The Lighthouse

By: Lew Folkerth, Principal Reliability Consultant

Jan/Feb 2020

40 Mile Point Lighthouse, Rogers City, MI – Photo: L Folkerth

Out-of-Band Management

In this recurring column, I explore various questions and concerns related to the NERC Critical Infrastructure Protection (CIP) Standards. I share my views and opinions with you, which are not binding. Rather, this information is intended to provoke discussion within your entity. It may also help you and your entity as you strive to improve your compliance posture and work toward continuous improvement in the reliability, security, resiliency and sustainability of your CIP compliance programs. There are times that I also may discuss areas of the Standards that other entities may be struggling with and share my ideas to overcome their known issues. As with lighthouses, I can't steer your ship for you, but perhaps I can help shed light on the sometimes stormy waters of CIP compliance.

What is Out-of-Band Management?

Out-of-band management is a method of managing computer systems that does not rely on having a physical presence at the computer system. This approach involves a network interface on the computer system that is used outside of the normal network connectivity, hence the term “out-of-band.” Since the purpose of out-of-band management is to manage the server remotely, almost all out-of-band management is a form of remote access.

Most datacenter-class servers have the capability for out-of-band management. For example, Dell offers its "integrated Dell Remote Access Controller (iDRAC)" and Hewlett Packard Enterprise offers the "integrated Lights Out (iLO)" controller. All server vendors that I've researched offer some form of this capability.

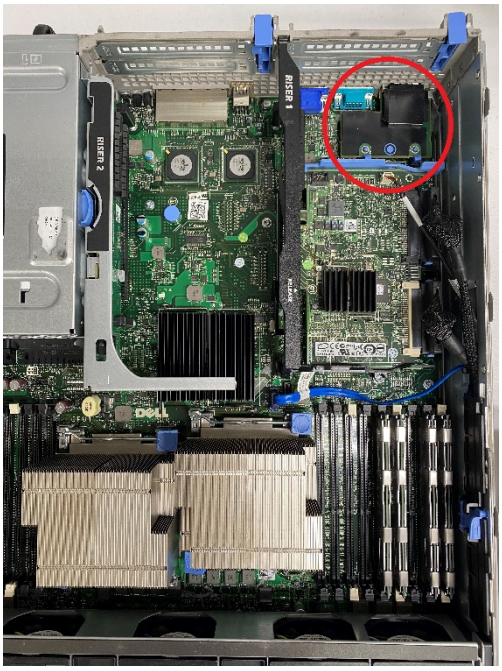

Figure 1 - The R710 Server with iDRAC 6 Enterprise Circled in Red

Out-of-band management is usually implemented by adding a controller with its own network interface to the server. The controller is an additional small computer with extensive monitoring and control capabilities for the server. Server vendors implement the management controller in different ways: as an integral capability of the core server board, as a daughterboard on the core server board (illustrated in Figure 1), or as a separate device outside the core server.

In this article I'll discuss the functionality of the Dell iDRAC 6 Enterprise (see *Figure 1*) installed in a Dell R710 server in a security testing environment (see Figure 2). If you're up on Dell technology, you'll realize that this equipment is at least three generations old. But the capabilities remain in modern hardware, and in most cases the capabilities have been enhanced. Please note that neither RF nor I endorse or criticize any individual vendor. I am using Dell because used equipment is readily available and that's one of the systems I have access to in the security testing environment. In the exercises below, I am using only documented features of the equipment. I do not use any exploits or other penetration testing techniques against the iDRAC.

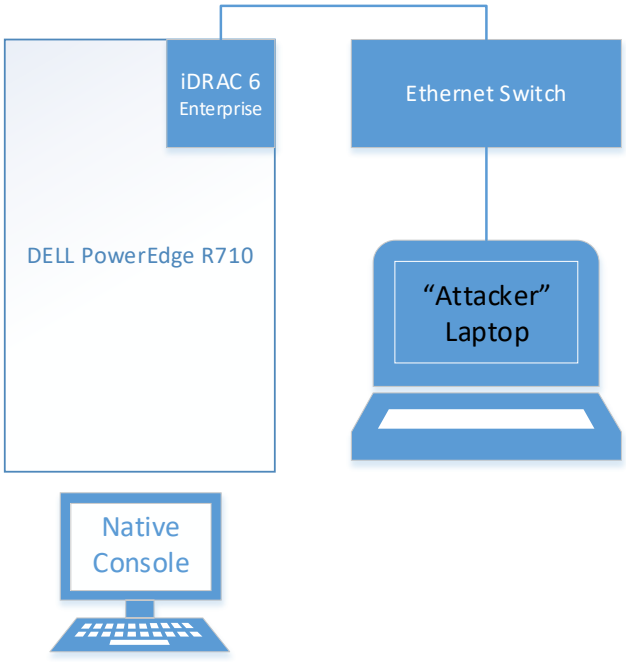

Figure 2 – Security Testing Environment

Remote Console

A significant feature of the iDRAC is the ability to access the server's hardware console remotely. This is not the same as using remote access client software to sign in to a Windows or Linux operating system. Once you have signed in to the iDRAC, you can bring up the iDRAC remote console and see the same display as the hardware video port on the server. The remote keyboard and mouse behave exactly like they are directly connected to the server. Why is this important? The remote access capability is available even before the system boots its installed operating system. On power up, the remote console sees the boot-up sequence and can enter BIOS and other console-only modes to

configure the system, possibly without further authentication.

Figure 3 shows the security testing environment setup with the R710 at the BIOS screen and the native console being mirrored by the iDRAC remote console running on the laptop. In this mode I can make BIOS changes and any other changes that can be made from the local console.

Figure 3 – Laptop Running iDRAC Remote Console at BIOS Screen

Web Interface

The iDRAC has many more capabilities. Many of these capabilities can be accessed through a web interface via the iDRAC port on the server. The web interface capabilities include:

- Monitoring server temperatures, voltages, and power consumption;
- Setting the device that the server will boot from next;
- Power on, power off, or perform a hardware reset to cause a reboot;
- Upload a disk image to the iDRAC internal storage, and then boot from that image;
- Create a blank disk image on the internal storage and make that image accessible to the server; and
- Download an image from the internal storage.

One of the exercises I've performed involved obtaining administrative access to the R710 through documented features of the iDRAC (and a little password cracking). With only the default iDRAC credentials, I was able to obtain files containing encrypted passwords. I then cracked the encrypted passwords on a penetration testing system, and was able to remotely sign in to the server's operating system with full administrative privileges.

Are out-of-band management capabilities inherently bad? Of course not. They can be very useful in managing a server at locations such as substations or control centers that do not have local IT staff to manage the IT-type systems. Use of out-of-band management capabilities can improve reliability by shortening downtime and by permitting monitoring of systems so preventive actions can be taken in a timely manner.

Compliance and Security Recommendations for Out-of-Band Management

The out-of-band management controller has extensive capabilities to remotely control, modify and operate its host computer. But the price of that functionality is risk. For any BES Cyber System, Protected Cyber Asset (PCA), Electronic Access Control and Monitoring System (EACMS), or Physical Access Control System (PACS), that risk must be mitigated. We need to apply our best security and compliance practices to protect the management controller of an in-scope system. We begin by appropriately identifying the management controller within the CIP scope, and then apply the

appropriate protections to the controller. The protections will include all applicable CIP Requirements, but may include additional protections as needed.

Identification – Integrated Management Controllers

Most servers will have the management controller built in to the server's motherboard or internally connected as a daughterboard (the R710 uses a daughterboard as seen in *Figure 1*). This means the controller is part of "the hardware, software, and data in those devices" per the definition of Cyber Asset.

The best approach I've seen in applying the CIP Standards is to identify the management controller as a Cyber Asset that is part of the hardware of the server. Since it is part of the server it must be classified the same as the server. For example, if the server is part of a high impact BES Cyber System, then the management controller would be identified as part of the same BES Cyber System. The controller would be tracked in your documentation as a separate Cyber Asset, even though it is actually part of the server.

Whether you use this approach or devise an approach of your own, be sure to identify and document ALL of these management controllers. Audit teams are aware that these capabilities, if not protected, can present a high risk to reliability and are actively monitoring for any of these interfaces you might miss.

Identification – Shared Management Controllers

Some servers employ shared management controllers. This can occur when more than one computer shares a single chassis. For example, Dell's VxRail Series G can contain four "nodes" (independent computers) in a single chassis. The iDRAC interface is housed in the chassis, not the node. The entire chassis, including the four nodes, is accessed through the single iDRAC controller in the chassis. If the chassis and all four nodes are considered to be one Cyber Asset, or all nodes are identified as the same classification, then the approach used for an integrated management controller could be used.

If the nodes are considered to be separate Cyber Assets with different classifications, the identification process becomes much more complex. You will need to demonstrate that internal access between nodes in the chassis is controlled, and that access through the management controller is controlled to the level required by the applicable Standards.

Let's look at some examples:

1. Assume a chassis has two nodes; Node 1 is classified as a BCA that is part of a high impact BES Cyber System and Node 2 is out of scope for the CIP Standards (see Figure 4). This is a case of mixed-trust within the chassis and will need to be carefully addressed. You will need to provide evidence that Node 2 has no access to Node 1 except through an Electronic Access Point (EAP). This will involve providing evidence that the management controller does not provide a communication path between the nodes, and the chassis does not provide a shared path to storage, memory, network, or other facility that could be used to transfer data to Node 1 or control Node 1. The management controller must be within an Electronic Security Perimeter (ESP) as it provides remote access to a BES Cyber System but cannot be an Intermediate System (an Intermediate System must be outside an ESP, so an EAP would need to be identified

between the management controller and Node 1). The management controller should be identified as an EACMS as it controls access to Node 1. The network interface of the management controller should be within an ESP, as the management controller does not have the ability to act as an EAP.

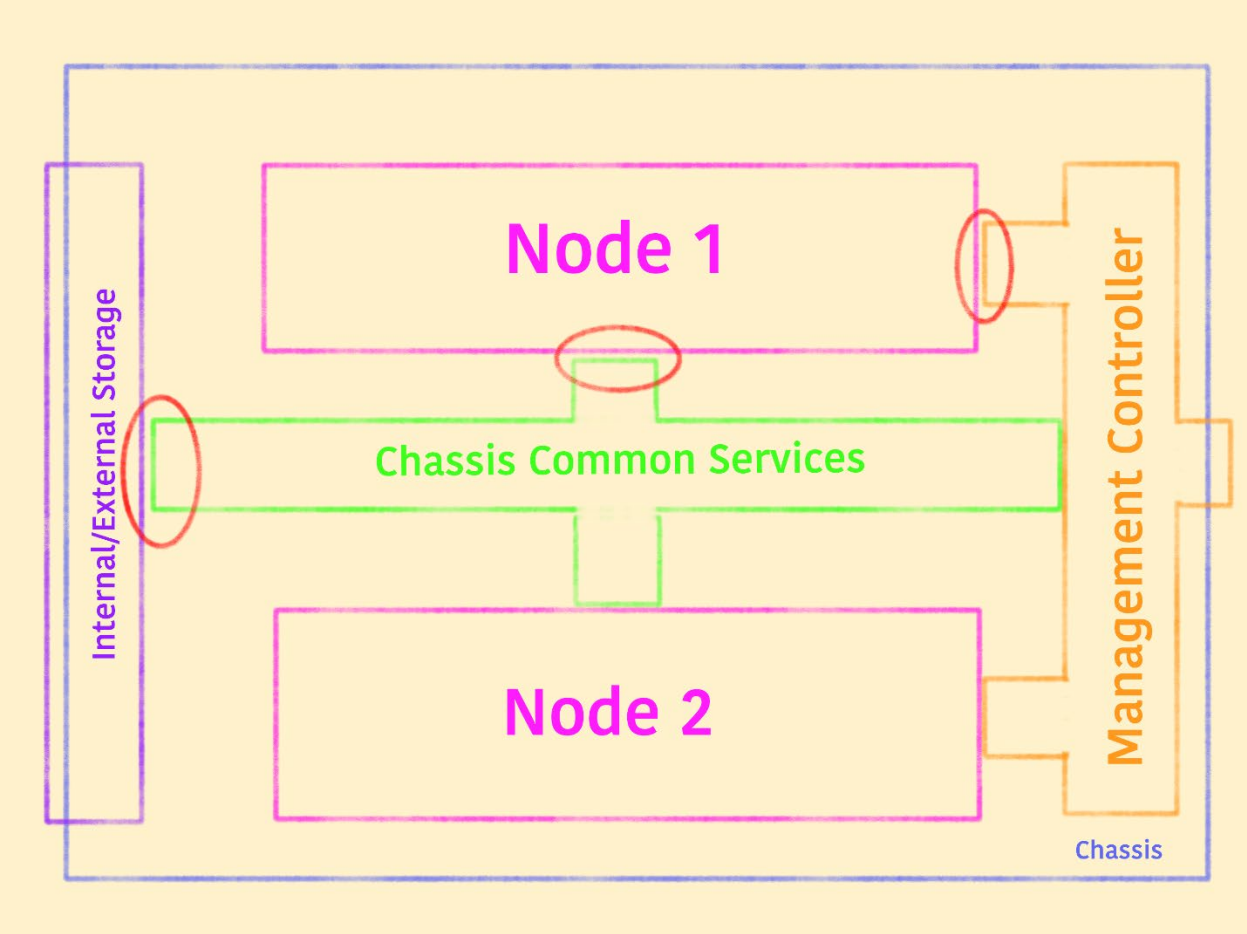

Figure 4 – Showing a two-node chassis with critical access points circled in red

2. While it may be possible to mix a node within an ESP and a node not inside an ESP on the same chassis and share a management controller, I do not recommend this approach. In order to demonstrate the logical separation of the nodes to an audit team, you will need to be able to provide evidence of the internal isolation of each node for each possible interface within the chassis (see Figure 4). I have not seen this done successfully in the field.

3. Now let's assume that Node 1 is an EACMS that is outside of an ESP and Node 2 is again out of scope (see Figure 4 again). Since we are outside of an ESP we don't need to have external network access to Node 1 or its management controller go through an EAP. How we identify the management controller in this situation gets very tricky. The controller itself does not meet the definition of an EACMS, since the definition only applies to a Cyber Asset that controls access to an ESP or BES Cyber System. That leaves no category within the CIP Standards with which to designate the management controller.

However, the management controller does control access to Node 1 which is an EACMS. This means that to control access to Node 1, we must control access to the management controller in addition to controlling all other forms of access. That brings CIP-004-6 R4, Access Management Program, and R5, Access Revocation, into scope for the management controller.

In addition, I strongly recommend voluntarily applying the full security protections of the CIP Standards to the controller, including restriction of network traffic (firewall), remote access through a jump host with multi-factor authentication, patch management, vulnerability assessment, security status monitoring, and change management. Above all, the default accounts and passwords for these devices are well known, so be absolutely sure you have changed the default passwords!

The remaining discussions assume an integrated controller for a system within an ESP.

Networking

Most server vendors recommend connecting the management controller to a network that is separate from the other networks connected to the server, hence the "out-of-band" designation. For servers within an ESP, this separate network must also be within an ESP. Otherwise the management controller would be an EAP, a role it is not suited to adopt.

Access Control

You must control access to the management controller at least as tightly as you control access to the server itself. Interactive Remote Access to a management controller within an ESP must be through an Intermediate System.

Baselines, Patching, etc.

The management controller should be subject to the same requirements as the server for baselines and change control, patch management, vulnerability assessment, ports and services, and password management.

Conclusion

Be sure to review all of your Cyber Assets within CIP scope and identify the out-of-band management capabilities of each. Document the presence of this capability on each applicable server, identity these devices in your Cyber Asset lists or baselines, and apply the appropriate CIP Standards to each. Be certain you have changed the default passwords.

In this short article I've only scratched the surface of management controllers and out-of-band management capabilities and concerns. As CIP and cyber security professionals, we must keep in mind the risks and benefits of using these capabilities, minimizing the risks and maximizing the benefits. We also need to monitor the Standards drafting efforts to ensure new CIP Standards meet our current and future needs.

Requests for Assistance

If you are an entity registered within the RF Region and believe you need assistance in sorting your way through this or any compliance related issue, remember RF has the Assist Visit program. Submit an Assist Visit Request via the RF website here.

Feedback

Please provide any feedback you may have on these articles. Suggestions for topics are always welcome and appreciated. I may be reached at lew.folkerth@rfirst.org.