NERC CIP: Understanding the Fundamental Security Requirements of an Electronic Access Control and Monitoring System (EACMS)

Best Practice Guidance - NERC CIP-005 | CIP-007 | CIP-010

The NERC CIP standards are the primary knowledge resource used by the Utility industry to ensure our nation's power grid is protected from unintentional (accidental) and intentional (malicious) disruption. This whitepaper looks at the specific capabilities required and best practice guidance of an effective Electronic Access Control and Monitoring System (EACMS) for Privileged Interactive Access Management, Logging and Monitoring (Situational Awareness), and Baseline Configuration Management as covered by portions of NERC CIP-005, NERC CIP-007, and NERC CIP-010.

The ConsoleWorks technology from TDi Technologies addresses all three of these key NERC CIP requirements, and others, in one solution providing situation awareness across IT and OT assets, as shown.

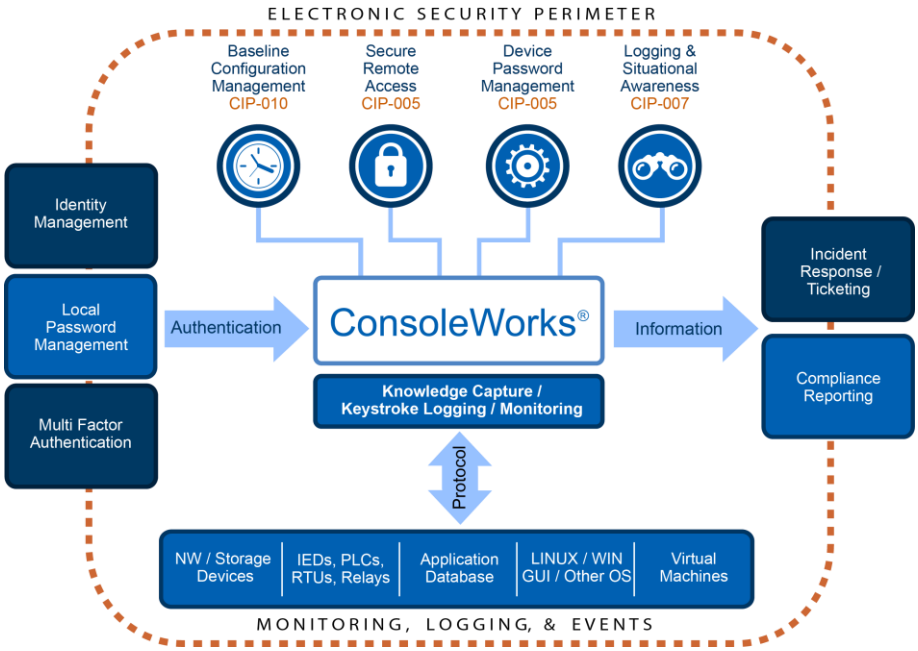

Secure, Interactive Remote Access and Situational Awareness CIP-005 and CIP-007

CIP-005-5 R2 is focused on ensuring that the security of the Bulk Electric System is not compromised by remote access. The general access control policy defined in section R1 is further augmented by the requirements of R2 for all remote access.

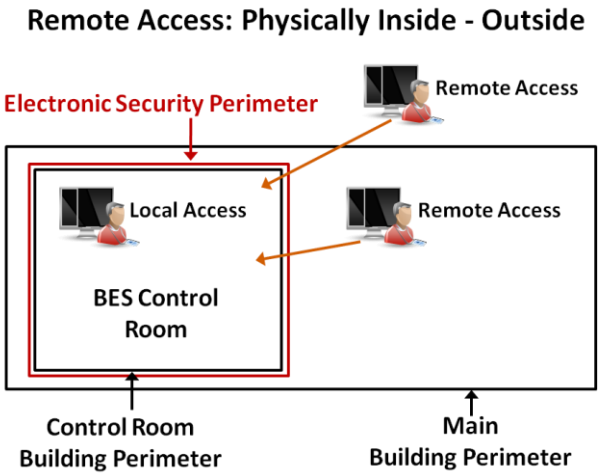

What is Secure, Interactive Remote Access?

Remote access occurs anytime an asset inside the electronic security perimeter (ESP) is accessed by a user that is outside of the ESP whether the asset is classified a Cyber Asset or not. This includes access from within a physical security perimeter and access from outside all physical security perimeters. The focus is on user access to the assets or cyber assets controlling the electric grid. As depicted in the diagram to the right, users may be inside a physical security perimeter yet outside the actual Electronic Security Perimeter or they may be at a remote location (i.e. – traveling, working from home, etc.) outside all physical and logical security. Only personnel inside the physical security of the ESP or control room (or other secured cyber asset location) would not be considered *remote*.

This is a new requirement for Utility organizations that was not included in previous versions of NERC-CIP. From NERC, it addresses “vulnerabilities for remote access methods and technologies that were previously thought secure and in use by several large electric security entities” (NERC-CIP-005-5 R2 - Rationale).

While NERC does not currently provide any requirements or guidance documents on how to accomplish secure remote access, NERC does define the key requirements that must be met by a secure remote access practice or solution in CIP-005.

The key requirements of CIP-005-5 R2 include:

- Implementing an Intermediate Device/System for Interactive Remote Access
- Encryption for all Interactive Remote Sessions
- Multi-factor authentication
- Up-to-date anti-malware software on user devices
- Up-to-date patch levels on user devices

Intermediate Device / System Explained

A firewall or other electronic access point (EAP) device provides access denial, unless authentication is accomplished, and limited access based on roles. Once authentication is accomplished, it allows the user to directly connect to one or more cyber assets, networks, or other logical elements. In the simplest terms, it is a locked door on the perimeter that must be opened to gain access.

With Remote Access Management there is another step that is required to meet the CIP-005 R2 requirements. Instead of gaining access through the EAP, the user gains access to an Intermediate Device (ID). The ID is connected to the cyber network inside the ESP. In this configuration, there is no direct connection between the remote user and the BES or cyber system. This kind of Intermediate Device is often called a jump box or bastion host. The jump box provides an added layer or buffer to the security, never directly exposing a cyber system to a remote cyber asset.

Compared to the EAP "locked door" analogy, CIP-005 R2 can be viewed as two locked doors, with the second door opening from a secure 'room' to the BES cyber network. Authenticated users enter the secure room where they can issue commands that the room can then carry out to the BES cyber asset. In this scenario, the user is never directly connected to the BES cyber system or network, as shown below.

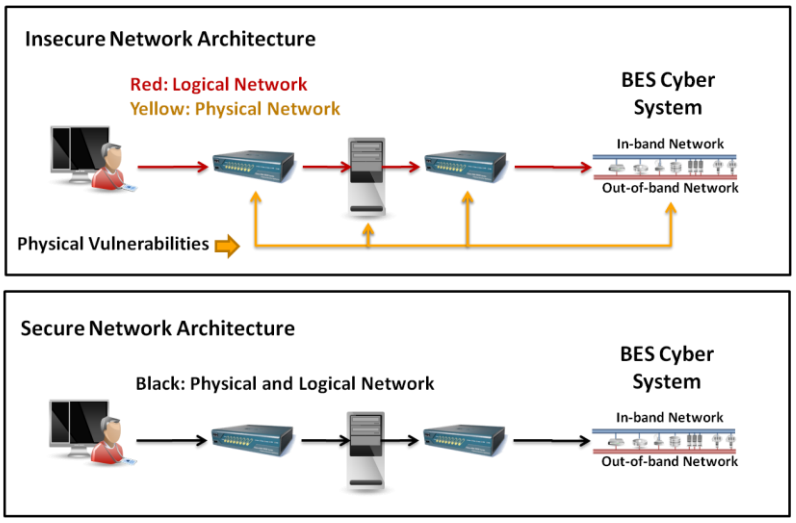

By deploying defense-in-depth with layering of firewalls, role-based access management, and high availability failover of security status monitoring and event logging, entities can be assured of data integrity, rapid incident response and disaster recovery.

Intermediate Device / System Advanced Capabilities

An Intermediate Device / System can provide advanced capabilities to harden the security footprint without impacting user performance – a major drawback in many Intermediate Device approaches.

Advanced capabilities in Intermediate Devices can include:

- Role-based access and control that limits each user's access to a predefined set of cyber assets in the BES
- Least privilege by user and/or role, limiting privileges to lowest level needed to perform the work the user is authorized to perform
- Capture of all user activity down to the keystroke (CLI).
- Capture of system messages from application logs, SNMP alerts, SYSLOG and other sources
- Management of all cyber assets in the ESP – not just servers. Support of more than just a single OS, application or interface.
- Support for normal and emergency operations including power reset, firmware management, BIOS Configuration as well as multi-user privileged access.
- Event detection and alerting – predefined and admin configurable
- Single pane-of-glass oversight
- Business rules that restrict, alert, or control user activity
- Coverage for all privileged interfaces, in-band (production) network and out-of-band (maintenance) network

The items in the list above (in particular items 1,2,3,7,8) are common security practices typically falling under the practice of Privileged Access Management (PAM).

★★★★★★★★★★

Security Perspective: Personnel remotely accessing the ESP must be managed per CIP-005 R2, but they also must be managed per CIP-007 (situational awareness through strong role-based access control, logging, and real-time event /incident detection).

From this perspective, even personnel inside the ESP should be utilizing the same controls for access as those outside the ESP. This ensures that a consistent view, method and process is used for control - all the while having command and control that is logged and audited in such a manner as to thwart an insider attack or insider unintentional impact.

CIP-005 R2 should be considered in the broader scope of the NERC CIP regulations when formulating an overall security strategy. Addressing the new requirements for Remote Access Management in isolation can result in a fragmented security solution with gaps that can have a significant impact on the ability of the Utility organization to support reliability and security - *its primary objective*.

★★★★★★★★★★

Privileged Access Management

Privileged Access Management is a highly appropriate value-added role for an Intermediate Device for NERC-CIP-005-5 R2. Restricting access to specific cyber assets, networks, or other logical elements at each EAP (the traditional approach) is valuable, but it does not provide fine-grained control over what a user can access or the privileges granted for each. Intermediate Devices should serve as the fine-grained control mechanism for the Remote Access Management practice.

Supporting the Role of People

The most secure access profile eliminates remote access altogether. This, however, is unreasonable as it would require that all staff - address security threats, perform IT maintenance, and respond to emergencies - would be required to be physically present inside the physical security of the ESP (24/7/365).

This brings up an extremely important point. When personnel are accessing the cyber systems remotely, they are typically doing so under conditions demanding fast response and expert skill. This most commonly occurs in one of the following two scenarios: 1) issues that threaten availability of the grid (alarm, outage, service issues, etc.) and 2) security threats.

An Intermediate Device can serve multiple purposes by addressing the context of the people needing remote access with supporting capabilities that improve their ability to resolve operational or security issues.

To do this, the Intermediate Device must have very good situational awareness of the complete cyber system inside the ESP – from hardware to the application and all points in-between. This way, when the remote or local user accesses the ID they will not need to look in multiple places to gain a forensic understanding of failures, degradation or other issues affecting availability (which cause remote access to be used by remote personnel). This directly supports the primary objective of Availability while providing the appropriate level of security.

Forensic capture and logging down to the keystroke of privileged user activity provides another important security and compliance function. Capturing privileged user activity actively deters out-of-policy behavior while the information it contains is often critical to resolving issues and mitigating security threats caused by human error or malicious intent by Insider threat.

Remote access to both in-band and out-of-band networks is a requirement. The out-of-band network is the only network and interface where emergency operations and actions can be taken to resolve hardware or software failures, including configuration issues related to hardware, operating system, network and often applications.

Where the Intermediate Device can capture system messages (application logs, SYSLOG, SNMP and privileged actions for single user and multi-user access et. al.), it can serve a dual purpose by automatically producing compliance records and for retaining information needed by remote users to troubleshoot issues, confirm operations, and be alerted to potential problems.

This ability also has impact on mean-time-to-repair (MTTR) since forensic information needed to repair or correct the BES configuration would be at-hand in the event of an outage or degradation. Having the information readily at-hand, eliminates the need to look in multiple places and find available logs messages (assuming they still exist or in some cases may have been a single last gasp alert not stored or no longer available).

From a management perspective, the Intermediate Device with the capabilities listed here provides comprehensive oversight and transparency. Because the Intermediate Device is effectively a single point of connection for remote and local users to the cyber system, it has access to all of the information needed to provide a single source of all Remote Access activity. This makes the Intermediate Device the ideal source for single pane-of-glass oversight and situational awareness.

Advanced Intermediate Device Solution

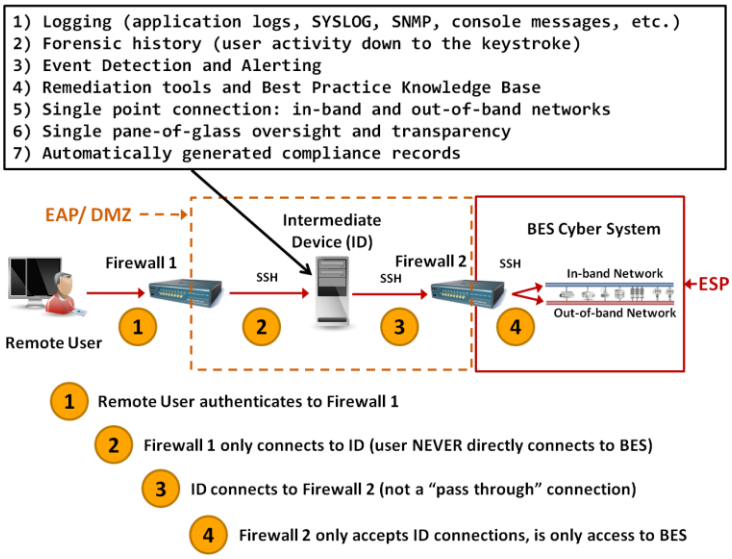

Ideally, the Intermediate Device would automatically confirm the user's device has met malware and patch level requirements before allowing the user to connect to it – although this may be instituted as a separate security procedure.

The challenge here is that the remote cyber asset is responding to queries by the Intermediate Device and as such, if infected, could very well be malware responding to queries by the Intermediate Device.

Instead of the Intermediate Device querying an agent on the remote device, it is better for the Intermediate Device to simply eliminate the ability to communicate (other than human communication) to cyber systems through its connections. In other words, no direct, outside protocol is allowed to communicate past the Intermediate Device.

Best Practice Guidance

The best practice guidance for Remote Access Management is to look at the problem holistically, including in that view, the perspective of Privileged Access Management. Utility organizations should recognize upfront that performance penalties are not a trade-off they must make, but that due diligence will likely be required on their part to ensure that performance degradation does not become part of the outcome of the Remote Access Management practice.

Specific consideration for Intermediate Device solutions should be given to:

- Define a strategy that will achieve security goals without negatively impacting performance or availability of the BES but instead increasing the efficiency of the remote user and their situational awareness of the BES
- Ensure the solution covers both in-band and out-of-band interfaces / networks – both REQUIRE remote access
- Institute a fine-grained, role-based access and control model for all users (restrict access, least privilege)
- Ensure the Intermediate Device can encrypt communications from the remote or local user to the BES cyber system

- Ensure the Intermediate Device supports multi-factor authentication
- Include all user activity logging in the solution. It provides additional security, can eliminate manual compliance reporting work, and enables oversight
- Review other potential benefits an Intermediate Device may have, such as automating portions of the security and compliance practice (event management, etc.)
- Consider a solution that can capture system messages. There may be an opportunity to further reduce compliance report generation costs while improving user performance and provide better situational awareness of the BES to the remote user responding to issues related to security, availability or performance of the BES
- Consider the solution from the Privileged Access Management perspective. A well-architected solution increases performance and can be applied to ALL users, reducing or potentially thwarting insider threats
- Consider the ability of the solution to perform configurable alerting and alarming and how that might be used to gain further oversight and proactive notification of security-related events
- Have a solution that has persistent agent-less monitoring – not polled solutions. Polling introduces windows of invisibility
- Look for a Vendor supported complete solution – Hardware, Operating System, Application. Training and installation – Turnkey
- Never provide command line access to the Intermediate Device.

In the case of Intermediate Devices, there are quite a number of different hardware and software solutions that can address at least a portion of the best practices outlined here. Some strategies may require multiple components to be used – noting that this may require additional work to properly harden the resulting Intermediate Device - while others may use a single solution (which may still need to be hardened).

The most important mistake to avoid is to take a minimalist approach that is focused on simply meeting the stated requirements of CIP-005 R2. While a properly designed strategy can deliver advanced security without negatively affecting performance (probably improving it), a less considered approach can have significant negative consequences to performance.

Controlling Access Through Configuration Ports

Configuration ports on critical and non-critical cyber assets are often misunderstood and overlooked in the overall cyber security strategy. This paper discusses the importance of configuration ports in the overall cyber security strategy and how they apply to the NERC CIP standard. An Industry Advisory from NERC with additional details on this subject is available here:

http://www.nerc.com/fileUploads/File/Events%20Analysis/A-2008-05-13-1.pdf

Configuration ports exist on almost every hardware device in the IT infrastructure. These physical ports provide a special level of privilege access that can be used to:

- Change Bios
- Upgrade Firmware
- Set Baseline Configuration
- Build-out devices that have components (like servers)
- Perform a variety of Administrative functions
- Perform emergency repair or failure recovery when no other port is accessible

Item six in the list above is very telling in respect to the important role these ports play in the cyber security strategy. Except for power supply or catastrophic electronic component failure, configuration ports are active at all times – even when conditions have degraded a device to the point that no other port can accept communications. They are the default emergency access point for every cyber asset device.

Per CIP-007 all ports should be either secured or disabled and documented for normal and emergency operations However, most cyber asset devices do not allow the disabling of these ports nor should these ports be disabled as they serve important purposes, including being the primary emergency access port. Instead, these ports must be secured.

A significant influence on the severity of the threat an access port presents to the Utility organization is the privileged capabilities the port presents to its user. Configuration ports present an extremely high set of privileges that can be used to change almost anything on the target device. This level of privilege is why access to configuration ports is often referred to as having the "keys to the kingdom."

The list of severe security threats over configuration ports is impossible to fully document due to the range of privileged commands these ports provide to its users.

Some of the more obvious threats are:

- communication ports can be changed or added
- data can be copied
- malware can be installed at multiple levels (Bios, Firmware, OS)
- user accounts and privileges can be added, changed or deleted
- device configuration can be changed
- ports are "discoverable" making them targets for malicious actors

The simple fact is configuration points are an extremely high security issue that can be exploited under a variety of scenarios where other security technologies, techniques, and practices cannot detect an active exploit.

Best Practice Guidance

The best practice guidance for configuration ports is that they should be treated just like any other security concern in regards to active monitoring and control. The steps that should be taken include:

- Insure that all configuration ports are connected to an out-of-band or management specific network
- Segregate the out-of-band network from the normal or production network(s)
- Institute role-based access and control over all configuration ports (restrict access, least privilege)
- Encrypt communications to configuration ports (where supported by devices)
- Use proper or multi-factor authentication to configuration ports
- Persistently monitor all configuration ports to ensure all access meets the security policy
- Log all access to configuration ports by each actor
- Log all privileged user activity over configuration ports
- Alert and ALARM on specific messages or events detected on the access port.

Practical Application of Baseline Configuration Management (NERC CIP-010)

The practical application of baseline configuration management involves a process to systematically control assets to a defined configuration state known to be the "most secure" configuration.

This process is defined with CIP-010 regulation, explicitly calling out primary security issues that include operating systems, security patches, software versions, logical ports, system services, and firmware.

The overriding purpose is to maintain cyber asset configurations at a known state that has the highest level of security. While the regulation calls out specific areas that need to be part of the baseline configuration management practice, it should not be construed as a definitive list. Any configuration elements that can impact or impose a security issue should be included in the overall practice.

The process behind baseline configuration management has these elements:

1) Establish a known secure baseline for each asset
2) Periodically collect the current configuration and compare it to the baseline
3) If available – verify that the volatile and permanent configurations are the same
4) Identify any changes between the two configurations
5) Create a change request for any changes that are not rolled-back
6) Assess cyber security controls that could potentially be impacted by the change
7) Review and Authorize the change (Change Advisory Board)
8) Update the baseline configuration per authorized changes

Best Practice Guidance

The best practice guidance for configuration ports is that they should be treated just like any other security concern in regards to active monitoring and control. The steps that should be taken include:

Manually collecting, recording, and assessing all of the configuration data on cyber assets is labor-intensive and may introduce inadvertent errors. In most cases, running a monthly or weekly baseline would be prohibitively expensive. With automation the baseline frequency can easily be supported as a routine security validation function to improve the security posture and be audit ready at any time.

The other best practice recommendations are:

- Capture, document and secure baseline data programmatically (eliminate human error)
- Perform comparisons programmatically and determine where differences exist (eliminate human error)
- Provide a change control portal for implementing configuration changes, logging who, what and how an asset was patched. A complete change session log
- Generate alerts when mismatches are found and review with the privileged actor that made the change (process improvement)
- Document the approval process
- Run comparisons frequently if technically feasible
- Force rerun immediately after approved changes are made (validation)
- Produce reports that detail all changes (and their dispositions)
- Review reports for patterns that can improve the practice

Various hardware and software solutions exist that can help meet these challenges. These solutions should be evaluated against organizational context and specific requirements for the different asset types in place in order to cover the broadest asset footprint possible – keeping in mind exceptions will require manual baseline configuration management.

About this Whitepaper

This whitepaper was written and produced by TDi Technologies, a software vendor that provides Secure Remote Access, Configuration Monitoring, Password Management, and Logging and Auditing solution to the Utility industry and other vertical markets. The whitepaper is intended to provide useful and educational content that can assist Utility companies in providing secure, dependable power to our Nation without interruption.

This whitepaper was written to help address NERC CIP V5 requirements in the Utility industry. The recommendations provided are believed to be accurate in their applicability and support for NERC CIP-005, CIP-007, and CIP-010.

If you would like to receive additional whitepapers on NERC-CIP from us as they become available, please email us at info@tditechnologies.com.