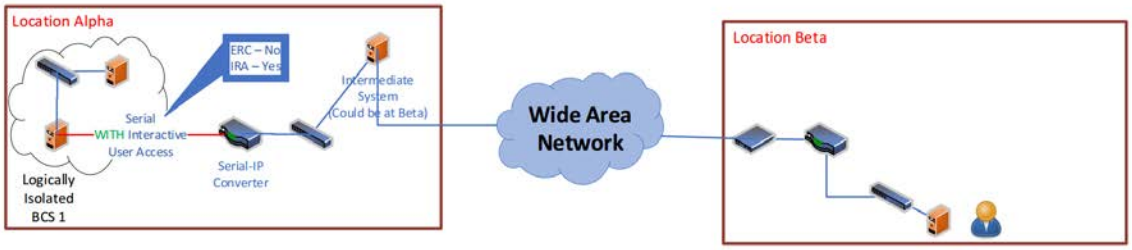

Remote Routable Data

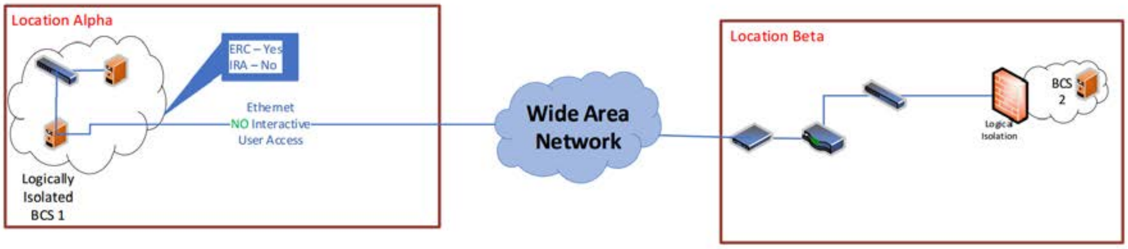

Remote Serial Data

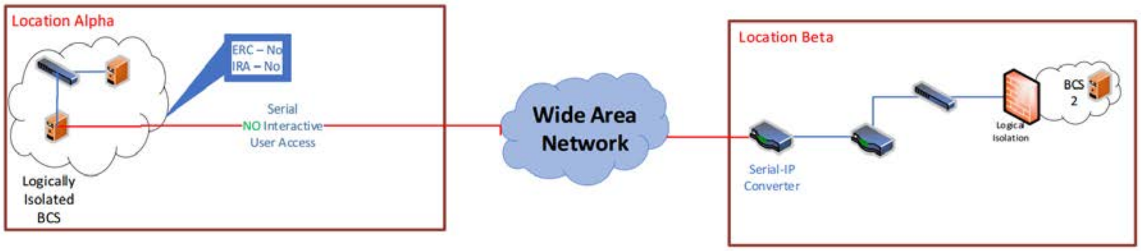

Remote Serial Access From Another BCS

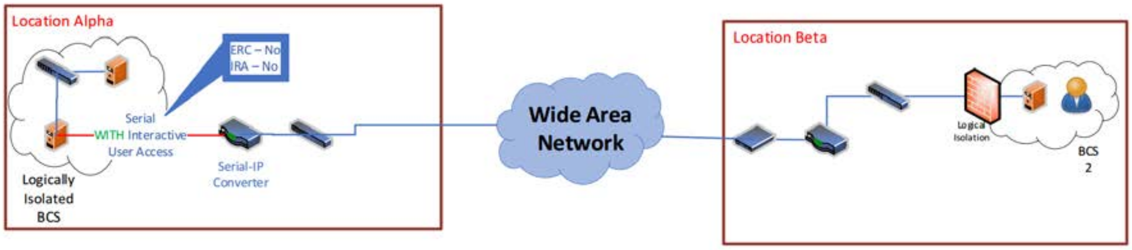

Diode

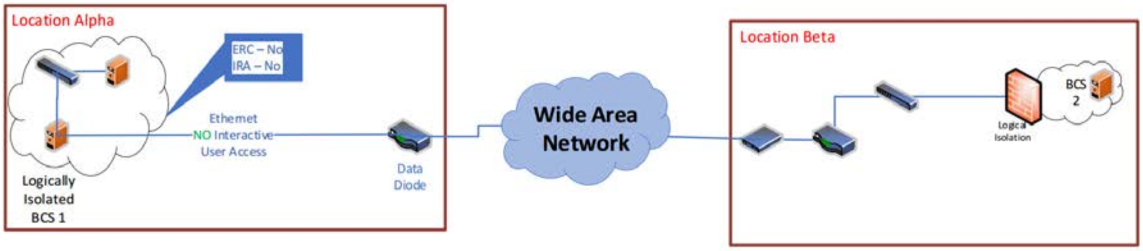

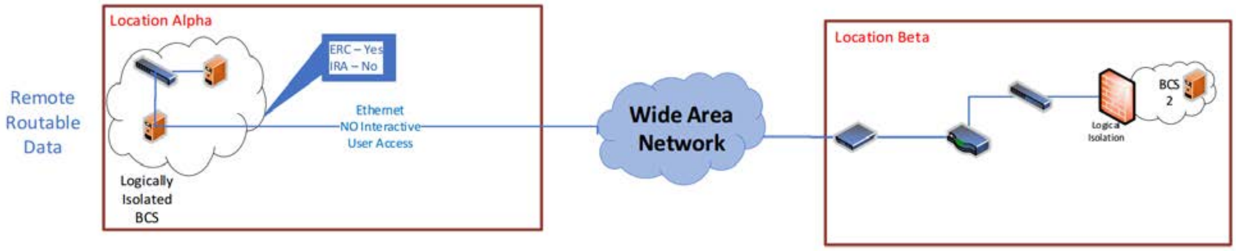

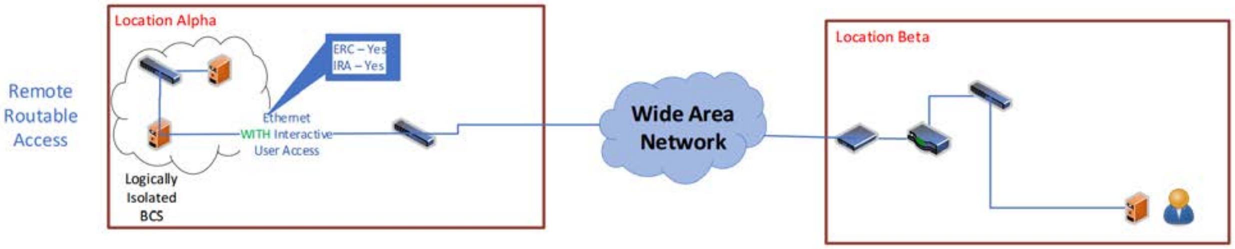

EACMS Changes to EACS and EAMS

As technologies and attacks have advanced and become more complex, entities are becoming more interested in partnering with outside and government security services. These includes services like NERC's Cyber Security Risk Information Sharing Program (CRISP), Cybersecurity for the Operational Technology (OT) Environment (CYOTE), and those of other external security services and internal monitoring centers. Going forward, these types of service and providers will become more cloud based. Security service providers have visibility into emerging threats and trends that come through their extensive collections of information. Analysis of this information can then be shared more broadly, improving the overall cybersecurity posture and reliability of the BES through early detection of compromise and the ability to monitor for threats and indicators of compromise (IOCs) at machine speeds.

Under the current body of CIP Standards, using the types of services that include electronic access monitoring data (not involved in the actual control of electronic access) may bring all Cyber Assets involved into scope as an EACMS. This may discourage or even preclude entities from using these services based on the cyber asset level requirements of an EACMS. These limitations affect personnel, physical security, patching, baselines, and other requirements that focus on a Cyber Asset. Entities may also be discouraged from providing and correlating security events across enterprise and control networks, even though most cyber-attacks against control systems today enter through business networks. There is great value in correlating security events seen across enterprise and OT networks that may be discouraged or precluded through the "M" in EACMS growing to include much an enterprise's other monitoring only Cyber Assets.

The cyber systems that do perform electronic access control will remain as they are today in the standards. Those cyber systems, such as firewalls and routers with ACLs and other systems that do perform access control and actively protect the networks to which BES Cyber Systems are connected should not change. However, the monitoring and logging aspect of EACMS presents a different though lesser risk. The creation of two different defined terms recognizes this. EACS represents those systems that do control electronic access and will essentially be a drop-in replacement for today's EACMS.