A Risk-based Approach to Determining Electronic Security Perimeters and Critical Cyber Assets

IOActive™
COMPREHENSIVE COMPUTER SECURITY SERVICES

701 5th Avenue, Suite 6850
Seattle, WA 98104

Toll free: (866) 760-0222
Office: (206) 784-4313
Fax: (206) 784-4367

Introduction

Our current understanding of Critical Infrastructure Protection (CIP) began in May 1998 when Presidential directive PDD-63 recognized certain parts of our national infrastructure as critical to our national and economic security, and outlined required steps to protect it. This was updated in December 2003 with Homeland Security Presidential Directive HSPD-7, which enlarged the scope of what was considered infrastructure to include:

> *the physical and virtual systems that are so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety.*

When analog infrastructure assets were physically isolated from one another the impact from an attack could be mitigated easily on site and contained, but digitization and virtualization have blurred the once-distinct lines between physical and virtual assets to the point that an attack in one sector directly affects those connected to it. Add to that the fact that it can be easier to break into an asset's network than it is to access the physical infrastructure, and you've got the possibility that one computer virus cripples an entire region's transportation, emergency services, and power.

To mitigate this possibility, the North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection Standards (CIPS) requirements 002–009 describe the cyber security standards with which bulk electric power providers must comply. As part of this compliance effort, power providers must identify their Critical Cyber Assets (CCA) and applicable corresponding Electronic Security Perimeters (ESP). This document provides a detailed methodology for determining ESPs and CCAs.

Overall Flow

As a concept, creating a list of CCAs and placing them within an ESP seems fairly straightforward because in most situations there is only one logical point at which to delineate a network and all stakeholders are in agreement. However, in some cases that point is not so clear, usually due to disparate stakeholder perspectives and needs: what a power engineer finds essential may be irrelevant to a networking technician. This document's intent is to provide an efficient, straightforward, repeatable method by which to view devices without the time commitment of a formal approach

The flow chart shown in Figure 1 follows the process outlined in the Critical Infrastructure Protection (CIP) standards to classify devices and construct ESPs—most devices can be classified by their nature. For example, it can be near-universally agreed that a remote terminal unit (RTU) is essential to the bulk of the electric system's operation; these devices will be classified early and with a minimum amount of work. If the device cannot be classified through any other means, a weighted score is generated based on a table of specific criteria; however, when applying this methodology, procedure should never replace professional judgment and adherence to good security practices.

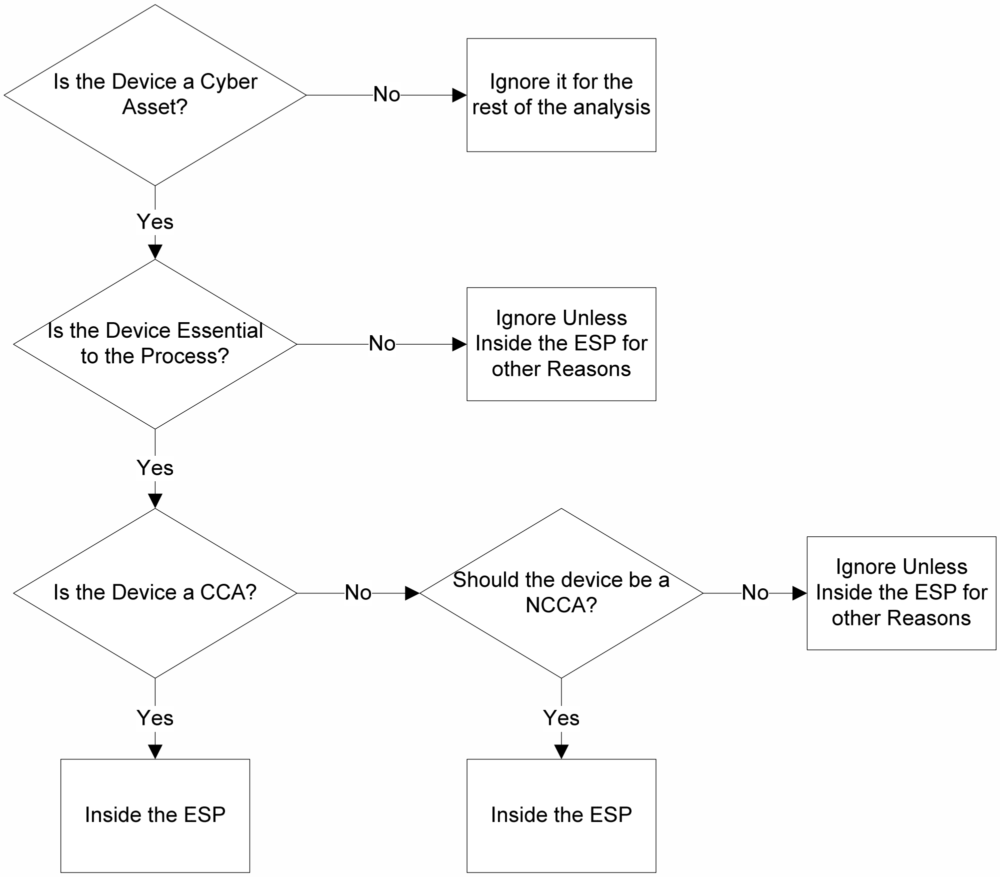
Is the Device a Cyber Asset?
No
Ignore it for the rest of the analysis
Yes
Is the Device Essential to the Process?
No
Ignore Unless Inside the ESP for other Reasons
Yes
Is the Device a CCA?
No
Should the device be a NCCA?
No
Ignore Unless Inside the ESP for other Reasons
Yes
Inside the ESP
Yes
Inside the ESP

Figure 1: Whether to place a CCA in the ESP

As shown in the flowchart, the first question to be answered is whether the device in question is a cyber asset; for most devices, this is an easy question. A Windows PC is definitely a cyber asset, but the wires running between two devices are not. The status of devices such as unmanaged switches is less clear, so this section describes a process with which one can make a clear, informed determination.

Once it is established that the device is a cyber asset, the next question is whether the device is essential to the electric system's operational process: a vibration controller is obviously essential because without it the generator cannot run. The CIP standards' intent is to help secure the electric system from cyber attack, and they define any device that gives an attacker control of the process as also essential to the process. Thus, leaving devices that have full control of the process unprotected violates the intent of the CIPs.

The three criteria for CIP-002 can then be applied to determine whether an essential device should be a critical cyber asset; however, a recursive problem exists within the CIP definitions. An ESP is supposed to contain all the CCAs, but if a device communicates outside of an ESP, that can make it a CCA. How does one determine the boundaries of an ESP if knowledge of the ESP is required to do so? Fortunately, there are a limited number of ways in which the logical border of most networks can be drawn, so one must consider only a few combinations. By definition, any device that is a CCA must be inside an ESP.

Finally, all devices that need to be inside the ESP—but are not CCAs—must be designated as non-critical cyber assets (NCCAs). While no CIP requirement exists that states these devices must be inside the ESP, to do otherwise will yield only short-term gains[1]. For example, some devices may be inside the ESP for purely practical reasons, such as a network printer. The printer is not essential to operational processes nor does it have control of any process, but putting it outside the ESP is impractical logistically. Identifying NCCAs minimizes the future cost of having to revisit or redefine the ESP's security boundaries.

Step 1: Determining Whether a Device is a Cyber Asset

Not all electronic components need to be considered as a device under the CIPs—the device must be programmable or configurable to qualify for consideration and can, in fact, be ignored during analysis if it does not. The primary question is “Can a hacker take over this device?” Sometimes the answer to that underlying question is ambiguous and depends on the professional judgment of an expert. Figure 2 provides a basis for decision making, but should not replace good judgment.

[1] There are several movements within the standards bodies to amend the language so that all essential devices are protected; only the final language's appearance is left to be addressed.

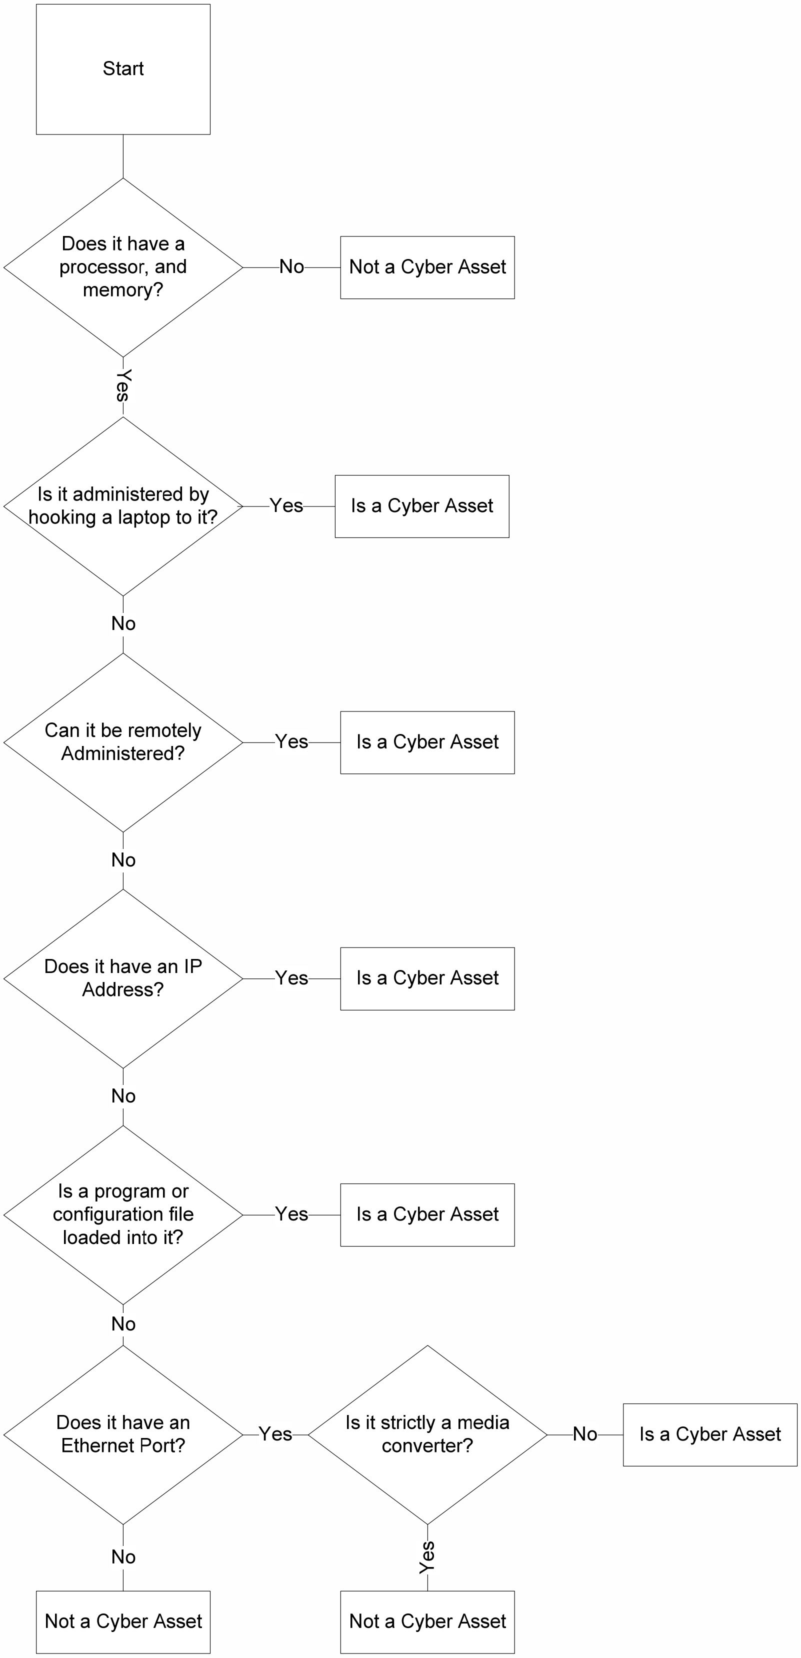
Start
Does it have a processor, and memory?
No
Not a Cyber Asset
Yes
Is it administered by hooking a laptop to it?
Yes
Is a Cyber Asset
No
Can it be remotely Administered?
Yes
Is a Cyber Asset
No
Does it have an IP Address?
Yes
Is a Cyber Asset
No
Is a program or configuration file loaded into it?
Yes
Is a Cyber Asset
No
Does it have an Ethernet Port?
Yes
Is it strictly a media converter?
No
Is a Cyber Asset
No
Not a Cyber Asset
Yes
Not a Cyber Asset

Figure 2: Determining a device's Critical Asset status

Does the device possess a processor and memory?

If the device doesn't have an obvious processor or memory then it does not meet requirements to be a cyber asset—for example, a cable is not a cyber asset. If the device's intelligence is unknown, assume it contains a processor and memory.

Does the device require a laptop be connected for administrative purposes?

If a device has a communication port and is programmed by way of a laptop interface then it is considered configurable and programmable. The existence of an interface port and software that communicates with the interface port represents sufficient criteria to be considered a cyber asset.

Can the device be administered remotely?

If the device can be remotely administered then it contains a communications stack and is configurable—this meets the minimum definition for a cyber asset.

Does the device have an IP address?

The Internet Protocol requires a device to exhibit intelligence and state—a device that has an IP address has at least rudimentary intelligence and meets the bar for being a cyber asset.

Is a Program or Configuration File loaded onto the device?

If the device has a configuration file then it possesses the rudimentary intelligence needed to parse that configuration file—this characteristic makes the device configurable and programmable, indicating a cyber asset.

Does the device have an Ethernet port?

The presence of an Ethernet port indicates either an intelligent, network-aware device as opposed to static device that simply moves packets—if the device does not have an Ethernet port it should be considered a static device.

Does the device function strictly a media converter?

The mere existence of an Ethernet port does not require intelligence. For example, dump hubs and media converters do not understand networking protocols; they simply shift bits around in memory. However, the device sometimes can have an embedded processor or an application-specific integrated circuit (ASIC) to help with protocol conversions since there is no significant way to interact with those internal processes.

Step 2: Determining Whether a Device is Essential

The CIP-002 does not elaborate on what it means to be essential to the operation of a critical cyber asset—this section is designed to provide some clarity on that issue. However, this process is not meant to replace the reasonable judgment of a professional; it is meant to drive the thought process through the necessary steps. Good judgment should always prevail.

At the base and obvious discernment level, if a device helps run the generator it is essential; a printer is not essential. If a device's essentiality is obvious, you may skip this decision process shown in Figure 3.

One thing to bear in mind when ascertaining a device's essentiality is the fact that attackers generally will take control of identical devices simultaneously and in a group. If an attacker is presented with a dozen front-end processors, he is likely to own them first and identify their utility later. Because of this, redundancy should never be a factor when determining whether a device is essential.

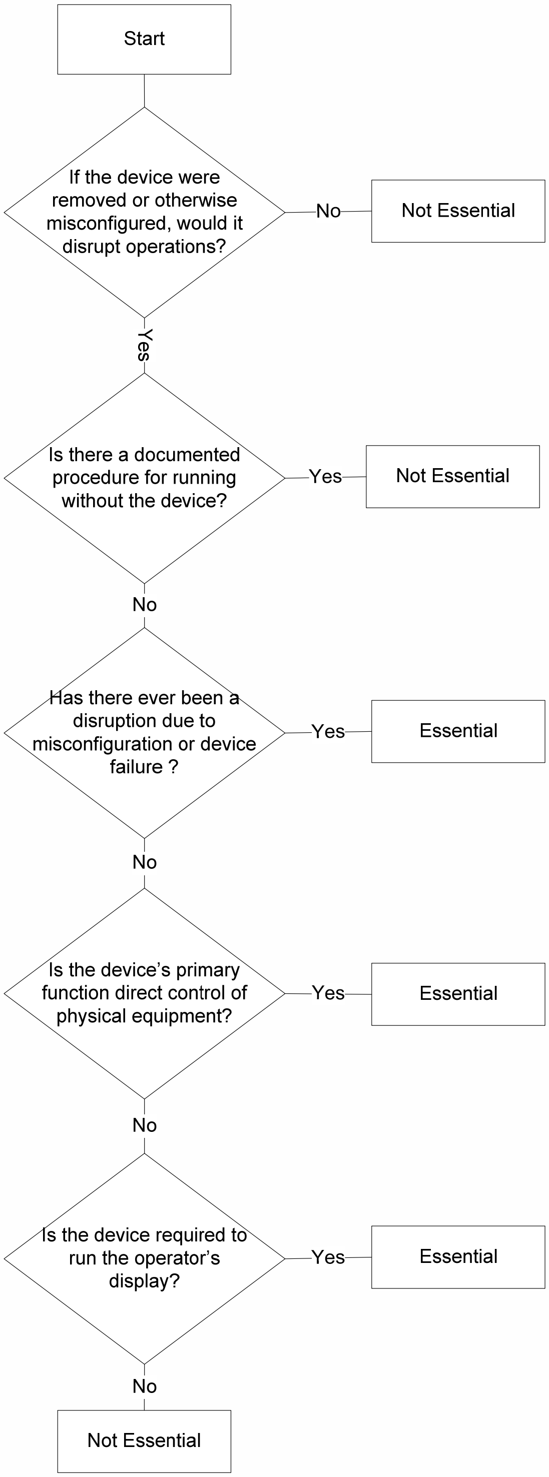
Start
If the device were removed or otherwise misconfigured, would it disrupt operations?
No
Not Essential
Yes
Is there a documented procedure for running without the device?
Yes
Not Essential
No
Has there ever been a disruption due to misconfiguration or device failure ?
Yes
Essential
No
Is the device's primary function direct control of physical equipment?
Yes
Essential
No
Is the device required to run the operator's display?
Yes
Essential
No
Not Essential

Figure 3: Determining a device's essentiality

Would the device's removal disrupt operations?

If the device and all its redundant counterparts were to be removed suddenly, would operations cease? If the operational process can continue without the device, then it is not essential. If the process needs to be shut down and then restarted in a different operating mode in order to function without the device then the device might be essential.

Is there a documented procedure for running without the device?

If there already exists a documented procedure for running without the device, then the device is not essential, indicating that in the case of a disruptive event, the process could be run according to the documented procedure.

Has there ever been an operational disruption due to misconfiguration or failure of the device?

Some devices are not essential based wholly on their face value; for example, disruptions have been reported around the unavailability of a SCADA system's central database. If there exists a history of the device causing operational disruption then that should be taken as proof that the device is essential.

Does the device control equipment directly as part of its primary function?

While the process may run satisfactorily without a local terminal or application server, any device that controls equipment directly should be considered essential. These devices are among an attacker's principal targets as their inherent function can include sending executable commands to the process; this includes most human-machine interfaces (HMI) and application servers. If a device generates or understands commands, it should be considered essential; if it simply forwards commands then it is not.

Is the device required to run the operator’s display?

When an operator loses their view it is considered a failure of the process and any device that is required to maintain an accurate view of the process should be considered essential. For example, an alarm server may not be essential to the process' operation, but alarms are essential for the operator to make informed decisions. Thus, deciding whether a device that supplies an individual metric is essential requires detailed knowledge of the process, and that is more information than can be covered in this document.

Step 3: Determining Whether a Device is a Critical Cyber Asset

After the device is declared essential, one should employ the three tests of CIP-002 R3 to determine whether it is also a critical cyber asset (CCA). If a device passes any of the three tests, it is a CCA and if it fails all three tests it is not a CCA. As quoted in the standard:

> *R3.1. The Cyber Asset uses a routable protocol to communicate outside the Electronic Security Perimeter; or,*
>
> *R3.2. The Cyber Asset uses a routable protocol within a control center; or,*

R3.3. The Cyber Asset is dial-up accessible.

The section R3.1 test should be ignored at this point—since the ESP has not yet been defined, it is impossible to tell whether a device communicates with a routable protocol outside the perimeter. This test will be revisited in Step 5.

The section R3.2 test contains two parts, the first of which determines whether the protocol is routable; Table 1 displays the status of the more ubiquitous protocols. If a protocol exists that is not represented in the table, an expert should be called in to make a determination regarding its status.

Table 1: Common protocols and their routable status

Protocol	Routable
TCP/IP	Yes
Modbus/Serial	No
DeviceNet	No
LNG	No
Ethernet/IP (CIP)	Yes
QNX	Yes
OSI	Yes
ICCP (TASE.2)	Yes
GE SNP	No
ISO	Yes
DNP/IP	Yes
DNP/Serial	No

The second part of R3.2 determines whether the device exists *within a control center*. The definition of a control center has been broadly interpreted and comprises almost any facility that contains control equipment. If the device is essential and uses one of the protocols listed in Table 1, it is a CCA.

The section R3.3 test assesses whether the device is *dial-up accessible*, which provides a narrow margin for interpretation: if the device is associated with a phone number that provides access to it, that device qualifies. Even if the device uses an internal phone system that cannot be reached from outside the network, the device still qualifies. If the modem is physically connected to the device and provides access to other devices, those secondary devices are not considered dial-up accessible. For example, if a modem is attached to a communications processor, the relays that are attached to the processor are not dial-up accessible.

Step 4: Determine Whether the Device Should be Inside the ESP as an NCCA

This process addresses potential future uncertainties in the standards themselves. The CIPs are vague on many points and are being updated rapidly—one of the only certainties is that the standards will be different in the near future. In general, the architectures that protect all essential assets will need little revision as the standards change.

This flowchart was designed to help provide a thought process for placing NCCAs inside the ESP—the goal is to minimize the continued cost of revisiting the architecture by placing the riskiest items inside the ESP.

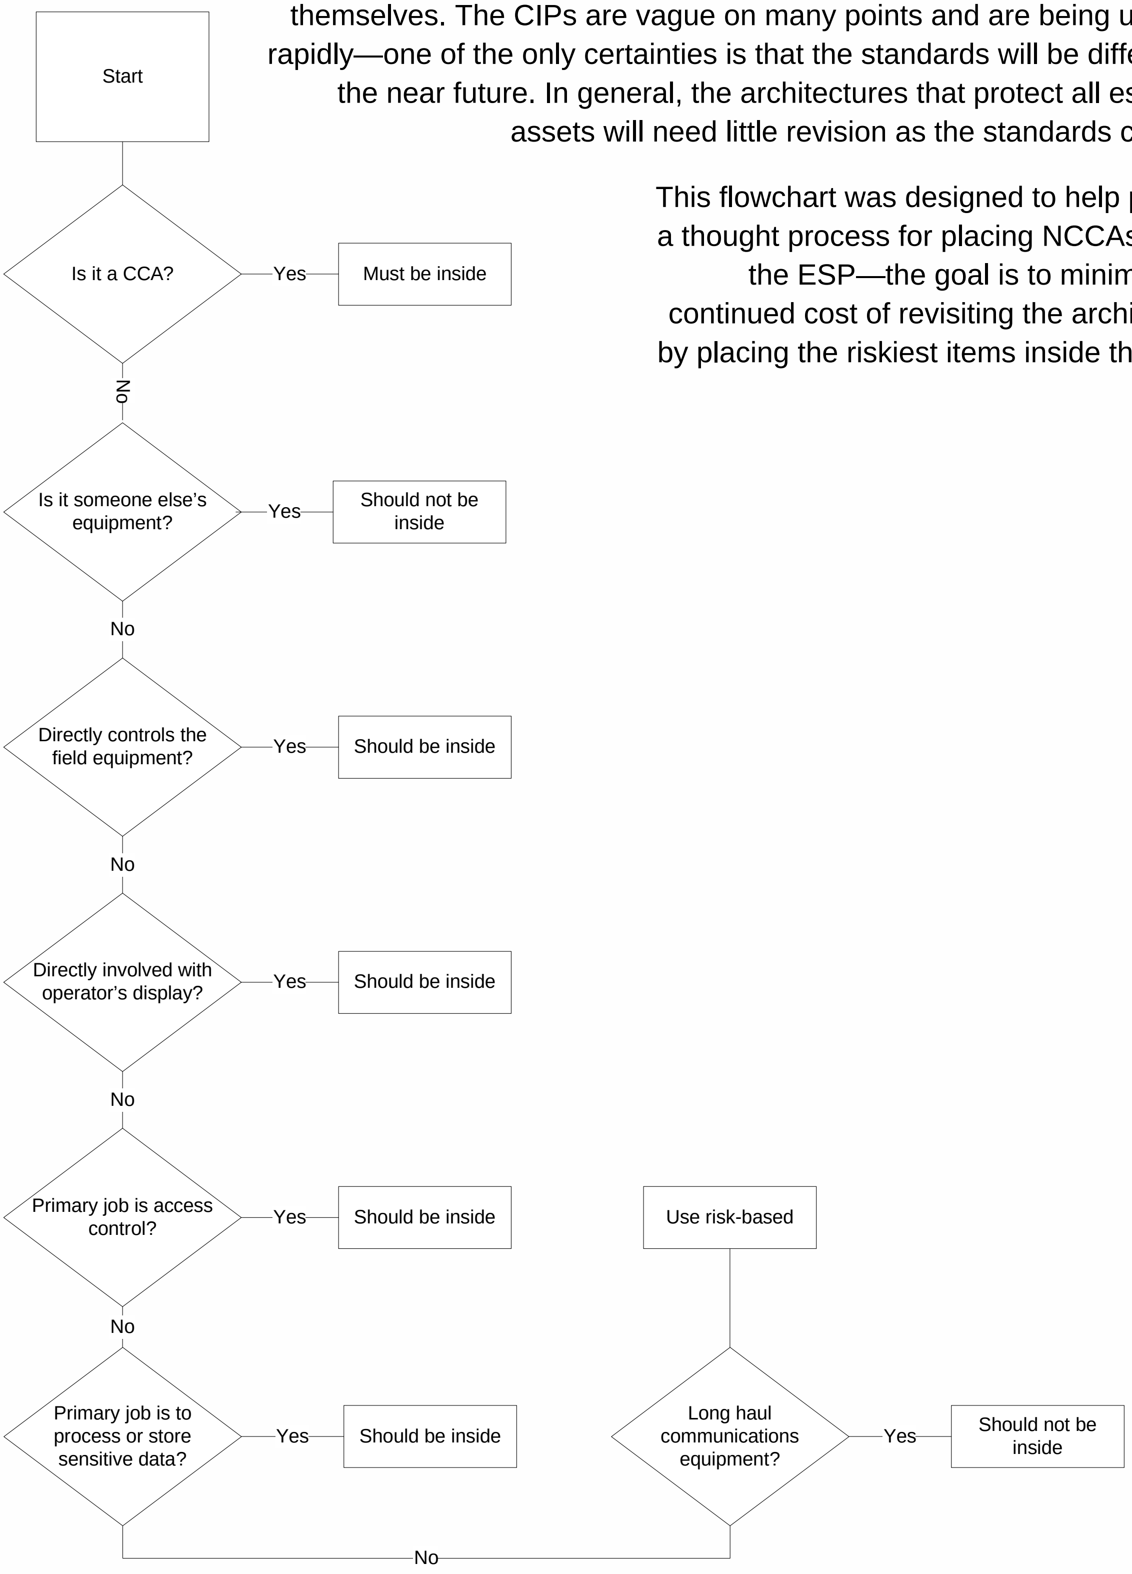
Start
Is it a CCA?
Yes
Must be inside
No
Is it someone else's equipment?
Yes
Should not be inside
No
Directly controls the field equipment?
Yes
Should be inside
No
Directly involved with operator's display?
Yes
Should be inside
No
Primary job is access control?
Yes
Should be inside
No
Primary job is to process or store sensitive data?
Yes
Should be inside
No
Long haul communications equipment?
Yes
Should not be inside
Use risk-based

Is the device a CCA?

If a device is already designated as a CCA then, per the standard, it must be inside the ESP.

Does the device belong to someone else?

If the equipment is owned, maintained, or administered by a third party then it probably should not be inside the ESP since good security practices dictate you should trust others less than you trust yourself. A good example of a third-party device are the HVAC systems, which are often monitored and maintained by third parties, but shouldn't be trusted on the network. For the most part, if an essential device is owned by a third party it should be placed in a DMZ.

Does the device control field equipment directly?

In general, anything that controls field equipment directly is already a CCA. If, for some reason, it is not designated as such, it should be included inside the ESP. A good example of this device type is equipment that controls a device that is not a CA.

Is the device involved directly with the operator's display?

If a device supplies data to the operator directly, but is not essential to the process, it should be included inside the ESP as an NCCA.

Does the device's primary job comprise access control?

The CIPs designate border access control devices (including firewalls) as access control devices and the ESP starts at the inside interface of the access control device; for example, a firewall is usually the first device outside the ESP. Any device used for access control that are not listed as access control devices under the CIPs should be positioned inside the ESP; Windows Domain Controllers are a good example.

Does the device's primary job comprise processing/storing sensitive data?

The CIPs require each utility to employ a process that protects sensitive data. Any device that processes or stores sensitive data and also is located on the same network as an ESP should be inside the ESP. Placing all of an enterprise's sensitive data-handling devices inside the ESP is understood to be impractical; however, where cost is not a significant barrier the devices should be protected.

Is the device considered to be long haul communications equipment?

If a device primarily is used to transfer data to a remote campus, it should not be inside the ESP. A future CIP standard will address communications equipment.

Using Risk to Determine a Device's NCCA Status

After completion of analysis, if uncertainty exists regarding whether a device should be located within the ESP, a simple risk score can be calculated—more complex risk-based analysis tools are available, but a simple total yields adequate results. Apply the scoring mechanism shown in Table 2; if the device's score is above 30, it should be considered a non-critical cyber asset.

Table 2: Evaluating a device's risk score

Condition	Weight
Device is commonly available	+10
Protocol used is common or documented	+5
Device is designated communications equipment	+1
Device is a computer	+5
Device is an embedded system	-2
Data from/to/through the device is used by primary control devices	+10
Devices has suffered known exploits in the past	+10
Protocol is custom	-5
Device is custom	-10
Device requires significant configuration before deployment	-2
Device configuration is shared with vendors, consultants, or other third parties	-2
Device processes or carries data from outside the ESP	-3
Device processes or carries data from outside the Utility	-7
Vendors or other third parties maintain the device	-3
Device is maintained electronically from outside the ESP	-3
Devices understands TCP/IP for either administration or primary function	+3
Device understands HTTP for either administration or primary function	+3
Device has a documented security assessment from a reputable source	-5
Device is deployed in a secure manner using a guide from a reputable source	-5

Step 5: Determine ESPs and Add Additional CCAs

Once a list of CCAs and NCCAs has been compiled, an ESP can be drawn around them and NCCAs can be removed from the ESP if the cost is prohibitive. In most cases there will be a natural demarcation between networks that works well for defining an ESP.

Only after the ESP is drawn can the CIP-003 R3.1 test be applied—any device that uses a routable protocol through the new ESP becomes a CCA. However, this does not apply to the access control devices at the network's edge since they are covered as access control devices and not as CCAs.

Conclusion

The intention of this document is to give energy providers a starting point from which to increase their security awareness and create a plan that hardens their critical asset security practices. It is not meant as a substitute for a detailed infrastructure assessment with attendant recommendations and guidelines.

If you are an energy provider or in charge of a similar SCADA-type infrastructure and would like more information on how you can protect your assets, IOActive can provide that level of guidance. We have performed security assessments for power/energy entities and worked with Smart Grid technology, so we understand the challenges that CIPS adherence can pose. IOActive's SCADA security assessments are built on information gained from direct penetration testing and architectural code review of utility and power control systems, as well as any related, third-party technologies.

Under the guidance of Jason Larsen—one of our nation's leading SCADA security experts—and Mike Davis—a recognized expert in reverse engineering metering platforms—IOActive technical deliverables offer unparalleled insight and our services facilitate deep-trust relationships with power and utility companies across North America.

For more information, please visit our website <www.ioactive.com>.