

Topic 8

Incident Management



Dr Diarmuid Ó Briain
Version: 3.0

Copyright © 2025 C²S Consulting

Licensed under the EUPL, Version 1.2 or – as soon they will be approved by the European Commission - subsequent versions of the EUPL (the "Licence");

You may not use this work except in compliance with the Licence.

You may obtain a copy of the Licence at:

https://joinup.ec.europa.eu/sites/default/files/custom-page/attachment/eupl_v1.2_en.pdf

Unless required by applicable law or agreed to in writing, software distributed under the Licence is distributed on an "AS IS" basis, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the Licence for the specific language governing permissions and limitations under the Licence.

Dr Diarmuid Ó Briain



Table of Contents

1 Objectives.....	5
2 Introduction.....	5
3 SANS Incident Response Framework.....	6
4 Preparation.....	7
4.1 Operational Technology Incident Response Team Resourcing.....	7
4.2 Team Roles and Responsibilities.....	9
4.3 Setting Policies and Procedures.....	13
4.4 Building the Cyber Incident Response Plan.....	14
4.5 Exercising the Plan.....	17
4.6 System State and Status Reporting.....	19
4.7 Incident Prevention.....	21
4.8 Vendor Interaction Considerations.....	22
5 Identification.....	23
5.1 Incident Detection.....	23
5.2 Detection by Observation.....	23
5.3 Automated Detection Methods.....	25
5.4 Incident Response Tools.....	27
5.5 Incident Categorisation.....	33
6 Containment.....	34
6.1 Methods of Malware Containment.....	34
7 Eradication.....	35
7.1 Remediation and Eradication.....	35
7.2 Methods of Malware Removal.....	35
7.3 When to Rebuild.....	35
8 Recovery.....	36
8.1 Recovery and Restoration in OT.....	36
8.2 Recommendations for OT Recovery.....	36
9 Lessons Learned.....	37
9.1 Conducting a Lessons Learned Exercise.....	37
9.2 Key Questions to Address.....	37
9.3 Preventing Incident Recurrence.....	38
10 ISO/IEC 27035 Information Security Incident Management.....	39
10.1 ISO/IEC 27035-1: Principles and Process.....	39
10.2 ISO/IEC 27035-2: Guidelines to Plan and Prepare.....	40
10.3 ISO/IEC 27035-3: Guidelines for ICT Security Operations.....	40
10.4 ISO/IEC 27035-4: Coordination.....	40

11 NIST SP 800-61r3 Incident Response Recommendations.....	41
11.1 Community Profile.....	41
11.2 Community Profile Example.....	43
12 Summary.....	44
13 Bibliography.....	45

Illustration Index

Figure 1: SANS Incident Management Stages.....	6
Figure 2: Incident Response Team.....	9
Figure 3: Datadog Network Performance Monitoring.....	27
Figure 4: SolarWinds NetFlow Analyser.....	28
Figure 5: Wireshark.....	29
Figure 6: Splunk Unified Security and Observability Platform.....	30
Figure 7: Splunk Data Pipeline.....	32
Figure 8: Incident Response Life Cycle based on CSF 2.0 Functions.....	42

Index of Tables

Table 1: SANS Incident Response Framework.....	6
Table 2: Summary: Incident Handling Frameworks.....	44

1 Objectives

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

- Explain the phases of the SANS Incident Response Framework (IRF).
- Identify and apply various methods for incident detection and containment.
- Formulate a comprehensive Incident Response Plan (IRP) that incorporates team roles, policies, procedures, and considerations for Operational Technology (OT) environments.
- Summarise the core principles of the ISO/IEC 27035 standards and NIST SP 800-61r3 and explain how they guide incident management practices.
- Describe the crucial steps for incident eradication and recovery.

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 SANS Incident Response Framework

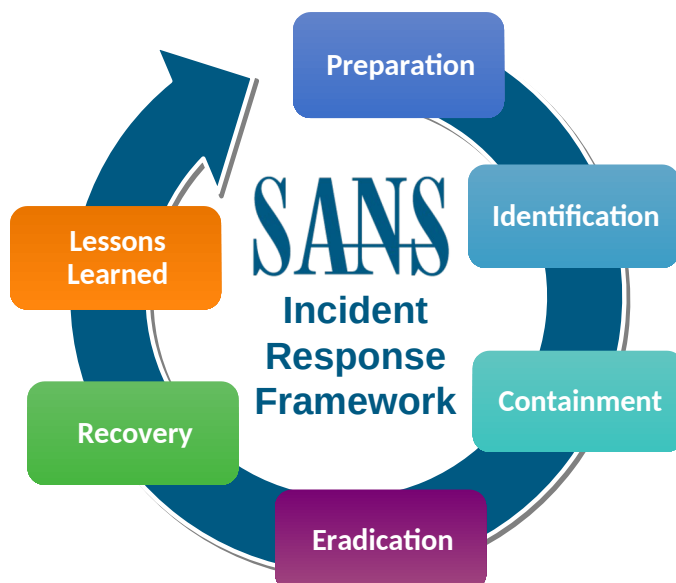


Figure 1: SANS Incident Management Stages

The SANS framework is known for being a very technical and practical approach to incident handling, providing hands-on guidance for security practitioners. The SANS Incident Response Framework (IRF) consists of the Preparation - Identification - Containment - Eradication - Recovery - Lessons Learned (PICERL) phases listed in Table 1:

Table 1: SANS Incident Response Framework

Phase	Description
Preparation	This phase involves getting ready for an incident before it happens by creating a plan, defining roles and responsibilities, and establishing a CSIRT.
Identification	The goal is to detect a security incident and determine its nature, scope, and severity by monitoring systems, analysing logs, and collecting evidence.
Containment	This phase focuses on stopping the incident from spreading and causing further damage by isolating affected systems or taking immediate actions to limit the impact.
Eradication	Once the threat is contained, you work to remove all traces of it from the environment, including removing malware, patching vulnerabilities, and addressing the root cause.
Recovery	After the threat has been eradicated, you restore affected systems to their normal, secure state, which may involve using backups, rebuilding systems, and monitoring.
Lessons Learned	This final phase involves a post-incident review to understand what happened, what went well, and what could be improved for future incidents.

4 Preparation

4.1 Operational Technology Incident Response Team Resourcing

The beginning point for creating a cyber-IR 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 IR 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.

4.1.1 Organising the Team

The first step in developing an IR 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.

4.1.2 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 IR related organisational policies and procedures
- 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
- 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.

4.1.3 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.

4.1.4 Centralised OT-CSIRT

A centralised cyber IR 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 IR 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.

4.1.5 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. IR, 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.

4.1.6 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.

4.2 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/manufacture. 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:

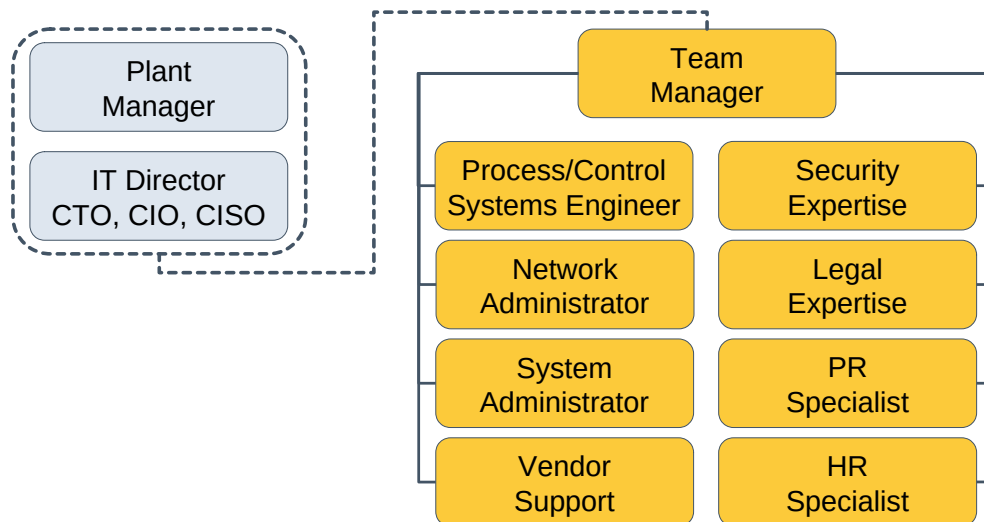


Figure 2: Incident Response Team

4.2.1 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.

4.2.2 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.

4.2.3 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.

4.2.4 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.

4.2.5 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.

4.2.6 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.

4.2.7 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.

4.2.8 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 IR issues.

4.2.9 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.

4.2.10 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.

4.2.11 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.

4.2.12 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 such as

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 Supervisory Control and Data Acquisition (SCADA) system outages, catastrophic equipment failure, or natural disasters such as floods or hurricanes.

4.3 Setting Policies and Procedures

While having policies and procedures are important in most business functions, IR 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 IR should be established and published within the OT organisation.

Clearly written, detailed operating procedures should be developed to implement the IR 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:

4.3.1 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.

4.3.2 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.

4.3.3 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.

4.3.4 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.

4.4 Building the Cyber Incident Response Plan

The cyber IRP establishes and documents the procedures and actions that implement the IR 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.

4.4.1 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.

4.4.2 Incident Description

Differentiating between IT and OT incidents is key. While many IT incidents are easily classified (e.g., DoS attacks, unauthorised access), defining what constitutes an OT security incident is more complex. It's crucial to determine if equipment failure or unexpected software behavior is a cybersecurity issue or a non-security related problem, such as mechanical wear.

Properly identifying the root cause is vital. If a failure is due to a cybersecurity vulnerability or malware, simply replacing the hardware or software won't solve the problem. The issue may persist or reinfect other systems. Accurate incident descriptions also prevent unnecessarily activating the OT-CSIRT.

For example, the response to a piece of equipment damaged by an employee with a crowbar would be vastly different from a similar incident caused by a remote attacker manipulating its controls. It's important to define each type of OT incident to ensure the appropriate response is followed.

4.4.3 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.

4.4.4 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.

4.4.5 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.

4.4.6 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.

4.4.7 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.

4.4.8 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].

4.5 Exercising the Plan

Although it may be inconvenient and disruptive to plan for, conduct and evaluate the results from an IR 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 IR 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.

4.6 System State and Status Reporting

Enabling system state and status reporting involves using automated mechanisms to report information about a system. This helps in detecting incidents, understanding their impact, and supporting quick resolution. Examples include network logging, database auditing, and custom-built or vendor-developed capabilities.

The primary purpose of adding reporting code to software is to aid in debugging or providing support information for troubleshooting. While it's valuable for forensics after an incident, its main use is for incident detection and resolution. The ability to detect equipment failure, improve work efficiencies, and resolve various system problems justifies the resources needed to enable it.

There are challenges in enabling status reporting in OT environments. Many devices use volatile memory, and their base code can be difficult to access. Vendors may also be reluctant to add new code due to cost or risk. Additionally, log data is often generated and overwritten too quickly to be stored practically. The extra network traffic from logging can even impair normal operations.

4.6.1 Networks Intrusion Detection Systems

Networks Intrusion Detection Systems (NIDS) 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.

4.6.2 Protocol-based Intrusion Detection System

A Protocol-based Intrusion Detection System (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.

4.6.3 Host-based Intrusion Detection System

An Host-based Intrusion Detection System (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.

4.6.4 Intrusion Prevention System

Because of their potential to cause OT failure, Intrusion Prevention Systems (IPS) aren't currently recommended for OT environments. They're similar to an IDS but actively block malicious activity. While not advised for OT systems, they are mentioned for their potential role in integrated business systems and for a comprehensive understanding of available technology.

If implemented, both NIDS and IPS would be primarily used on the network with limited application on server components. Due to the risks, extensive preliminary testing for OT compatibility is highly recommended before any deployment. An active system such as an IPS can inadvertently prevent legitimate activity, so it's critical to establish a baseline of approved activities before use.

4.6.5 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.

4.6.6 Configuration of Data Generators

Successful data gathering with commercial systems requires careful consideration of several key elements. You must understand all settings, properly configure devices, and regularly monitor alert notifications. An alert is useless if no one is assigned to receive or act on it, or if there are so many false positives that real incidents are overlooked.

For custom logging and monitoring, useful settings increase a device's value. For example, continuously measuring and reporting the state of field devices allows a server to test for out-of-range conditions or unusual traffic, which can then trigger alerts. The key is to analyse the specific devices and apply either vendor-provided or custom monitoring. When direct access to older or proprietary devices isn't possible, you can still monitor signals going to and from them to validate their state.

OT network traffic is limited and specific compared to IT systems, allowing baselines to be established and signatures created for abnormal activity. Be careful when enabling state and status reporting on OT systems, as some tools, such as poorly configured IDS and antivirus, can introduce operational issues. These tools can be intrusive, potentially disabling or shutting down legacy control systems or slowing critical data communications. Any deployment plan for these tools must be checked with the OT vendor and tested for compatibility.

When deploying these systems, you must consider:

- Where and for how long log files will be stored?
- Will older logs be deleted or archived?
- What parameters are being investigated? (e.g., ports, login times, abnormal traffic).

4.7 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.

4.7.1 Patch Management Considerations

Patch management is a key part of an effective cybersecurity programme, especially in OT environments. It is crucial for both preventing incidents and responding to vulnerabilities to prevent re-exploitation.

Key considerations for OT patch management include:

- **Scheduling:** Finding maintenance windows for production systems can be difficult.
- **Unsupported Equipment:** Patches may not be available for older equipment.
- **Third-Party Patches:** Relying on patches from sources other than the original vendor.
- **Testing:** It's essential to test patches in a non-production environment before deployment.
- **Backup & Recovery:** Creating backups and having a disaster recovery plan is crucial in case a patch causes issues.
- **Roll-back Procedures:** Developing procedures to revert a patch if it interferes with operations.
- **Compatibility:** Patches can cause issues with adjacent OT applications.
- **Vendor Interaction:** Receiving timely patches and understanding the vendor's testing processes.
- **Risk:** Assuming the risk that a patch won't negatively impact the production system.
- **Timing:** Knowing how long it takes to deploy or remove a patch.
- **Embedded Software:** Managing patches for software embedded in OT components.

4.8 Vendor Interaction Considerations

Effective cybersecurity in an OT environment requires a strong partnership with vendors due to the proprietary nature of OT software, the industry's immaturity in cybersecurity, and a limited customer base.

IT and OT vendor models are vastly different. In the IT world, a small number of vendors serve millions of users, leading to mature patching processes and a customer expectation of timely support. It's also common practice for vendors to withdraw support for older products, encouraging customers to upgrade. In contrast, OT vendors sell numerous products with long service lives, often to a small customer base. This pressure to support multiple versions for a limited number of customers means vendors may not provide timely patches or fixes.

To ensure prevention and response, targeted interaction between customers and vendors is crucial. A unified customer voice can pressure vendors to address security issues. It is essential to establish and maintain Service Level Agreements (SLA) to ensure ongoing patch support. Customers should also provide continuous feedback to vendors to help them understand priorities.

When an incident occurs, the relationship with the vendor's technical staff is critical. It may be necessary to include the vendor's technical personnel as an extension of your OT-CSIRT. For this to work, you must maintain up-to-date contact information, including expertise and availability, and they should be aware that they may be called upon in an emergency. It's too late to set up new contracts during an incident, so SLAs should be in place to define expectations and costs for assistance. Where possible, include a turnaround time for patches or fixes in the agreement.

5 Identification

5.1 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.

5.2 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.

5.3 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 such as 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 programme:** so that analysts are increasing their effectiveness in defining algorithms for detection, and operators are trained to better understand the data.

5.4 Incident Response Tools

IR tool examples include:

5.4.1 Network Performance Monitoring

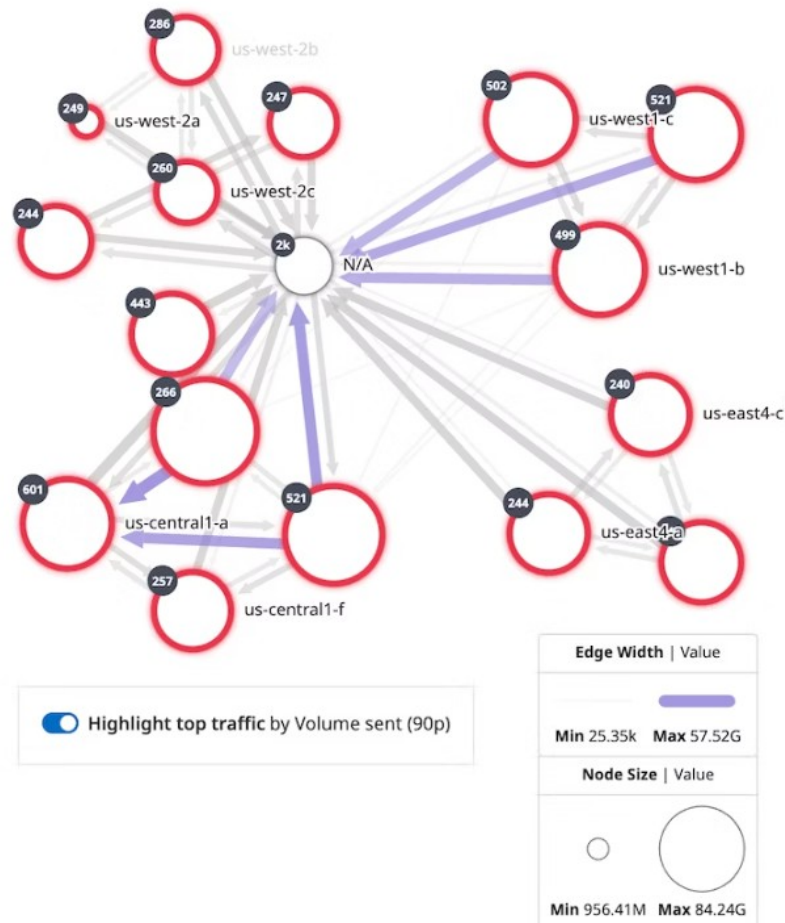


Figure 3: Datadog Network Performance Monitoring

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.

5.4.2 Network Traffic Analysis

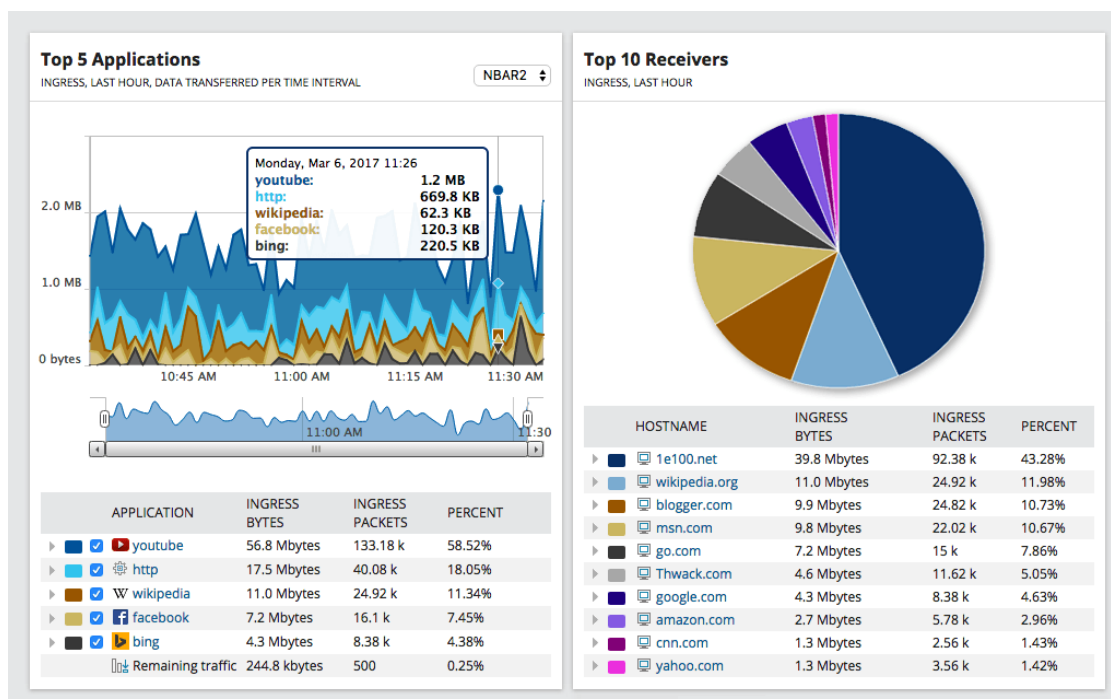


Figure 4: SolarWinds NetFlow Analyser

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.

5.4.3 Network Troubleshooting

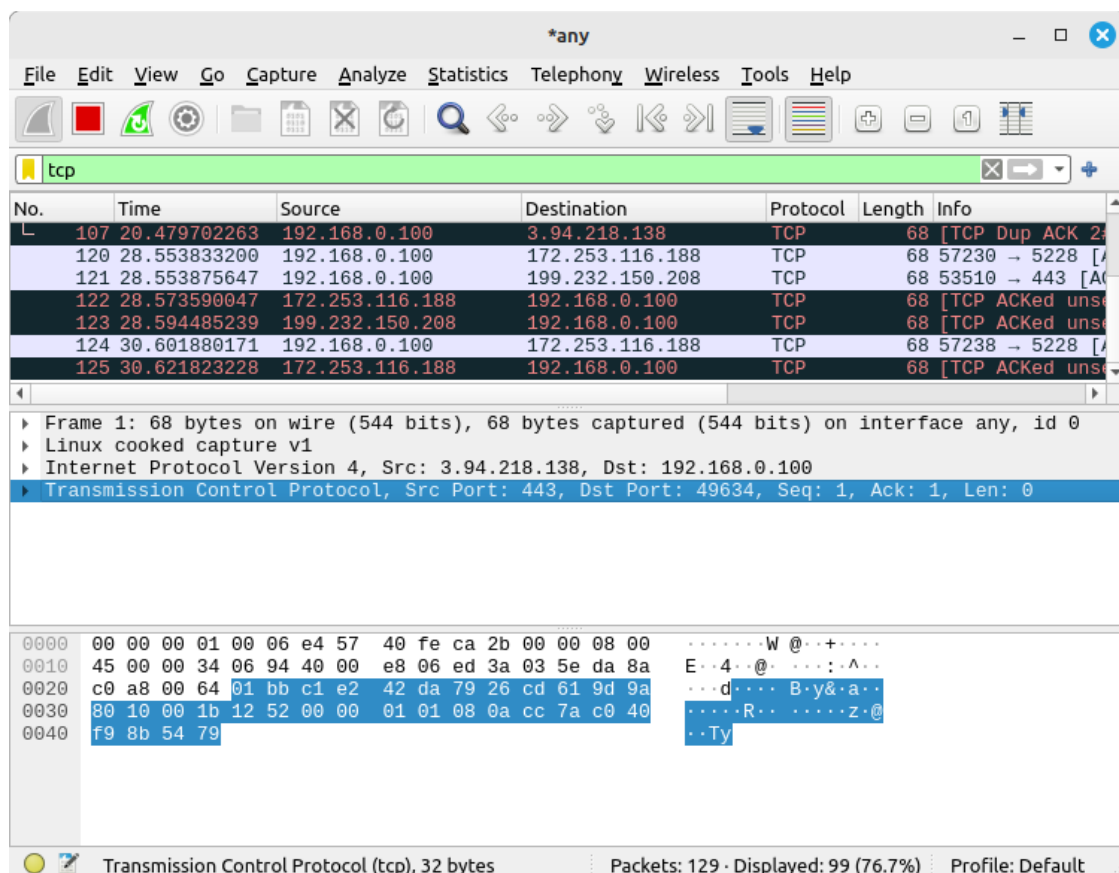


Figure 5: Wireshark

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 <<<

```

5.4.4 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

SIEM Example – Splunk

Splunk is a distributed SIEM system that enables an organisation to collect, index, and analyse machine-generated data at scale.

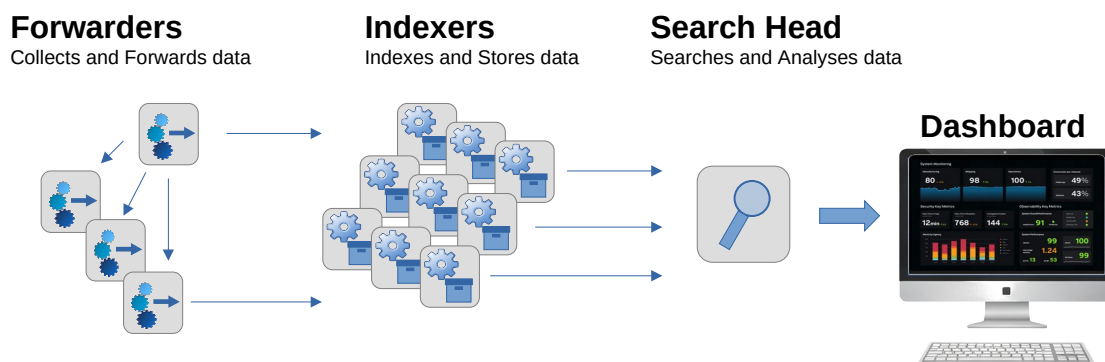


Figure 7: Splunk Data Pipeline

As illustrated in Figure 7, the core Splunk architecture is built on three main components that work together in the form of a data pipeline:

- **Forwarders:** These are lightweight agents installed on the systems where the data is generated such as servers, network devices, and applications. They collect the raw log data and securely forward it to the indexers. The most common type is the Universal Forwarder, which is designed for minimal resource consumption.
- **Indexers:** Indexers receive the data from the forwarders, parse it into individual events, and store it in an efficient, searchable format. They create and manage indexes, which are the data repositories, and handle the execution of search requests.
- **Search Heads:** This component provides the user interface for interacting with Splunk. Users connect to a Search Head to run queries using the Search Processing Language (SPL), create dashboards and reports, and visualise data. The Search Head sends search requests to the Indexers, aggregates the results, and presents them to the user.

The Splunk data pipeline works in a sequential flow with:

- **Data Ingestion:** Forwarders collect data from various sources and send it to the Indexers.
- **Indexing:** The Indexers process the incoming raw data, extract key fields such as timestamps and host information, and store it in an optimised format for fast retrieval.
- **Searching and Analysis:** A user, via a Search Head, runs a query using SPL. The Search Head distributes the query to the Indexers, which then retrieve the relevant data. The results are sent back to the Search Head, where they are aggregated and presented to the user as a searchable list, chart, or dashboard.

Separation of roles in this manner allows the system to be highly scalable and performant.

5.5 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.

6 Containment

Containment is not limited to malware; it also applies to other types of incidents, such as an employee using a stolen password to access unauthorised information. In this case, containment involves removing the employee's access and enforcing disciplinary action. For an attacker who directly accesses OT components without using malware, containment would involve blocking the intruder, restoring affected equipment, and applying protective measures.

This section focuses primarily on containing malware that has been left on OT systems to create an access path for intruders or cause independent harm. The two main goals of containment are to stop the malware from spreading and to prevent it from causing further damage to the OT. Even if malware is isolated, it can still cause harm within the contained segment.

Containment strategies are not one-size-fits-all; they depend on the type of malware, the importance of the affected system, and the organisation's risk tolerance. Every organisation must define and document its own containment criteria, which should be well-understood by the OT-CSIRT.

6.1 Methods of Malware Containment

There are three primary methods for containing malware:

1. **Automated Technologies:** This method uses tools such as antivirus software to detect and respond to known malware. While effective as a first line of defence, it cannot address zero-day vulnerabilities. A key challenge in OT is finding automated applications that work with unique or dated components.
2. **Halting Services:** This is a more drastic measure that involves temporarily disabling specific services at the application or network level. The goal is to stop the spread of infection while keeping unaffected components operational. To prepare for this, organisations must maintain a list of all network and component services.
3. **Disabling Connectivity:** This method involves restricting network connectivity to infected systems to prevent malware from downloading or spreading. The intention is to isolate critical systems by removing network communication and then verifying the isolation without disrupting other critical services. This method should be identified and tested during the incident planning phase.

7 Eradication

7.1 Remediation and Eradication

Before a full system recovery, remediation efforts must be performed to fix the source of the problem. This includes the eradication of any malware, removal or replacement of vulnerable equipment, patching, and possibly revoking access for certain personnel.

If the incident involved unauthorised access, efforts should be made to close the access path. This may require changing all passwords and certain usernames, blocking access from identified IP addresses, and reconfiguring firewalls. A careful analysis of the OT system should be performed to verify the intruder's path, not only to expose the specific weakness but also to highlight similar vulnerabilities that may need attention.

If malware was left on the system, eradication is necessary. The goal is to remove the malware with the least amount of disruption. This process can take time depending on the type of malware, the severity of the infection, and the containment method used.

7.2 Methods of Malware Removal

Several techniques can remove malware from an infected system:

- **Automated Tools:** The most common method uses automated eradication tools such as antivirus software, spyware detectors, and patch management software. These tools can be effective against known threats but may not detect malware on specialised control systems. There is also a risk that they could remove or alter legitimate system files. In these cases, manual removal may be necessary, often with the help of a vendor.
- **System Restoration:** Options also include restoring a system to a pre-infection state or reloading key system files.

7.3 When to Rebuild

A complete system rebuild may be required for more severe infections. This involves reinstalling and securing the operating system and applications, followed by restoring data from backups. A rebuild should be considered if:

- The intruder gained root or administrator-level access.
- Undetected backdoors were created.
- System files were replaced by malware or the intruder.
- The system remains unstable or malfunctions after eradication efforts, indicating that the malware was not completely removed or that it caused irreparable damage.

After eradication is complete, it is highly recommended to perform testing to verify that the OT system is working as intended. This includes reviewing incident detection logs for any signs of remaining rogue code.

8 Recovery

8.1 Recovery and Restoration in OT

While some recovery steps are common to both IT and OT, such as removing malware and restoring backups, the OT environment introduces additional complexities. Many OT services cannot be shut down during an incident, so alternative approaches must be taken. This includes switching to fail-over systems, using temporary backup equipment, or isolating system components from the network. In these situations, vital equipment continues to operate in a temporary state with reduced functionality and integration, which presents a higher risk.

Having redundant systems is expected, but triple redundancy is often not feasible due to high costs. If backup systems fail, production stops, creating significant pressure on the OT-CSIRT to restore operations quickly.

8.2 Recommendations for OT Recovery

Specific recommendations for OT recovery include:

- **Contingency Plans:** Establish contingency plans with identified backup equipment, including portable options, before an incident occurs.
- **Maintain Backup Systems:** Ensure all backup systems are patched and maintained at the same level as the primary systems.
- **Regular Testing:** Conduct regular, planned tests to verify that fail-over systems will work properly when needed.
- **Isolation Plans:** Establish plans to run segments of the OT in isolation to understand component interdependencies, which helps in making decisions about necessary isolation.
- **Realistic Time-frames:** Test backup equipment against realistic worst-case scenarios, such as needing to power a system for days, not just hours.
- **Acceptance Tests:** Establish and run acceptance tests to ensure systems have been restored to their pre-incident state.
- **Defined Authority:** Define procedures within the IRP for who has the authority to accept the tests and declare the OT fully operational.

The final stage of recovery is not just to restore the system, but to make it more secure. The goal is to have the same operational capabilities while also protecting against the specific exploit that caused the incident.

9 Lessons Learned

After an incident, it is essential to conduct a lessons learned exercise. This post-incident analysis is a crucial opportunity to improve your security posture and response capabilities. The exercise helps to identify and document weaknesses, leading to recurrence prevention and stronger security programmes. This process should happen as soon as possible after the recovery phase to prevent the OT from being vulnerable to similar exploits.

9.1 Conducting a Lessons Learned Exercise

Every incident, successful or not, offers valuable information. For example, a "near-hit" where an external reconnaissance attempt is detected can provide useful data. By extensively reviewing logs from firewalls, routers, switches, and servers, you can establish a baseline of normal activity and understand how unauthorised access was attempted.

To get the most out of the exercise:

- **Participation is Key:** All members of the internal OT-CSIRT should participate to provide diverse perspectives.
- **Structured Process:** The OT-CSIRT Team Manager should organise the exercise, take notes on the discussion, and document action items.
- **External Input:** Seek information from external sources, such as vendors, integrators, and national IR teams, to get additional details on the exploit and how others have mitigated it.

9.2 Key Questions to Address

The lessons learned exercise should answer key questions to inform future actions. These questions include:

- What components were affected?
- How was access gained, and what damage was done?
- What network vulnerabilities allowed access?
- What standards, solutions, or procedures could have prevented the incident?
- How was the incident detected, and could it have been found earlier?
- Have vendors provided patches or solutions, and were they implemented in a timely manner?
- What were the breakdowns in the IR process, including communication, lines of authority, and vendor interactions?
- What areas need improvement, and have processes been changed?

Based on the identified weaknesses, specific assignments should be given to participants to systematically address each concern. The OT-CSIRT Team Manager should ensure all actions are completed in a timely manner to prevent further exploits.

9.3 Preventing Incident Recurrence

Once a vulnerability is discovered, it remains a risk until preventive action is taken. The lessons learned exercise is a foundation for proactive recurrence prevention.

- **Identify Access Methods:** Understand how the intruder gained access. For an insider threat, solutions might include stronger background checks and better access control. If malware was involved, the solution might include additional antivirus support and user training. If the access method is unclear, it may be necessary to bring in outside experts or systematically strengthen all possible access paths.
- **Understand Intruder Motivation:** Knowing the attacker's motive, whether it's to steal information, cause physical damage, or financial loss, allows you to focus security resources on the most likely targets.
- **Assess and Strengthen Specific Components:** The incident may expose vulnerabilities in specific hardware or software. This analysis can justify replacing outdated equipment, patching components, or strengthening boundaries around systems that cannot be easily replaced.
- **Review Detection Methods:** An incident often reveals that detection methods were not strong enough to identify the threat in its early stages. Solutions might include stronger intrusion detection systems, new software applications, or more frequent log reviews.

10 ISO/IEC 27035 Information Security Incident Management

The ISO/IEC 27035 series is a collection of international standards that provides comprehensive guidance for an organisation's information security incident management programme. The series breaks down the topic into several parts, each focusing on a specific aspect of incident management.

10.1 ISO/IEC 27035-1: Principles and Process

ISO/IEC 27035-1 [3] is the foundational document of the series. It establishes the core concepts, principles, and a five-phase process for managing information security incidents and vulnerabilities. The phases are: Plan and Prepare, Detect and Report, Assess and Decide, Respond, and Learn Lessons. It provides a structured, generic approach that can be applied by any organisation and is the basis for the more detailed guidance in the other parts of the series. Here's a breakdown of its key components and purpose:

10.1.1 A Structured, Generic Approach

The standard provides a generic and structured process that can be applied by any organisation, regardless of its size, type, or industry. This makes it a universal reference for building an incident management programme. It recognises that while an organisation may have information security controls in place, incidents are inevitable, and a planned approach is essential.

10.1.2 A Five-Phase Incident Management Process

The standard outlines a clear five-phase process for incident management. These phases are:

1. **Plan and Prepare:** This involves establishing a formal policy, creating an incident response team, and putting a plan in place. It also includes training and technical support preparation.
2. **Detect and Report:** The focus here is on identifying and reporting information security events as they occur.
3. **Assess and Decide:** This phase involves evaluating a reported event to determine if it is indeed an information security incident that requires a full response.
4. **Respond:** Once an incident is confirmed, this phase covers the technical investigation, containment, eradication, and recovery actions.
5. **Learn Lessons:** This is a crucial, continuous improvement step. It involves analysing the incident to understand what happened, identifying weaknesses, and updating policies and controls to prevent a recurrence.

10.1.3 Focus on the Full Lifecycle

ISO/IEC 27035-1:2023 does not just focus on the moment of the attack. It covers the entire lifecycle, from the proactive planning and preparation phase to the reactive response and the subsequent post-incident review and improvement.

10.1.4 Integration with Other Standards

The standard is designed to be used in conjunction with other parts of the ISO/IEC 27000 series, particularly ISO/IEC 27001 [4] (which sets the requirements for an Information Security Management System) and ISO/IEC 27002 [5] (which provides guidance on security controls). This ensures that incident management is not a stand-alone process but an integrated part of an organisation's overall information security strategy.

10.1.5 Beyond Just Incidents

While its primary focus is on incidents, the standard also covers the management of information security vulnerabilities, which are often the root cause of incidents. This holistic view helps organisations not only respond to attacks but also proactively address the weaknesses that could lead to them.

10.2 ISO/IEC 27035-2: Guidelines to Plan and Prepare

ISO/IEC 27035-2 [6] provides more specific, actionable guidance on the *Plan and Prepare* and *Learn Lessons* phases introduced in Part 1. It details how to set up the necessary organisational structure, technical support, and procedures for an IR team. It also provides guidelines for how to conduct post-incident reviews to learn from incidents and continuously improve the programme.

10.3 ISO/IEC 27035-3: Guidelines for ICT Security Operations

ISO/IEC 27035-3 [7] focuses on the operational aspects of IR within Information and Communication Technology (ICT) security operations. It provides guidance on how to detect, analyse, contain, eradicate, and recover from incidents within a Security Operations Centre (SOC) or a similar environment. This part is more focused on the hands-on, technical activities involved in IR.

10.4 ISO/IEC 27035-4: Coordination

ISO/IEC 27035-4 [8] provides guidelines for how multiple organisations can handle information security incidents in a coordinated manner. It addresses the challenges and best practices for inter-organisational communication and collaboration during a major incident, which is increasingly important in today's interconnected world.

11 NIST SP 800-61r3 Incident Response Recommendations

The current NIST publication in this area is Special Publication 800-61 Revision 3 [9] which was developed in an environment where incidents occur frequently and cause significant damage, and recovering from incidents often takes weeks or months due to their breadth, complexity, and dynamic nature. IR is now considered a critical part of Cybersecurity Risk Management (CRM) that should be integrated across organisational operations. The lessons learned during IR should often be shared as soon as they are identified, not delayed until after recovery concludes. Continuous improvement is necessary for all facets of CRM in order to keep up with modern threats. NIST SP 800-61r3 has moved away from a separate set of IR functions as was recommended in earlier revisions and it now dovetails with the NIST Cybersecurity Framework 2.0 (NIST CSF 2.0) [10]

11.1 Community Profile

NIST SP 800-61r3 is NIST CSF 2.0 Community Profile because it provides specific, detailed guidance for a particular community. In this case, organisations that need to manage cybersecurity IR, and organisations that guidance around the common structure of the NIST CSF 2.0.

The Community, in NIST CSF 2.0, is any organisation that needs to establish, maintain, or improve its cybersecurity IR capabilities. While the NIST CSF 2.0 is a general framework applicable to all organisations. NIST SP 800-61r3 focuses specifically on the needs and practices of this group.

11.1.1 Common Taxonomy

NIST SP 800-61r3 aligns its recommendations with the six functions of the CSF 2.0: Govern (GV), Identify (ID), Protect (PR), Detect (DE), Respond (RS) and Recover (RC). Instead of presenting IR as a separate, isolated task, the publication integrates it into the broader CRM activities defined by the CSF. For example, it explains how the GV function is critical for establishing IR policies and how the RC function fits into the overall restoration of business operations.

11.1.2 Specific Guidance

A Community Profile is designed to be a more detailed, actionable version of the high-level CSF. NIST SP 800-61r3 does this by providing concrete recommendations, considerations, and examples for IR that align with each CSF function, helping organisations apply the framework's principles in a practical way.

By framing IR guidance in this manner, SP 800-61r3 moves IR from being a reactive, technical process to a strategic, organisation-wide function that is an integral part of CRM.

The NIST CSF 2.0 defined functions to organise cybersecurity outcomes at their highest level as follows:

- **Govern (GV):** The organisation's CRM 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 IR.

- GV, ID, and PR help organisations prevent some incidents, prepare to handle incidents that do occur, reduce the impact of those incidents, and improve IR and CRM practices based on lessons learned.
- DE, RS, and 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 8 illustrates the IR life cycle model based on the six NIST CSF 2.0 Functions. The top half reflects that the preparation activities of GV, ID, and PR are not part of the IR life cycle. Rather, they are much broader CRM activities that also support IR. The new response life cycle for each incident is shown in the bottom half of the figure: DE, RS, RC. Additionally, the need for continuous improvement is indicated by the Improvement Category within the ID Function and the dashed green lines.

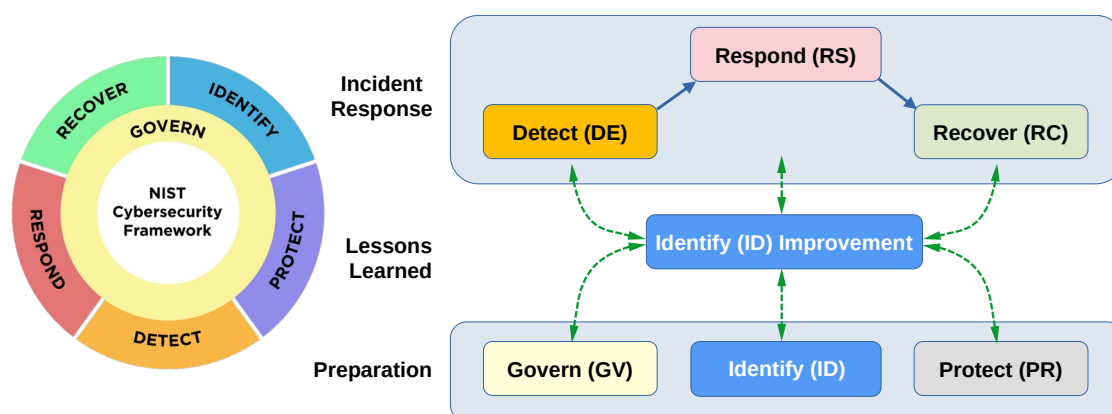


Figure 8: Incident Response Life Cycle based on CSF 2.0 Functions

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 and communicating those lessons to the appropriate personnel so that the organisation's IR and other cybersecurity risk management policies, processes, and practices can be adjusted as needed.

11.2 Community Profile Example

Imagine a large hospital network with dozens of facilities. The hospital's executive leadership wants to improve its cybersecurity posture using the NIST CSF 2.0. However, the CSF is a high-level framework. While it provides the six core functions, it doesn't give specific, detailed instructions on how to handle a data breach or a ransomware attack, which are major concerns in the healthcare community.

The hospital's CISO and IT team need practical guidance for building an IR programme that aligns with the CSF's principles. Simply knowing they need to Respond is not enough; they need to know what that actually looks like in a hospital environment.

11.2.1 Solution: Using NIST SP 800-61r3 as a Community Profile

The hospital team can employ NIST SP 800-61r3 as a Community Profile. It provides specific, actionable guidance they need, all within the structure of the NIST CSF 2.0.

- **Governing IR:** The CISO can use the guidance in NIST SP 800-61r3 to establish the hospital's IRP, define roles and responsibilities for legal, HR, and clinical staff, and ensure IR is part of the broader risk management strategy. This directly supports the NIST CSF 2.0's GV function.
- **Preparing to Respond:** The IT team can follow the guidance to create IR playbooks for common threats such as ransomware. They can use the document to identify critical medical devices and patient data (ID) and then implement multi-factor authentication and network segmentation to protect them (PR). This moves them from high-level CSF functions to concrete actions.
- **Detailed Response Steps:** When an incident occurs, the IT team does not have to start from scratch. NIST SP 800-61r3 provides a clear lifecycle for DE, RS, and RC from the incident. For example, it provides detailed steps on how to contain a virus on a medical device, eradicate the threat, and then restore the system from a clean backup, a specific scenario in the medical community.
- **Continuous Improvement:** After the incident is resolved, the hospital can use the Lessons Learned guidance from NIST SP 800-61r3. This helps them understand how the attacker got in and how they can improve their security to prevent a similar event in the future. This feedback loop is a key part of the CSF's continuous improvement model.

In essence, NIST SP 800-61r3 translates the generic, six-function language of the NIST CSF 2.0 into a detailed, practical roadmap for a specific group of people, those responsible for cybersecurity IR.

12 Summary

The three frameworks discussed in this topic all provide a structured approach to IR; however, they differ in their scope, target audience, and level of detail.

Table 2: Summary: Incident Handling Frameworks

Feature	SANS IR Framework	NIST SP 800-61r3	ISO/IEC 27035 Incident Management
Primary Audience	Incident handlers and technical responders.	Organisations establishing or improving an IR capability; security managers and planners.	Senior management, policy-makers, and those responsible for governance and risk management.
Focus	A tactical, hands-on guide for responding to incidents. It is highly practical and focused on the actions taken during an incident.	A strategic and programmatic guide for creating a comprehensive IR programme. It provides recommendations for the entire lifecycle, including policy and programme development.	A high-level, management-oriented standard for the entire incident management process, including prevention and continuous improvement. It is less technical and more about policy and governance.
Phases / Lifecycle	6 steps: Preparation, Identification, Containment, Eradication, Recovery, and Lessons Learned.	4 phases: Preparation, Detection and Analysis, Containment, Eradication, and Recovery, and Post-Incident Activity.	5 phases: Plan and Prepare, Detect and Report, Assess and Decide, Respond, and Learn Lessons.
Key Characteristics	<ul style="list-style-type: none"> • Actionable: Provides a clear, sequential process for responders. • Technical: Offers detailed guidance on tasks like evidence collection and system hardening. • Simple: The six-step model is straightforward and easy to remember. 	<ul style="list-style-type: none"> • Comprehensive: Covers all aspects of an IR programme, from policy to technical implementation. • Flexible: The guidance is general enough to be adapted by organisations of any size or sector. • US Government-backed: Widely adopted in the US public sector and considered a best practice globally. 	<ul style="list-style-type: none"> • Policy-Driven: Emphasises top-down management commitment and policy development. • International: A globally recognised standard, suitable for multinational corporations and for meeting compliance requirements. • Integration: Designed to be integrated with other ISO 27000 series standards.
Emphasis	Strong emphasis on the technical steps of containment, eradication, and recovery to minimise damage and restore normal operations quickly.	Strong emphasis on pre-incident preparation and post-incident analysis to continuously improve the programme. It also stresses the importance of communication and coordination.	Strong emphasis on the entire management lifecycle, including proactive measures like vulnerability handling and continuous improvement through <i>lessons learned</i> to prevent future incidents.

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, June 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, 'NIST SP 800-82r2 Guide to Industrial Control Systems (ICS) security', National Institute of Standards and Technology, Obsolete Standard SP 800-82 Revision 2, 2022. Accessed: Aug. 08, 2023. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r2.pdf>
- [3] ISO/IEC 27035-1:2023. *Information technology — Information security incident management — Part 1: Principles and process*, Geneva., Feb. 2023. Accessed: Sept. 10, 2023. [Online]. Available: <https://www.iso.org/standard/78973.html>
- [4] ISO/IEC 27001: 2022 *Information security, cybersecurity and privacy protection — Information security management systems — Requirements*, Jan. 10, 2022. Accessed: Sept. 10, 2023. [Online]. Available: <https://www.iso.org/standard/27001>
- [5] ISO/IEC 27002: 2022 *Information security, cybersecurity and privacy protection — Information security controls*, Feb. 2022. Accessed: Sept. 10, 2023. [Online]. Available: <https://www.iso.org/standard/75652.html>
- [6] ISO/IEC 27035-2:2023. *Information technology — Information security incident management — Part 2: Guidelines to plan and prepare for incident response*, Geneva., Feb. 2023. Accessed: Sept. 10, 2023. [Online]. Available: <https://www.iso.org/standard/78974.html>
- [7] ISO/IEC 27035-3:2023. *Information technology — Information security incident management — Part 3: Guidelines for ICT incident response operations*, Geneva., Sept. 2020. Accessed: Sept. 10, 2023. [Online]. Available: <https://www.iso.org/standard/74033.html>
- [8] ISO/IEC 27035-4:2023. *Information technology — Information security incident management — Part 4: Coordination*, Geneva., Dec. 2024. Accessed: Jan. 01, 2025. [Online]. Available: <https://www.iso.org/standard/80973.html>
- [9] A. Nelson, S. Rekhi, Souppaya, Murugiah, and K. Scarfone, *NIST SP 800-61r3 Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile*, Standard NIST SP 800-61 rev 3, Apr. 2025. Accessed: May 29, 2025. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r3.pdf>
- [10] NIST, 'NIST CSWP 29 Cybersecurity Framework 2.0 (CSF2.0)', National Institute of Standards and Technology, NIST CSWP 29, Aug. 2023. Accessed: Aug. 22, 2023. [Online]. Available: <https://doi.org/10.6028/NIST.CSWP.29.ipd>

This page is intentionally blank