Topic 7

Incident Management

Dr Diarmuid Ó Briain

Version: 2.0

Ollscoil Teicneolaíochta an Oirdheiscirt
South East Technological University

Dr Diarmuid Ó Briain

Table of Contents

Illustration Index

1 Objectives

By the end of this topic, you will be able to:

- Analyse the benefits and drawbacks of different Operational Technology (OT)/Computer Security Incident Response Teams (CSIRT) structures.
- Evaluate the effectiveness of various OT-CSIRT roles and responsibilities.
- Design an OT-CSIRT resource allocation plan that optimises cost and performance.
- Develop a cyber Incident Response Plan (IRP) that is tailored to the specific needs of an organisation.

2 Introduction

OT-CSIRT are responsible for developing and implementing IRPs, conducting regular exercises, and investigating and remediating OT security incidents. The team should be cross-functional and include representatives from all relevant departments, such as IT, security, operations, and engineering.

The OT-CSIRT should be organised in a way that allows it to respond to incidents quickly and effectively. The team should have a clear chain of command and well-defined roles and responsibilities. Typical OT-CSIRT roles and responsibilities include OT-CSIRT Team Leader, Process/Control System Engineer, Network Administrator, System Administrator, Security and Legal Expertise, Public Relations Specialist, Human Resources (HR) and Vendor Support.

The OT-CSIRT should develop and implement policies and procedures to guide its response to incidents. These policies and procedures should be based on best practices and the specific needs of the organisation.

The OT-CSIRT should also develop a cyber IRP. The IRP should document the team's responsibilities, procedures, and communication plan. It should also include a list of key contacts and resources.

The IRP should include an overview of the team's goals and objectives, a description of the types of incidents that the team is responsible for responding to, incident detection and notification procedures, incident investigation and response procedures, a communication plan, forensics procedures, and a plan for exercising the plan.

The OT-CSIRT should also play a role in incident prevention, patch management, and vendor interaction during an incident.

Overall, the OT-CSIRT is responsible for ensuring that the organisation is prepared to respond to and recover from OT security incidents.

3 OT Incident Response Team Resourcing

The beginning point for creating a cyber-incident response capability is the planning and preparation phase. All the elements are brought together to prevent an incident if possible or to be ready to respond to one if it occurs. A cyber incident response capability consists of several core building blocks that include the organisation of the response team, establishing the organisation's policies and procedures, developing the response plan itself, defining reporting and communications within and external to the team, verifying that the plan works as expected, and then enabling state and status reporting to support the team if and when an event occurs.

Figure 1: Incident Management Stages

3.1 Organising the Team

The first step in developing an incident response capability is team organisation. Most groups are organised into a OT-CSIRT. The OT-CSIRT may be composed of specialists dedicated to this effort or part-time staff with other day-to-day responsibilities. In this topic, the OT-CSIRT will refer to the internal response team that is directly supporting the OT. Other external response teams are organised around specific technical areas or along geographical or organisational boundaries.

3.2 Team Responsibilities

The responsibilities of the OT-CSIRT will vary depending on the asset owner's organisational size and structure. The responsibilities also may be shared among different departments that have not traditionally provided support to the OT security team. Third party involvement can be employed through vendor Service Level Agreements (SLA) with equipment vendors or with consultants or other specialists. This option may be necessary for asset owners with limited resources. The OT-CSIRT's responsibilities will include:

- Acting as an expert resource on cybersecurity threats and vulnerabilities
- Serving as a clearing house for incident prevention, information, and analysis
- Developing organisational policies and procedures related to incident response
- Understanding safeguards on the OT
- Identifying operational impacts to the organisation in the event of an incident
- Creating and testing the IRP

- Acting as a single point of contact for all internally reported incidents or suspected incidents
- Responding to the incident when one occurs
- Reporting to key stakeholders and external agencies after the incident such as the National Cyber Security Centre (NCSC) and the Gardaí or police
- Gathering forensic information to support analysis and as evidence for legal actions
- Implementing safeguards to prevent a recurrence of the incident
- Remediating the OT after the incident.

3.3 Team Organisation

Various models have been identified for organising a OT-CSIRT. The most applicable OT-CSIRT model for OT environments is either a centralised or a distributed response team.

3.3.1 Centralised OT-CSIRT

A centralised cyber incident response team may be found in various size organisations and is made up of individuals with various backgrounds. Its distinguishing feature is the close geographic proximity to the OT. In this approach, servers, networks, monitoring equipment, engineering workstations, and the controlling devices connected to physical equipment's are all typically found at one facility. This single team works on site and handles all the incident response activities. This model is the recommended approach, where possible, because it will reduce the overhead associated with multi-team interaction and allow for onsite access, control, and analysis.

3.3.2 Distributed OT-CSIRT

A distributed response team may include a central OT-CSIRT, but because of the separate physical locations of the organisation, multiple teams may exist or be required. This model applies where facilities are spread across multiple regions, or countries and a single team would not be in a position to respond in a timely manner to any specific incident. It is also necessary in large organisations that are geographically dispersed, where the remote teams may include contracted specialists or even part-time staff. This approach requires more emphasis on communications and coordination between teams, but it also allows for a remote team to be onsite at the source of the incident. It is recommended that distributed organisations have strong centralised OT-CSIRTs with self-contained, individual OT-CSIRTs in the remote locations. Planning, prevention, analysis, and forensics can all come from the central group, allowing for efficiencies of scale. Incident response, however, must be a hands-on experience with the local OT-CSIRT taking the lead on an incident, with the support of the organisations central staff.

3.3.3 Key Considerations when organising a OT-CSIRT

Information Technology (IT) environments undergo dynamic change with commonality in network configurations, operating systems, and equipment. By comparison the OT environment tends to have static configurations typically consists of unique and even deprecated devices with site/operations-specific configurations. When dealing with a common item of OT equipment, its use, and the impact as a result of failure is almost always unique to the particular organisation. Unfortunately, this environmental knowledge is often limited to a few key control systems engineers. This aggravates the problem of attempting to provide continuous coverage with a limited pool of resources. If allowed to continue, it can result in employee burnout and higher turnover, both of which are detrimental because specialised knowledge is required to maintain and operate these systems. In organising the team, consideration must be made to assignments and may include delegation of as many tasks and responsibilities to non-key staff, or to subcontractors, as possible.

Staffing decisions must address division of authority. In IT, decisions typically roll up to a Chief Information Officer (CIO), Chief Information Security Officer (CISO), IT director, or equivalent. OT operational responsibilities will often fall on the plant manager who is highly sensitive to interrupting the process. The plant manager also may come from a traditional engineering background and not have adequate awareness of cybersecurity issues. Upper management may pressure the plant manager to prevent any work stoppages. An understanding, with agreed upon authority must be established between the OT-CSIRT, operations, engineering, and IT management prior to an incident. Each of these organisations can bring important knowledge and skills to the team, but the OT-CSIRT must have the proper level of authority from the beginning, otherwise, valuable time will be lost determining authority while plant operations are at risk.

3.4 Team Roles and Responsibilities

Though every organisation will not be able to staff each position directly, each role should be identified and assigned, even if it is part-time, with staff having multiple roles, or with personnel from the OT integrator or OT vendor/manufacturer. For larger organisations where the demand might be greater, or to ensure redundancy, it may be necessary to have several people assigned to a particular role. This is especially true for process and operations engineers with unique knowledge and experience. Each OT-CSIRT role is described below:

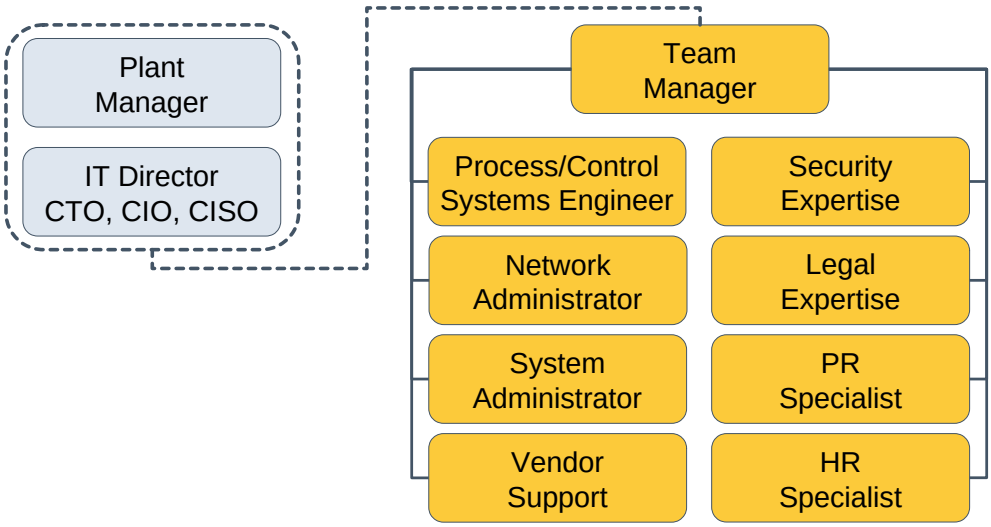

Figure 2: Incident Response Team

3.4.1 Plant Manager

The plant manager, including OT and Control Centre Managers, may not be involved in many of the details of the IRP, the plant manager must be involved in assigning authority to interrupt operations, being part of the risk assessment process when an incident is identified, funding OT-CSIRT tasks, and acting as a liaison to executive management and external parties, including the press.

3.4.2 IT Director, CIO, CISO, or Chief Engineer

This role is similar to the plant manager in terms of responsibilities. These two management positions are essential and must communicate and coordinate delegation of authority and what resources can and will be applied to an incident. A modern control system is typically integrated into existing IT networks, business systems, and communication equipment.

3.4.3 OT-CSIRT Team Manager

It is necessary to assign one person the responsibility of seeing that the team is organised and accomplishes its objectives. This person may act as a technical lead, or a separate technical lead may be designated from someone on the OT-CSIRT. The manager should have the authority granted by senior management to act in the best interest of the company. If functions of the OT-CSIRT are outsourced, then this person must oversee the actions, tasks, and contracts of subcontractors. This person is critical to assembling key resources to mitigate, contain, and resolve computer incidents in a timely and successful manner.

3.4.4 Process/Control System Engineer

This person should be the Subject Matter Expert (SME) on the control system architecture and should understand the system components and products being produced or supported by the OT. S/he provides important information on normal and abnormal equipment functions and functional cycles as well as the potential impacts when a component in the OT is removed from service. The process engineer is key player to the OT-CSIRT's understanding of how to resolve or work around equipment failures and how to resume operations when necessary.

3.4.5 Network Administrator

The network administrator can provide a key role in the OT-CSIRT if the incident involves a cyber attack originating from the computer network. This person will be knowledgeable on network access, including security vulnerabilities, patching, intrusion detection, and system monitoring. Knowledge and availability of activity logs from network switches, routers, and firewalls before, during, and after a cyber event are crucial in determining the scope and complexity of the incident and provide insight on how to resolve and remediate any vulnerability discovered. Most cyber related incidents will involve a network, and thus a knowledgeable network administrator is the key to finding and resolving an incident.

3.4.6 System Administrator

This is primarily the control system administrator, but it also may include IT administration because of the high degree of integration in modern organisations. The system administrator should have a high level of skill and knowledge relating to access permissions and system operation logs on affected servers. Administrators may be familiar with process control operations and operational cycles. These administrators should be aware of what is happening on their respective systems and should be cognisant of potential vulnerabilities. They also may interface with vendors and suppliers.

3.4.7 Security Expertise

Security expertise in this topic deals primarily with cybersecurity expertise. These individuals may play dual roles, but someone needs to be available with in-depth knowledge of vulnerabilities, exploits, prevention techniques, and especially an understanding of how to prevent incidents and how to recover if they occur. They also may, on occasion, be involved in supporting identification and prosecution of criminal activities.

3.4.8 Legal Expertise

Legal expertise is necessary in several areas including: ensuring compliance with all national, European and international laws and regulations; explaining what evidence is admissible when taking action; specifying how evidence can be collected; third-party maintenance liability exposure; and helping the team understand what pitfalls, such as privacy rights violations, should be avoided. These individuals can be very useful when the team is preparing the IRP, enabling state and status reporting, and in forensics and data collection. Larger organisations may have legal departments in house. Smaller organisations may require outside legal help, in which case, legal firms should be contacted that have had specific experience with incident response issues.

3.4.9 Public Relations Specialist

This person should be involved as necessary. He or she will play a critical role if the incident causes noticeable disruption to service or impacts the organisations ability to deliver a product. This can be especially important if the organisation supplies services directly to the public, such as in the generation of power or treatment of

waste water. This person is responsible for ensuring the appropriate information and messaging is sent to the public via the news media.

3.4.10 Human Resources Specialist

The Human Resources (HR) specialist will be involved in OT-CSIRT activity if the incident is being attempted or carried out by someone inside the organisation. Legal issues, policies and procedures, and punitive actions will typically be handled by this person.

3.4.11 Vendor Support Engineers

Because of the specific and essential knowledge held by the vendor's technical staff, individuals from the vendor facility should be identified that can provide technical support to the asset owner on the equipment and systems involved in the incident. These individuals can provide information and understanding that may not be found in the OT-CSIRT. For example, their expertise would be valuable in restoration of the asset and also for the creation of custom patches, if necessary.

3.4.12 Other Support Staff

Support personnel can be incorporated into the OT-CSIRT as additional expertise is required. These could include legal or Gardaí/police, computer forensics specialists, risk management specialists, database administrators, application developers, platform specialists, and governmental agencies if warranted. For daily tasks like organisation support and scheduling or preparing policies and procedures, secretarial and technical writing personnel are valuable.

If the OT-CSIRT model is distributed, as many of the above-mentioned roles as possible, should be filled at the central office with specific technical staff available at each remote location. At a minimum, someone with process engineering, system administration, and network experience should be available at each distributed location. Communications must remain effective and reliable when an incident occurs, recognising that the incident itself may disrupt normal communication paths.

Logistical elements will not be discussed at length, but recommended infrastructure for the team would include some type of permanent or temporary incident or "war" room mobile communication devices, laptops, and available documentation, including policies, plans, procedures, phone lists, etc., all residing in locations that are less likely to be compromised by an incident.

While the primary focus of the OT-CSIRT is to handle cyber-related incidents, the response team could be used for non-cyber events such as OT or SCADA system outages, catastrophic equipment failure, or natural disasters such as floods or hurricanes.

3.5 Setting Policies and Procedures

While having policies and procedures are important in most business functions, incident response is important because decisions are being made under pressure of production stoppage, high financial cost, often at the most inconvenient times, and in situations where those with authority may not be readily available. Development of procedures and supporting policies while team members are not under pressure is crucial. At that time, team members can discuss and weigh options, test the approach, analyse impacts and alternatives, and obtain management input and approval. Many types of general cybersecurity policies are valuable for both IT and OT protection. In the context of this document, policies related to incident response should be established and published within the OT organisation.

Clearly written, detailed operating procedures should be developed to implement the incident response policy. The procedures found in an IRP are similar to those found in non-cyber emergencies and should be tested before the event occurs. Problems in the mechanics, accuracy, and timeliness of the procedures should be discovered during the development phase, when adjustments can be made, rather than in the middle of an actual response.

The initial IRP should direct the establishment of the OT-CSIRT and lay the foundation for the IRP. The IRP should define the authority of the OT-CSIRT. The policy forms the backbone for the procedures and actions defined in the plan. Although many additional security-related policies exist that should be considered, those that relate more directly to OT are as follows:

3.5.1 Human Resources

Policies should be included that address actions taken against employees or contractors when the incident is caused by someone inside the organisation. These would apply to immediate response and actions during a discovered incident, how the investigation is conducted, and any related punishment policies.

3.5.2 Information Disclosure

Policies must be defined to address the organisation’s position on disclosure, and what actions it will take in the event of an information breach. Policies should include who to contact and what time constraints exist on reporting. The plan must address information that may be stolen and potentially sensitive. This may include security classification levels, private personal data, business or engineering process information, or even vendor proprietary data or code that may reside on a control device.

3.5.3 Communications

If an incident occurs, policies should be in place regarding media interaction and communications. The policy should define who will speak on behalf of the organisation. It also may define interaction with vendors and customers.

3.5.4 Authority Assignments

As already mentioned, in the OT environment, a tendency exists to have dual organisational responsibility. The plant manager is responsible for operations, and the CIO is primarily concerned with the networks and computer-related equipment connected to or even used in the OT. Policies should address escalation lists and division of authority as well as delegation, including backup, when a specific manager is not available.

4 Building the Cyber Incident Response Plan

The cyber IRP establishes and documents the procedures and actions that implement the incident response policy for the OT. It defines the security incident and outlines the steps that should be taken to respond to the incident and mitigate damage to the organisation. A variety of IT-related IRP templates and examples are available, some of which are included in the references. They can serve as a good starting point for building the plan. The following key sections should be considered when creating the plan.

4.1 Overview, Goals, and Objectives

These sections of the plan define what will be accomplished. In these sections, the organisation can provide direction and guidance for overall business objectives in comparison to the response options to the incident.

4.2 Incident Description

Many IT-type incidents are fairly easily classified. These include Denial-of-Service (DoS) attacks, unauthorised access to networks, accessing protected and private information, defacing web pages, misuse of services, etc. In the OT environment, clear definitions of what is a security incident must be identified and communicated to the extent possible. This is particularly important when considering if equipment failure or unexpected software behaviour is caused by a cybersecurity incident, due to mechanical failure because of wear, environmental conditions, or other non-security related factors. It is important to understand and differentiate between a cybersecurity and non-cybersecurity incident. If an isolated case of equipment or software failure exists, a replacement may resolve the problem. If the failure is the result of a compromised vulnerability, corrupt or untested patch deployment, or if malware remains somewhere on the system or network, then the original, or other similar, equipment may be at risk. Replacing the hardware or software will not resolve the problem or prevent reinfection if the configurations are the root cause of the incident. Accurate descriptions of an incident will also prevent unnecessarily activating the OT-CSIRT.

With OT, differentiating between cyber-based incidents and those caused by other sources is critical. For example, the reaction to equipment damaged by a disgruntled employee with a crowbar would be vastly different than damage to the same piece of equipment caused by an unknown attacker who manipulated controls on the equipment. It’s important to identify and define each incident type so that the appropriate response can be followed for that unique situation is important.

4.3 Incident Detection

This is also called *discovery* and includes ways in which an incident is identified and reported. While few cases of obvious incidents (an intruder is found logged onto the OT network or a website is defaced) exist, detecting most incidents will require automated analysis tools, system behaviour patterns, and an awareness of what to

look for among operators, supervisors, and other staff. The operators and the process engineers are usually critical to detection of unusual operations and are the first to note a difference in system behaviour. This difference is the key to understanding what is happening in the OT. The IRP must address automated systems, expectations for staff, contractors, and partners when suspicious activity is detected; and procedures for help desk and call centre staff.

4.4 Incident Notification

Once an abnormal event is identified, it needs to be prioritised to determine the cause and whether this is a minor system event or if it requires immediate escalation. This section of the plan should identify the contact information for incident reporting. The section should include basic work phone, mobile phone, e-mail, instant messaging, and pager information for internal staff, including system and network administrators. It also should address the following circumstances:

- After-hours phone and pager
- Offsite contact numbers
- Contact information for customers and partners
- Phone or pager numbers for backup staff
- Contact information for management and rules for escalation
- Criteria for filtering out false positives
- Contact information for any relevant regulatory authorities
- NCSC contact numbers and information
- Vendor/integrator responsibilities and contact information.

This contact information should be publicised to everyone that might identify a potential incident. A weekly and monthly duty call list issued to operations may be of help to let all employees know who is available to call for assistance in the event of a cyber incident. Because external agencies may be reporting a potential incident, based on events at other sites, the contact information should be available to all necessary external organisations as well.

4.5 Incident Analysis

Procedures in the plan should address how to evaluate and analyse a reported incident. The incident might be reported by internal or external sources and could happen at any time. In this stage of incident management, those receiving the report must determine:

- What dangers or effects on the facility or facility personnel safety may be caused by the event
- If the reported incident is real or a false positive
- What stage the incident is in—beginning, in process, or has already occurred
- What the impact might be to the organisation
- The specific type of incident

- What systems and equipment are or may be affected by the incident
- If the system has failed over to an available backup system
- If the incident has the potential to spread across other networks or even outside to partners or customers
- What organisations will be affected and who should be part of the response.

4.6 Response Actions

This section is essential to the plan because it defines the procedures to follow for each type of incident detected. An incident will typically occur at the most inopportune time; there will be increased stress and pressure on staff, little time for testing options, and every action will be watched and measured by upper management, stakeholders, and perhaps even by the public. It becomes essential that well thought out actions be defined and tested before the incident occurs. When defining the response actions, consider the following:

- The response must be directly associated with the incident type; one approach will not fit all situations, and new attack vectors should be considered on a regular basis.
- The plan must account for contingency situations including nights, weekends, holidays, unavailable staff, and non-functioning communications equipment. External factors affecting the plan, such as deliberate or accidental power loss, also should be addressed.
- The actions identified in the plan must include a comprehensive response covering containment of the problem, restoration of operations to a functional state, and prevention of a reoccurrence. As mentioned above, the actions will be dependent on the type of incident and its severity.
- The response procedures should be tested in a situation as realistic as is practical to determine elements that were missing, misunderstood, incomplete, or inaccurate. Corrections can be made and then retested until all concerns have been addressed.
- The response actions must be weighed against business impact and approvals secured while in the planning stages. Some remediation activities may cause more harm to the business than the incident itself.
- All available perspectives should be involved in preparing the plan. This includes technical, legal, communications, management, operations, engineering, and human resources.
- The actions must take into consideration any forensics requirements. It will not be necessary in all cases, but some incident types will require that the procedures accommodate the need to identify and preserve information for potential criminal or other legal actions.

4.7 Communications

While elements of communications could be included in the response actions, the topic is unique enough that it could be addressed in a separate section in the IRP. The communications section should include:

- Lists of all necessary contacts in the media, emergency responders, civil authorities, and local and global organisational contacts.
- A designated point of contact with one or more alternates who are prepared to speak for the organisation when an incident occurs.
- Prepared and vetted statements and press release information that would be available for immediate use. This is particularly important when the organisation provides a product or service on which the public depends.
- Reporting chains both internal and external to the organisation.
- A current list of contact names with the respective skill sets at key vendors for critical systems and components in the overall OT.
- A description of alternate physical methods to handle impaired communications through the telephone lines, cellular networks, or the internet. This would include contingencies if any or all the methods were non-functional.

4.8 Forensics

Cyber forensics focuses on collecting, examining, and analysing data related to an incident along with protecting incriminating evidence for use in legal action against a suspected offender. This data can be found in available logs (network, server, and workstations), physical components (hard drives and bitmap images of affected Real Time Operating System (RTOS) if possible), emails, voicemail, texts, and telephone records. While the information gathering can be useful in understanding the incident and helping in preventing further actions, the approach has nuances related to data integrity and protection that go well beyond just learning about an incident. A recommended practice is available that focuses completely on cyber forensics related to OT. This recommended practice should be consulted when preparing the forensics section of the IRP [1].

4.9 Exercising the Plan

Although it may be inconvenient and disruptive to plan for, conduct and evaluate the results from an incident response drill; considering the stakes involved, it is essential. Even the best response plans cannot anticipate all the obstacles that will be faced when a real incident happens, nor can they anticipate, in all cases, how people will react to unforeseen situations. The people who were expected to be available and fill certain roles will often be inaccessible. New people may have replaced previously trained workers. Unanticipated events may occur where decisions need to be made with little or no time for analysis.

Many problems that would occur in a real incident also will be present in the test exercise or drill. This means that an opportunity is available to review, analyse, and change the procedures without suffering the effects of catastrophic decisions or even lost production. This is only true, however, if the plan is tested in an environment closely replicates the production system.

To conduct partial tests of the IRP is also productive to evaluate unexpected behaviour. These partial tests allow adjusting and making the plan more effective and

streamlined prior to a full test. Partial testing can be a good training exercise for new OT-CSIRT members without incurring the cost and disruption of a full test.

The following are items that may be considered when setting up the incident response simulation.

- Some aspects of the IRP will be similar for all incident types, but others will be vastly different. Different incidents may require different levels of response, for example, an intruder scanning the OT network but not altering equipment settings would require a lower level of response than someone overriding safeguards to lock up pumps or valves that control the processing of toxic chemicals. The drills should address as many critical scenario types as possible and the nature of the drill adjusted accordingly.
- The exercise should mimic real-world conditions as much as is practically possible in order to discover weaknesses in the IRP. The closer the exercise is to the actual circumstances of the operating environment, the more problems will be found and resolved before a real event occurs. Actual equipment should be used if possible in order to gain accurate insight into how the IRP plays out. This may mean working with a vendor to provide temporary equipment specifically for the exercise.
- The drill should simulate worst-case conditions. An intruder who is intent on causing the most damage possible or who is seeking widespread publicity may intentionally strike at the worst possible time. Depending on the desired outcome, this may be at the peak of the workday when the maximum numbers of people are on site, or it may be in the middle of the night on a weekend or holiday when key technical staff and decision-makers are gone.
- The drill should involve all those who may be involved in the response and mitigating efforts. Having trained one set of people will not be helpful if the actual workers that face the incident are not knowledgeable on what to do if an event happens on their shift.
- Drills should be held on a regular basis to accommodate staff changes, changes in the facility or equipment, and new information gained from previous drills and actual events.
- Circumstances surrounding the drill should be designed to cause the staff to think through unusual situations. This can reveal weaknesses in the decision-making process and potential unintended cascade effects and consequences.
- The OT-CSIRT should, wherever possible, draw upon the experience of other facilities in preparing for the drills and potential incidents.

4.10 System State and Status Reporting

Enabling system state and status reporting refers to associating automated mechanisms with the hardware or software that report information about the system, including abnormal behaviour, intrusion attempts, or any other data that would be useful in detecting an incident, understanding impact, and quickly supporting resolution. Examples include network logging and database auditing, customised applications developed in-house for specific networks or equipment, or vendor-developed capabilities built into supplied equipment.

When programmers apply state and status reporting to software applications (or build code into the program solely to provide status or state information unrelated to its intended purpose), it is almost always done to help debugging the program or in providing support information if problems are reported. When considering justification for expending resources to enable the system, consider that it can be helpful for resolving any type of system problem, including debugging software, detecting pending equipment failure, or just improving efficiencies in the work processes.

Adding code to report state and status information can be very valuable in supporting forensics after an incident has occurred. However, its primary purpose is not forensics, but rather incident detection and resolution.

While there are real advantages to enabling status information about the OT, challenges exist. Because of the nature of OT, many devices are designed with volatile memory, the base code may be difficult to access. Vendors may be reluctant to add new code because of cost or risk. In addition, data that is generated and available is often replaced so quickly that log data cannot be written or stored in a practical way. Network traffic loads can be affected by the additional logging, even to the degree of altering or impairing normal operations.

A variety of ways to approach automating system components are available to collect useful information. Several key types of approaches are:

4.10.1 Networks Intrusion Detection Systems (NIDS)

These applications, which include both hardware appliances and software solutions, reside on the network and are useful in detecting attempts to access the network. They have been around for many years in IT and are equally useful in the OT environment. A NIDS will act to alert the network administrator of intrusion attempts and record all alert information, according to parameters set by the administrator.

4.10.2 Protocol-based Intrusion Detection System (PIDS)

A PIDS is associated with a component rather than the network. Typically it would reside between a server and a connected device and analyse communication protocols between the two. A variation of PIDS is the Application Protocol-based Intrusion Detection System, which is placed between several servers, all communicating with application-specific protocols.

4.10.3 Host-based Intrusion Detection System (HIDS)

An HIDS resides on a host system and analyses data unique to the applications on the host. It may include analysis of log files, file systems, database changes, etc.

4.10.4 Intrusion Prevention System (IPS)

Because of the immaturity of IPS technology and the high risk of inadvertently causing OT failure, these systems are not currently recommended for OT environments. They are mentioned so-as to provide a more comprehensive understanding of available technology and for their potential role in integrated business systems. An IPS is similar to an IDS with the exception that it actively reacts to malicious activity and blocks or prevents the activity if possible. If implemented in the OT environment, both the NIDS and IPS will be most closely associated with the network with some limited

application to server-type components. Detailed information is readily available on the Internet for IPS and IDS-type products. Extensive preliminary testing to ensure OT compatibility is highly recommended before system deployment. An active system like the IPS can prevent legitimate activity, so the establishment of approved activities is critical before this approach can be used.

4.10.5 Network and Device Logging

Mature products are available on the market for network logging including the IDS types mentioned above. This is not always the case with the variety of control system devices being used. Device logging will vary based on age, vendor, device type, and available settings. Administrators should enable auditing and logging capabilities whenever they are available and in circumstances that will not interrupt operations. Vendors should also be encouraged to provide self-monitoring capabilities with new products or upgrades to existing hardware.

4.10.6 Configuration of Data Generators

Several key elements should be considered when using commercial systems for successful data gathering. It is important to know and understand all settings, properly configuring the device, and regularly monitoring alert notifications. A perfectly operating detection system will be of no use if an alert is sent but no person receives or acts on the notice. This can be the case when a duty officer has not been assigned or when so many false positives are published that actual incidents may be easily overlooked. For customised logging and monitoring, having useful settings will increase the value of the device. For example, the state of field devices in normal operations may be measured and reported to a server on a constant basis. The server may have an ongoing test for out-of-range conditions or unusual traffic, which would be reported via e-mail, pager, alarms, etc. The key is to analyse the specific devices involved and apply either vendor provided or custom-monitoring capabilities to the device. With custom monitoring, no direct access to devices may exist, especially those that are older or have proprietary software. In these situations, monitoring internal to the device may not be available, but there may be an opportunity to test signals going to and from the device. External ways to accurately validate the state and status of the component may be available. Differences exist between the OT and IT systems in regard to network traffic. Because OT traffic is limited and specific, as compared with the IT systems, signatures can be created based on what is outside of range or is abnormal after the baseline has been taken.

Take care when enabling state and status reporting on the OT because some systems and applications have the potential to introduce operational issues. For example, some legacy control systems can be disabled or shut down because of the very intrusive nature of some IDS and antivirus tools. Poorly configured IDS and antivirus tools have slowed critical data communications to the point the OT becomes inoperable. Any plan to deploy these tools must be checked with the OT vendor and tested for compatibility with the OT and additional supporting applications that are co-resident on these systems. Some newer software may be incompatible with existing support software (various Java versions for example). Questions regarding these types of systems include:

- Where will the log files be stored?
- How long will the log files be stored?
- Will older log files be deleted or archived?
- What parameters are being investigated? (Ports, login/logout times, abnormal traffic cycles and times, etc.)

5 Incident Prevention

Preventing a cyber incident is preferable to responding to one, but prevention takes on a whole new dimension in the OT environment. This is because compared with typical IT, beyond the network there are far fewer, and in some cases, no detection capabilities available in system devices. In addition, working components may have vulnerabilities that may never be fixed, and the results of the most severe attacks could include injury, loss of life, and severe financial loss. Because the relative vulnerability and consequences are both high, the facility should put sufficient resources into incident prevention.

5.1 Patch Management Considerations

Patch management is only one of many areas of consideration in an effective cybersecurity program. Patch management and vendor interaction are specifically highlighted in this document because of the unique requirements related to OT. Patch management is important to incident response in two ways. First, and foremost, it is an essential means of preventing an incident from occurring. Second, patching is a way to respond to vulnerabilities and prevent reoccurrences of the exploit. Without patching, systems can be left in the same vulnerable state they were before the incident. The following issues related to patch management of OT must be considered:

- Difficulties in scheduling maintenance windows on production systems to perform the patch
- Equipment that is no longer supported and no patches are available
- Patches that were issued by a third party—not the original vendor or supplier
- Testing of a patch in a non-production environment before implementing it on the production systems, especially where equipment is unique and expensive
- Creating a test bed or simulated environment
- Creating a viable backup of the system configuration as a Disaster Recovery (DR) point of the working system, if the last known good configuration needs to be deployed
- Development of patch roll-back procedures, should it be discovered that a patch interferes with proper OT operation
- Patches that cause issues with adjacent applications in the OT
- Receiving patches from vendors in a timely fashion
- Accepting the testing processes used by the vendor, including both unit and integrated system tests

- Assuming the risk that the patch will not bring down or impact the production system
- Knowing the time it takes to deploy the patch, or knowing how long it takes to remove the patch if necessary
- Working with and patching software embedded in OT components.

5.2 Vendor Interaction Considerations

The need to work with vendors on cybersecurity is important in the OT environment because of the proprietary nature of the software, the lack of maturity of the industry in relation to cybersecurity, and the more limited customer base of the vendors.

The business IT model has literally millions of users for a very limited number of operating systems, (or networking) vendors. This means that vulnerabilities in a single product such as GNU/Linux or Microsoft Windows, would impact nearly the entire customer base. Patching processes are more mature and well established, and vendor response is expected, even taken for granted. In addition, support for a single product or version of a product can be withdrawn with the expectation that the customer will upgrade to later versions. This is accepted practice in the IT environment and is often desirable so that new features are available to the customer base.

A single vendor is selling numerous products in the OT model, and of those products, many versions are actively being used in the field. These products can have a long service life extending 20 years or more. In addition, the number of customers is relatively small when compared with products in the IT environment. In some cases, there may be only tens or hundreds of customers, depending on the product, its age, and how unique it is. With the pressure on vendors to support multiple products and versions of products, and with smaller numbers of customers demanding a fix, vendors cannot guarantee provisions for patches, state and status reporting, or fixes in a timely manner, if at all.

To ensure the highest degree of both prevention and response targeted interaction between the customer and the technical staff of the vendor are significant. A unified voice of all customers will be helpful in putting pressure on the vendor to address security issues with appropriate patches. SLAs must be established with vendors to ensure ongoing patches and related support. These agreements should not be allowed to lapse, or legal influence will be lost. Customers also can provide direction on priorities and customer needs. From the perspective of the product user, this would involve user groups and provide ongoing feedback to the vendor’s technical and sales staff.

When responding to an incident, the relationship of technical or support staff at the vendor site is critical. Depending on the criticality of the OT component, it might be necessary to include the vendor’s technical personnel as an extension of the OT-CSIRT or even part of it. This means that names, expertise, and contact information should be maintained. These people should know that they may be called on to assist in the event of an emergency. This arrangement may require contracts with SLAs that define what help can be expected and what the cost for that assistance would be. When an incident is happening, it is too late to be trying to set up new contracts with

vendors. It is not practical in all cases, but advisable where possible, to include a turnaround time for patches or fixes in the agreement.

6 Incident Management

This section discusses the four key primary activities: detection, containment, remediation, and recovery and restoration related to managing a cybersecurity incident.

6.1 Incident Detection

Detecting an incident early will help to limit or even prevent possible damage to the OT and reduce the downstream efforts to contain, eradicate, recover, and restore the affected systems. This section focuses on the methods of detecting cybersecurity incidents by discussing warning signs to indicate when a cybersecurity incident is pending, how to categorise and prioritise cybersecurity incidents and responses, and recommended detection steps.

6.2 Detection by Observation

Two general approaches can detect an OT cybersecurity incident. The first is through user observation of abnormal system or component behaviour. An observation can come from any member of the organisation, including operators, process engineers, or system administrators. The second is through automated detection via applications or routines, such as network monitors, network traffic analysis applications, IDSs and antivirus programs that can detect and flag malware, intrusion attempts, policy violations, and exploits, as well as component failure. These automated approaches still require some human interaction for configuration, review, analysis, and action.

The approach requiring user observation is essentially an after-the-fact approach and can carry a number of adverse risks. After-the-fact means that an intrusion and cyber attack is currently taking place or has already occurred. Thus, this method provides no initial protection or prevention capability to a cyber incident. Some of the adverse effects associated with this approach are listed as follows:

- Damage to the physical system or equipment
- Extraction of critical control system operations data
- Alterations to the software configuration algorithms to produce future undesired system actions
- Injection of malware, such as viruses or worms, which compromises the confidentiality, integrity, and availability of the system or system data.

Every effort must be made to identify warning signs that could be observed prior to a system or equipment failure. Means other than a cyber attack can trigger many warning signs, but they are still worth considering as possible precursors to an incident.

The following list of symptoms to be considered as possible indicators of an attack [2].

- Unusually heavy network traffic
- Out of disk space or significantly reduced free disk space
- Unusually high CPU usage
- Creation of new user accounts
- Attempted or actual use of administrator-level accounts
- Locked-out accounts
- Accounts in use when the user is not at work
- Cleared log files
- Full log files with an unusually large number of events
- Antivirus or IDS alerts
- Disabled antivirus software and other security controls
- Unexpected patch changes
- Machines or intelligent field devices connecting to outside Internet Protocol (IP) addresses
- Requests for information about the system (social engineering attempts)
- Unexpected changes in configuration settings
- Unexpected system shutdown.
- Other possible indicators of a cyber incident include:
- Stoppage or displayed error messages on a web, database, or application server
- Unusually slow access to hosts on the network
- Filenames containing unusual characters or new or unexpected files and directories
- Auditing configuration changes logged on the host records, especially disabling of auditing functionality
- A large number of bounced e-mails with suspicious content
- Unusual deviation from typical network traffic flows
- Erratic OT equipment behaviour, especially when more than one device exhibits the same behaviour
- Any apparent override of safety, backup, or failover systems
- Equipment, servers, or network traffic that has bursts of temporary high usage when the operational process itself is steady and predictable.
- Unknown or unusual traffic from corporate or other network external to control systems network
- Unknown or unexpected firmware pulls or pushes.

This list provides examples of symptoms to monitor for but is not exhaustive. It is recommended that a proper operational state be understood and documented if possible. Any deviation from the expected functionality could be considered a warning.

Operator experience may be the best source of detecting deviations from normal, because subtle differences in equipment behaviour may create a *just doesn't feel right* situation that is difficult to identify. Very experienced operators will know when things are not working right and can detect potential cyber problems as well as non-security related equipment wear and tear.

Management should provide specific contact and reporting instructions to operators and any other plant personnel that may be in a position to detect unusual system or equipment behaviour. This should include pager, phone, and e-mail information to allow the operator to contact the OT-CSIRT. These instructions should also include a checklist of information to gather and report to assist the OT-CSIRT in analysing and accessing the unusual behaviour. The contact information and checklist instructions should be posted in convenient and easily accessible locations.

6.3 Automated Detection Methods

Automated methods of incident detection can be extremely valuable in preventing exploits to the OT. The nature of attacks, the number of attempts, and the round-the-clock timing of the attempts create an environment where manual observation is very difficult, if not impossible. Most networked OT of any substance will have some type of automated detection capability. This may include sophisticated, commercial IDS attached to the OT networks or it may be simple firewall logging. It is essential that a proper balance of automation for the application be configured properly, be working as intended, and include the appropriate human review and interaction.

The concept of system state and status reporting puts emphasis on using both commercial and customised methods to let the components of the system report on status and state information. This information is useful in preventing an incident, but is also valuable in post-incident analysis and forensics.

All automated detection systems have at least three components that are necessary for them to work properly:

i. A programmed method to detect an out-of-range or targeted event. This may include the detection of a character string that matches a known virus signature or certain network behaviour such as a denial-of-service attack. It also may detect attempts to access certain restricted ports, or it may recognise a known rogue IP source address. With individual OT components, it could be a customised application that detects when the equipment or software behaviour goes outside pre-set thresholds.

ii. The ability to capture and report the event or change. Detecting an event is the beginning; but to be of value, the application must organise and present the data in a useful format. More advanced systems will include filtering and reporting; others may just write log information to a text file. To be useful, specialised components must be able to write out or state changes in some form of audit or log file. Some processes cannot continuously be writing out a constant flow of log data without affecting equipment operations. In these cases, the ideal situation would be to set ranges and report only when outside the range.

iii. Communication of flagged events to an operator. Some sophisticated systems like an IPS may be able to take some preventative actions without human intervention. However, the IPS is designed for well understood IT applications and not for a production OT where inadvertent shutdowns could have undesirable results. In a more typical situation, a human must be involved to decipher false-positives and to separate maintenance issues from potential cyber attacks. The human also must be able to respond to the data and initiate the appropriate response, including activating the OT-CSIRT when necessary.

Each of the three components of an automated detection system must work properly or the system will fail. While the first two items have certain limitations, the major challenge seems to be with the practical aspects of the third, related to human observation and response. Some of the most significant challenges are related to availability, training in finding real events, and initiating a proper response when an actual event is discovered. The suggestions provided below address ways to support OT personnel:

- Use centralised logging that consolidates a variety of data sources, allowing administrators to see a unified set of information presented in one place and in a consistent format. This may require interfaces to and data pulls from log files or audit tables.
- Develop necessary algorithms and business rules to filter and process raw log data (some IDS already do some of this and is referred to as log reduction). The objective is to simplify and automate the logic as much as possible so the operator does not have to constantly be reviewing raw data.
- Create effective communications capabilities between the automated, central program, and the staff. This may include automated e-mail or page notification, and even audible alarms when necessary. The capability should be planned around both normal operations and times when experts may be gone, such as nights, weekends, and holidays.
- Setup an ongoing improvement program so that analysts are increasing their effectiveness in defining algorithms for detection, and operators are trained to better understand the data.

6.4 Incident Response Tools

Incident response tool examples include:

6.4.1 Network Performance Monitoring

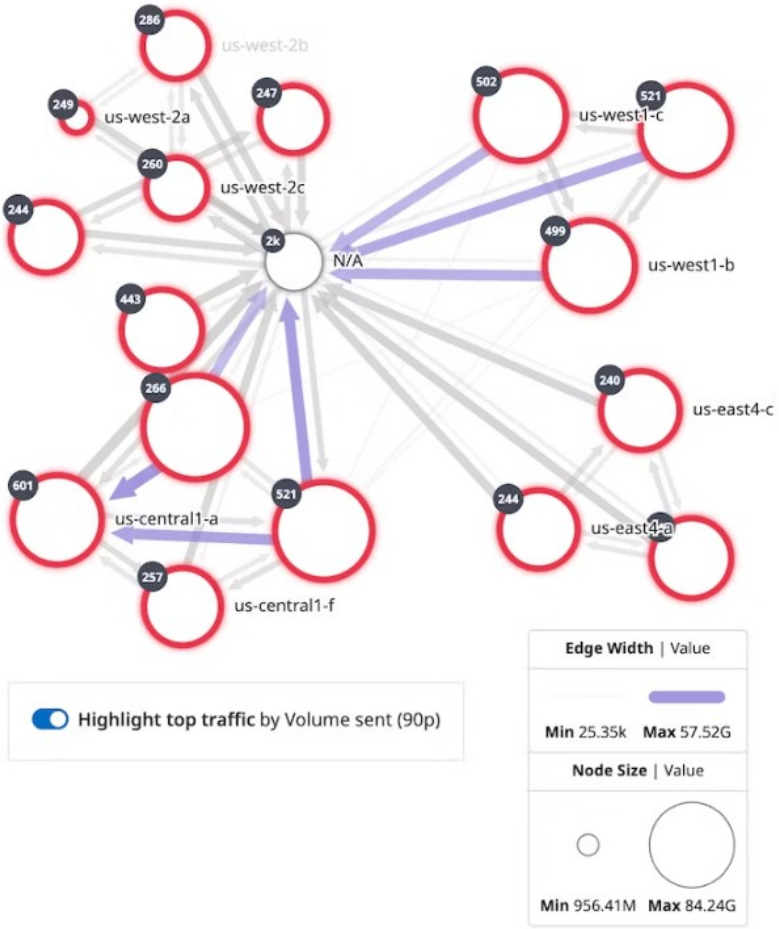

Figure 3: Datadog Network Performance Monitoring

Network Performance Monitors

They provide additional insight into network performance and can help identify where out-of-normal performance is occurring. They may also include bandwidth monitoring and analysis as well as network routing analysis.

Availability Monitors

These tools can assist in determining if network devices are available with advanced “ping” capabilities such as displays of real-time response rates.

Application Monitors

A specific application can be monitored if there is suspicion of unauthorised access or manipulation. These tools allow a more granular analysis of a suspected application as compared with overall network monitoring.

6.4.2 Network Traffic Analysis

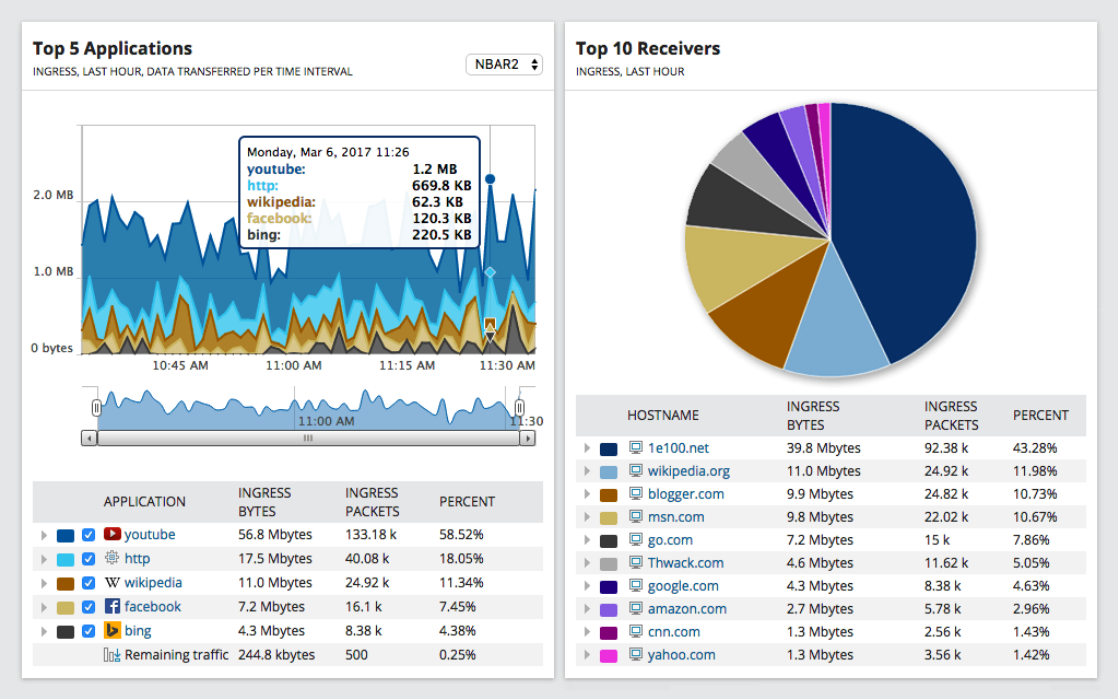

Figure 4: SolarWinds NetFlow Analyser

Netflow Capture and Analysis

These tools provide methods to capture and display the type of traffic crossing the network, including inbound and outbound traffic. These tools can isolate data by applications, conversations, domains, endpoints, and protocols. Many of these tools will also store data for both analysis and forensic work.

Packet and Traffic Reconstructors

Often associated with, or bundled as part of a network traffic monitor, these tools reconstruct files back into their original format on the network, capturing a static image of the network and the associated traffic.

6.4.3 Network Troubleshooting

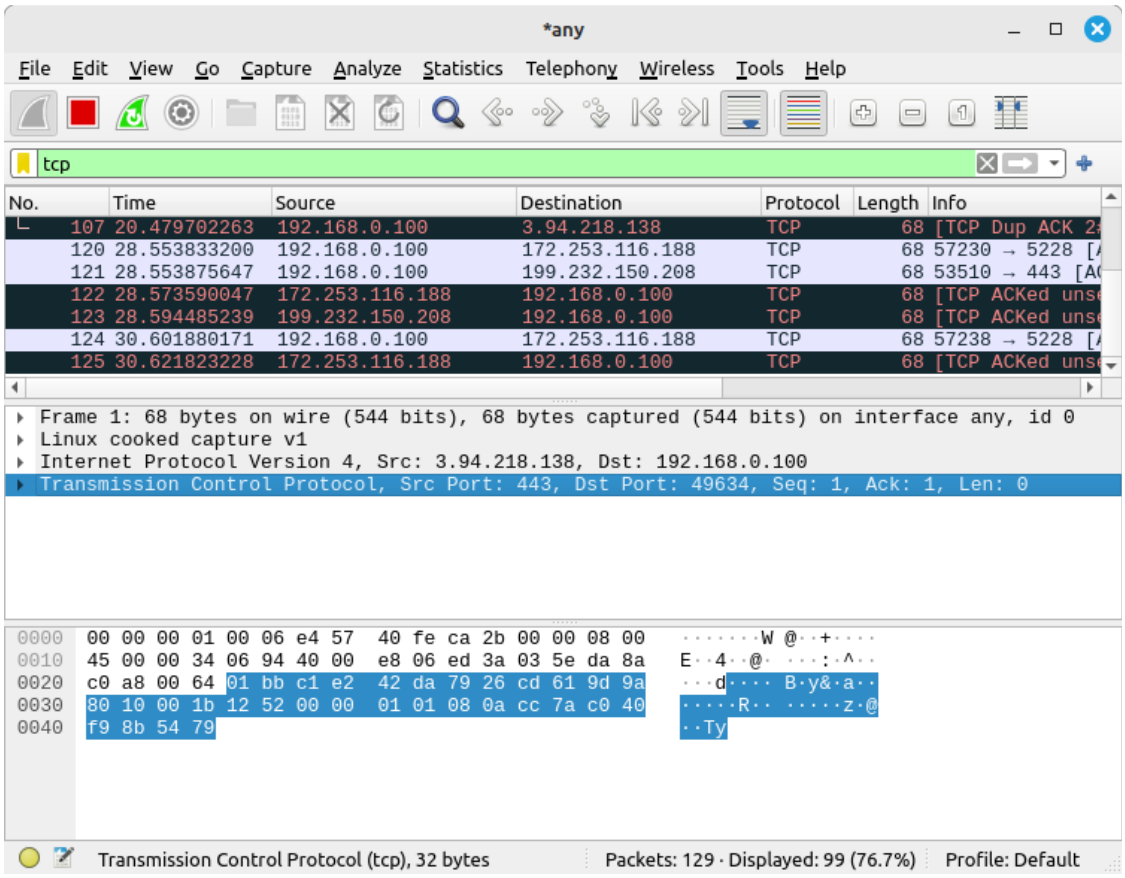

Figure 5: Wireshark

Protocol Analyser

Similar to other tools mentioned above, this tool/feature captures and stores for potential forensic analysis packet information, including consolidated statistical information.

Trace Route and Whois tools

These can be helpful in tracing an intruder to the location of the source computer. Associated functions allow IP address blocking and reporting.

```
~$ traceroute www.setu.ie
traceroute to www.setu.ie (172.67.41.36), 30 hops max, 60 byte
packets
 1  _gateway (192.168.0.1)  1.811 ms  2.758 ms  2.653 ms
 2  109.255.186.1 (109.255.186.1)  10.439 ms  17.298 ms  18.767 ms
 3  109.255.251.158 (109.255.251.158)  15.093 ms  14.993 ms  16.888
ms
 4  162.158.36.15 (162.158.36.15)  24.042 ms  23.957 ms  23.866 ms
 5  172.67.41.36 (172.67.41.36)  22.222 ms  23.690 ms  23.606 ms
```

```
~$ whois setu.ie
Domain Name: setu.ie
Registry Domain ID: 700080-IEDR
Registrar WHOIS Server: whois.weare.ie
Registrar URL: https://www.heanet.ie/services/hosting/domain-registration
Updated Date: 2023-03-04T14:58:07Z
Creation Date: 2011-01-18T00:00:00Z
Registry Expiry Date: 2024-01-18T14:51:28Z
Registrar: HEAnet
Registrar IANA ID: not applicable
Registrar Abuse Contact Email: noc@heanet.ie
Registrar Abuse Contact Phone: +353.16609040
Domain Status: ok https://icann.org/epp#ok
Registry Registrant ID: 543066-IEDR
Registrant Name: REDACTED FOR PRIVACY
Registry Admin ID: 541920-IEDR
Registry Tech ID: 546768-IEDR
Registry Billing ID: REDACTED FOR PRIVACY
Name Server: eleanor.ns.cloudflare.com
Name Server: kurt.ns.cloudflare.com
DNSSEC: unsigned
>>> Last update of WHOIS database: 2023-10-14T07:01:04Z <<<
```

6.4.4 Security Information and Event Management (SIEM)

Figure 6: Splunk Unified Security and Observability Platform

SIEM tools can fit into multiple of the categories just described. SIEM tools collect and analyse log data from a variety of sources, including networks, applications, and security devices. This data can be used to monitor network performance, detect threats, and troubleshoot problems.

Here are some specific examples of how SIEM tools can be used in each category:

Network Performance and Monitoring

SIEM tools can be used to monitor network performance and availability by collecting and analysing log data from routers, switches, and other network devices. This data can be used to identify performance bottlenecks and troubleshoot network problems.

Network Traffic Analysis

SIEM tools can be used to analyse network traffic by collecting and analysing log data from firewalls, intrusion detection systems, and other security devices. This data can be used to identify suspicious activity and detect threats.

Network Troubleshooting

SIEM tools can be used to troubleshoot network problems by collecting and analysing log data from a variety of sources, including networks, applications, and security devices. This data can be used to identify the root cause of network problems and develop solutions.

SIEM tools are versatile tools that can be used for a variety of network security and monitoring tasks. In addition to these categories, SIEM tools can also be used for other purposes, such as forensic analysis and compliance reporting.

For OT there are a number of vendors that offer SIEM products, examples include:

- Industrial Defender
- LogRhythm NextGen SIEM for OT
- Siemens Industrial Security Sitewatcher
- Waterfall Security Industrial Control System Security Suite
- Dragos Industrial Security Platform

7 Incident Categorisation

Once positively identified, a cyber attack should be categorised, and the response prioritised based on that categorisation. The categorisation should be based on the type of incident and the potential damage to the OT. The type of incident will drive the appropriate level of response. The IRP should outline in detail what the level of response (and level of effort) should be for each type of incident. As mentioned earlier, this planning should occur well in advance of an actual event.

The prioritisation of the response should be based on the current and potential effect to the OT, and the criticality of the effected equipment and system to company operations.

The following questions will aid in determining the categorisation/prioritisation criteria:

- How did the exploit occur and can it happen again? In what timeframe?
- Was this internal or external to the organisation?
- What type of attacker tools were placed onto the system, if any?
- What networks and systems are affected by the attack vector, and can the problem spread to other sites and customers?
- Are there legal or safety issues caused by the attack?
- How much does the impact increase if the incident is not contained within hours or days?
- Can systems safely fail-over or continue operating?
- How important are the effected components to the OT and to operations in general?

The following are recommended categorisation/prioritisation steps to take:

i. Assign a principal investigator responsible for identifying and mitigating each incident.
ii. Validate if the incident is a malicious or non-malicious occurrence. If the event is non-malicious, the full OT-CSIRT will not be required, though some resources may be used to solve the problem.
iii. Identify and evaluate the evidence in detail and keep accurate documentation with controlled access to the evidence.
iv. Coordinate with the specific personnel that provide operating business unit network services to the effected system.
v. Specific steps unique to the organisation should be included. They should be clearly defined in the IRP and should guide the actions of the OT-CSIRT when categorising and prioritising an incident.

8 Incident Containment

While containment often focuses on preventing the spread and effects of malware, several other types of incidents will require other actions related to containment. An example would be an employee who accesses unauthorised information by using another person's user account and password. Containing the situation would require removing the employee from access to the information and then enforcing disciplinary action as necessary. For an attacker who did not leave malware on the system, but was directly accessing OT components, containment would include blocking the intruder, restoring the equipment, if affected, and then applying protective steps.

The primary case for containment is where malware in some form has been left on the OT. This section will focus on containment issues related to software that has been placed on servers or other components that will either create an access path for an intruder, or will independently run to cause harm to the OT [3].

There are two main purposes in the containment of malware. The first purpose is to stop the spread to other parts of the system. The second purpose is to prevent continued damage to the OT. Even if the malware is isolated from spreading to other components or networks in the OT or across facilities, it may and can continue to cause damage in the isolated segment.

The containment of malware does not follow a standard approach for each organisation. It will vary based on the type of malware, the importance of the effected system, and the acceptable level of risk. Thus, every organisation must determine its proper containment actions based on its unique system requirements. The containment criteria need to be well documented and understood by members of the organisation and the OT-CSIRT.

Several methods to malware containment are available. The first method uses automated technologies such as virus removal programs to eliminate the problem and restore system functions. The second method halts services while the incident is being handled, and the third method blocks certain types of network connectivity by using a filtering process.

Using automated technologies provides immediate detection and response if the user chooses to program the application in this manner. This method can only act against known malware and cannot remediate Zero Day vulnerabilities. Zero Day exploits target vulnerabilities for which there is no available patch. These tools can significantly reduce cyber threats by acting as a filtering process or first defence, which can save organisation resources and reduce system downtime. One of the challenges to the control system engineer is finding automated applications that handle unique OT components, especially those that are using dated or unique protocols.

Temporarily halting services is a more drastic and potentially disruptive measure typically executed at the application level such as disabling a service. This could occur on a server or at the network level such as using firewalls to block IP addresses or ports associated with a service. Halting specific affected services stops and prevents the rapid spread of the infection while maintaining operation of the unaffected components to avoid complete loss of service. The desired goal is to contain the incident effectively with the least amount of loss in functionality. To

effectively prepare for halting services, an organisation must maintain a list of network and component services used along with their associated Transmission Control Protocol and User Datagram Protocol ports.

Using containment through disabling connectivity is an effective and quick means of temporarily restricting network connectivity to infected systems attempting to establish connection to an external system. This can prevent malware from downloading and prevent the spread of that system's infection to other internal networked systems. The intention is to isolate the critical control system from the network by removing the networking communication point and then to test and verify isolation without disrupting other critical services. This method of disconnecting critical OT network components should be identified and tested in the incident planning and preparation stage.

9 Incident Remediation

Prior to full system recovery, remediation efforts should be performed to fix the source of the problem. This may include eradication of any malware left on the system, removal or replacement of vulnerable equipment, reconfiguration and patching of equipment or software, and possible access cancellation for certain personnel.

If the incident involved unauthorised access then efforts should be made to close the access path. This may include changing all passwords and certain user names. Efforts may also include blocking access from identified IP addresses and changing of port configurations on firewalls.

Careful analysis should be performed on the OT to verify the path taken by the intruder. This should not only expose the actual weakness, but it can also highlight similar areas that may need attention. A specific dial-up device may have been the culprit, but other comparable devices may be scattered throughout the OT that are just as vulnerable.

If the incident involved malware left on the system, then removal or eradication will be necessary. Ideally, eradication will remove the malware with the least amount of disruption to the facility's operations. This process of removing malware could take some time to successfully accomplish, depending on the type of malware, severity of the infection, and containment method used.

Many techniques can remove malware from an infected system. The most common method is using automated eradication tools such as antivirus software, spyware detection and removal utilities, and patch management software. Other options include restoring a system to a set point before the infection or reloading key system files. These tools can quickly find and remove malware if they have detected the infection. Unfortunately, most antivirus type programs focus on typical IT systems and would not detect malware on more specialised control systems. There is also the danger that these utilities will remove or alter legitimate system or data files. In these situations, manual removal may be necessary with help from the vendor, or the vendors themselves may be able to provide removal software that has been tested against the target system.

For more severe cases of malware infection, a rebuild may be required. This technique would encompass reinstallation and securing of the operating system and application followed by restoring data from backup files.

A complete rebuild should be considered if the following system characteristics are present:

- The intruder gained root or administrator-level access to the system.
- Back-door type access has been granted that is not readily identified. The risk is that one back door may be found, but others may go undiscovered.
- System files were replaced by the malware or directly by the intruder.
- The system is unstable or does not function properly after antivirus software, spyware detection and removal utilities, or other programs or techniques eradicate the malware. This indicates that either the malware has not been eradicated completely or that it has caused damage to important system or application files or settings.

When the eradication efforts are finished, it is highly recommended that testing be conducted to verify that the OT is working as intended. This includes not just observable behaviour, but also reviewing any incident detection information to look for underlying signs of remaining rogue code.

10 Incident Recovery and Restoration

The OT environment introduces additional complexities related to recovery and restoration that would not be found in typical IT systems. However, some commonalities with traditional IT include removal of malware, restoring backup data to databases, systematically removing temporary containment actions, and restarting all operational systems and applications.

The additional complexities in the OT are related to the manner in which systems must be managed as part of the incident response. Because many of the services provided by the facility cannot be shutdown during the response, other approaches have to be taken. These include switching the control functions to fail-over systems, moving to backup equipment that is temporary or has limited capabilities, or isolating system components from network access. In these situations, the vital equipment and processes continue to operate, but in a temporary state with limited integration and, in some cases, reduced functionality.

Because of the demand for continuous operation, this temporary operational state is a higher risk for the enterprise. Having redundant systems in place is expected in most critical situations but, triple redundancy, while ideal, is not always possible due to high costs and architectural complexities. As a result, if the backup systems fail, production stops, which puts great pressure on the OT-CSIRT and operational staff to restore operations as soon as possible.

Information on restoring traditional IT components can be found in computer security documents, such as NIST SP 800-61 [4]. Specific recommendations for OT follow:

- Establish contingency plans with available equipment (even portable equipment if necessary) identified before the incident. This will allow operations to continue while primary systems are being restored.
- Patch and maintain all backup systems to the same level as the primary systems.
- Conduct regular and planned testing at a planned specific time to verify that the fail-over systems will work properly when called upon.
- Establish plans to run segments of the OT in isolation prior to an incident. This will provide the engineers a realistic picture of interdependencies between components, allowing them to make decisions on isolation, if necessary.
- Test backup equipment against realistic time-frames found in a worst-case scenario. For example, backup generators may need to power a system for days rather than hours, depending on the circumstances of the facility.
- Establish and run acceptance tests and procedures to ensure that systems have been restored to the pre-incident state. These may include both automated and manual tests.

- Define procedures as part of the IRP to provide for the proper authority to accept the tests and declare the OT fully operational.

The final stage of recovery is to not just restore the system to where it was, but rather to make it better and more secure. The system should have the same operational capabilities, but it also should protect against the exploit that caused the incident in the first place.

11 Post Incident Analysis and Forensics

Post-incident analysis and forensics consists of three subject areas. The first area is lessons learned where an attempt is made to analyse the incident, the response, and the impact to discover and document what could have been done differently to improve the response. The second area is recurrence prevention, or applying what was learned in remediating discovered weaknesses in the cybersecurity programme, including preventing a similar incident. The third area is forensics, which includes capturing and protecting data as evidence for potential legal action.

11.1 Lessons Learned

Unfortunately, cyber attacks are dynamic, with attackers learning quickly from their successes and capitalising on failed or incomplete defences. Every cyber event provides an opportunity to see clearly the weaknesses in the security posture of the control systems. It also reveals any weaknesses in the way the organisation handles its response. Performing a lessons learned exercise is essential to identifying weaknesses and preventing the reoccurrence of mistakes.

Any incident, whether successful or not, should be used as a chance to gain additional information to secure the OT. For example, a near hit, where an outside reconnaissance effort is detected yet not exploited, can provide valuable information. Much useful data can be discovered by extensive review of the logging functions of firewalls, routers, switches, servers, and workstations. This allows the analyst to determine a baseline of normal activity and how unauthorised access is attempted or successfully completed. An incident need not be limited to only a physical attack on a system. Other attempts to gain access include non-cyber related activities such as social engineering or email phishing attempts to get recipients to reveal data, passwords, or account configuration information.

A lessons learned exercise should be performed after every identified incident. Doing so will allow the incident to be reviewed so that access paths in system security can be identified and closed. If the problem is not found and fixed, the attack can be repeated, only with greater ease and frequency.

It is highly recommended that other incidents, beyond those of the facility, be considered for review. This is an ideal opportunity to continuously improve the security posture of the OT without having to suffer the damages of an actual incident. A lesson learned exercise should be held as soon after the incident as possible. This will typically follow the recovery and restoration phase. Any delay in conducting this exercise will leave the OT vulnerable to additional, similar exploits. Guidelines for conducting this exercise are as follows:

i. All members of the internal OT-CSIRT should participate if at all possible; the different perspectives and experience base will produce valuable perspectives.
ii. The OT-CSIRT Team Manager should assume responsibility to call and organise the lessons learned exercise. Notes should be taken of both the discussion and the action items.
iii. Information should be sought from external sources, including vendors, integrators and other national and subject-specific incident response teams such as the NCSC. This will provide additional details on the exploit and ways that others have mitigated the vulnerability.

Key questions should be answered, including:

- What components were affected—type, manufacturer, etc.?
  - What operating systems, including embedded ones, were affected?
  - How access was gained?
  - What damage was done and what potential damage could have been done?
- What network vulnerabilities, if any, allowed access to the OT?
- What standards and technical solutions might have prevented the incident?
  - What procedures and policies might have prevented the incident?
  - What training is necessary to prevent additional exploits?
  - How was the incident detected?
  - Could it have been found earlier or prevented?
  - Are we still vulnerable and for how long?
- Have vendors provided any patches or other solutions, and if so, were they implemented in the OT in a timely manner?
- What were the breakdowns in the incident response, including equipment, communications, lines of authority, vendor interactions, analysis, decision-making, and recovery?
- What areas need to be improved and have processes changed?
  - Can this information be shared with trusted partners?
- Can this information be shared with appropriate government agencies, including response teams?
- Based on the identification of weaknesses, specific assignments should be given to participants to systematically address each concern. The OT-CSIRT Team Manager should take responsibility to see that all actions are completed in a timely manner to prevent additional exploits.

11.2 Incident Recurrence Prevention

Once a vulnerability has been discovered, it will remain an open door until preventive action is taken. One of the primary purposes of the lessons learned exercise is to analyse the incident and initiate action to prevent a recurrence of the exploit. Considerations following on the incident:

11.2.1.1 Identify access methods

Identifying access methods may be simple or difficult, depending on the incident. An incident that was caused by an employee or contractor with inside access would be easy to identify but difficult to resolve because legitimate access must be provided. Preventive actions may include increased background checks, better training, and access control based on a stronger need to know and role responsibility in the organisation. Other incidents may involve malware that was loaded on a server. The solution might include removal of the malware followed by additional antivirus support and more detailed user training on social engineering and prevention techniques. A more difficult situation, where the access method is hard to discover, might include a skilled intruder that erased network logs, spoofed timestamps, or used compromised accounts that had necessary or administrative access to the OT. If the method cannot be discovered by the internal OT-CSIRT, it may be necessary to call in expert help to discover the method. If the access path cannot be found, then, as a last resort, the facility should conduct a systematic effort to strengthen all possible access paths.

11.2.1.2 Understand intruder motivation

Because it would be difficult to assign resources to every aspect of the OT, it is practical to increase security in specific targeted areas. For example, if the motive is to steal information, then databases would most likely be the target and securing key database servers and management systems would be the immediate priority. If public chaos, harm, and publicity are the desired outcomes, certain OT components may need attention. If the motive is financial damage to the company, production processes might be targeted. Understanding the motives allows more immediate and focused attention on specific aspects of the OT environment.

11.2.1.3 Assess and strengthen specific OT components

Access methods typically involve the network, but the incident might expose vulnerabilities related to specific models or types of components. This assessment can expose unpatched components, outdated equipment, or open communications between components. Solutions include equipment replacement, patching, and strengthening boundaries around components in the OT where components cannot be easily replaced. This analysis may provide cost justification for replacement of dated system components.

11.2.1.4 Review detection methods

When an incident occurs in a facility, the detection methods typically were not strong enough to identify the attempt in the early stages. For example, reconnaissance activities may have been going on for days or weeks before the actual exploit. Solutions might include stronger intrusion detection methods and software applications or a greater need for log reviews and analysis.

12 Incident Response Recommendations SP800-61r3

The current NIST publication in this area is Special Publication 800-61 revision 2 [4]; however, NIST now have a Special Publication 800-61 revision 3 in Draft status [5] incidents were relatively rare, and the the scope of most incidents was narrow and well-defined, and incident response and recovery was usually completed within a day or two. Under those conditions, it was realistic to treat incident response as a separate set of activities performed by a separate team of personnel and to depict all incident response activities as part of a circular life cycle. Formal post-incident activities would identify needed improvements and feed them into the preparation stage, thus starting the cycle again. Incident response activities were typically intermittent rather than continuous.

However, NIST now considers this model as no longer reflecting the current state of incident response. Today, incidents occur frequently and cause far more damage, and recovering from them often takes weeks or months due to their breadth, complexity, and dynamic nature. Incident response is now considered a critical part of cybersecurity risk management that should be integrated across organisational operations. The lessons learned during incident response should often be shared as soon as they are identified, not delayed until after recovery concludes. Continuous improvement is necessary for all facets of cybersecurity risk management in order to keep up with modern threats.

The Cybersecurity Framework 2.0 (CSF 2.0) [6] functions organise cybersecurity outcomes at their highest level:

- **Govern (GV)**: The organisation's cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored.
- **Identify (ID)**: The organisation's current cybersecurity risks are understood.
- **Protect (PR)**: Safeguards to manage the organisation's cybersecurity risks are used.
- **Detect (DE)**: Possible cybersecurity attacks and compromises are found and analysed.
- **Respond (RS)**: Actions regarding a detected cybersecurity incident are taken.
- **Recover (RC)**: Assets and operations affected by a cybersecurity incident are restored.

All six Functions have vital roles in incident response.

Govern (GV), Identify (ID), and Protect (PR) help organisations prevent some incidents, prepare to handle incidents that do occur, reduce the impact of those incidents, and improve incident response and cybersecurity risk management practices based on lessons learned.

Detect (DE), Respond (RS), and Recover (RC) help organisations discover, manage, prioritise, contain, eradicate, and recover from cybersecurity incidents, as well as perform incident reporting, notification, and other incident-related communications.

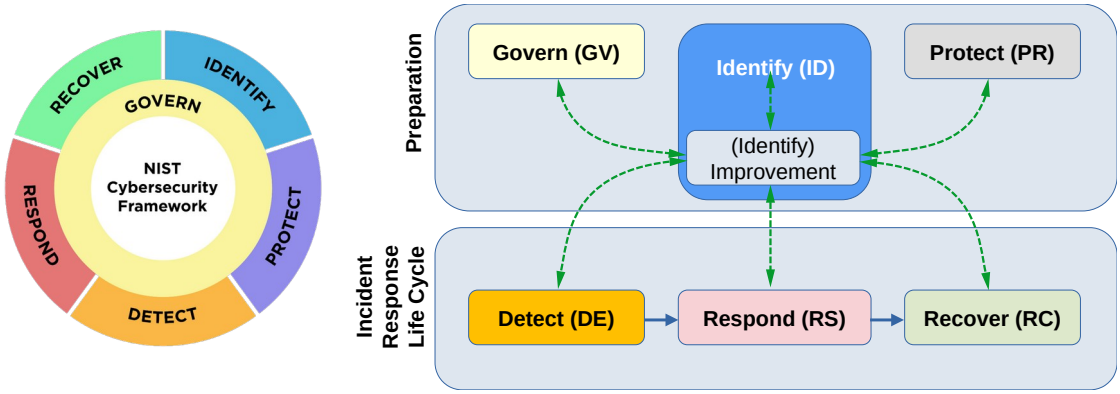

Figure 7: Incident Response Life Cycle based on CSF 2.0 Functions

Figure 7 illustrates the newly proposed incident response life cycle model based on the six CSF 2.0 Functions. The top half reflects that the preparation activities of Govern (GV), Identify (ID), and Protect (PR) are not part of the incident response life cycle. Rather, they are much broader cybersecurity risk management activities that also support incident response. The new response life cycle for each incident is shown in the bottom half of the figure: Detect (DE), Respond (RS), and Recover (RC). Additionally, the need for continuous improvement is indicated by the Improvement Category within the Identify (ID) Function and the dashed green lines.

Lessons learned from performing all activities in all Functions are fed into Improvement, and those lessons learned are analysed, prioritised, and used to inform all of the Functions. This reflects that organisations should be learning new lessons at all times (e.g., detecting the presence of a new threat and characterising its behaviour) and communicating those lessons to the appropriate personnel so that the organisation's incident response and other cybersecurity risk management policies, processes, and practices can be adjusted as needed.

13 Bibliography

[1] Eran Salfati and Michael Pease, ‘Digital Forensics and Incident Response (DFIR) Framework for Operational Technology (OT)’, National Institute of Standards and Technology, NISTIR 8428, Jun. 2022. Accessed: Aug. 22, 2023. [Online]. Available: https://doi.org/10.6028/NIST.IR.8428

[2] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, ‘Guide to Industrial Control Systems (ICS) security’, National Institute of Standards and Technology, SP 800-82 Revision 2, 2022. Accessed: Aug. 08, 2023. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r2.pdf

[3] M. Souppaya and K. Scarfone, ‘Guide to Malware Incident Prevention and Handling for Desktops and Laptops’, National Institute of Standards and Technology, NIST SP 800-83 Rev. 1, 2013. Accessed: Aug. 08, 2023. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-83r1

[4] P. Cichonski, T. Millar, T. Grance, and K. Scarfone, ‘NIST SP 800-61 Computer Security Incident Handling Guide’, National Institute of Standards and Technology, NIST SP 800-61 Rev. 2, Jan. 2020. Accessed: Aug. 08, 2023. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-61r2

[5] A. Nelson, S. Rekhi, Souppaya, Murugiah, and K. Scarfone, ‘NIST SP 800-61r3 Incident Response Recommendations and Considerations for Cybersecurity Risk Management’, National Institute of Standards and Technology, Initial Public Draft NIST SP 800-61 rev 3, Apr. 2024. Accessed: Oct. 31, 2024. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r3.ipd.pdf

[6] NIST, ‘Cybersecurity Framework 2.0 (CSF2.0)’, National Institute of Standards and Technology, Aug. 2023. Accessed: Aug. 22, 2023. [Online]. Available: https://doi.org/10.6028/NIST.CSWP.29.ipd