INDUSTRIAL
CONTROL
SYSTEMS
Community of Interest

Where to store logs from Industrial Control Systems/Operational Technology environments

Introduction

This article provides guidance and considerations for Critical National Infrastructure (CNI) operators when considering where to store the logs captured within their Industrial Control System (ICS)/Operational Technology (OT) environment. It is aimed at architects, Security Operation Centre teams, third-party integrators, service providers, and vendors. It is the fifth of a series of articles to help operators undertake better logging and monitoring within their ICS/OT environments that the ICS COI has published. This guidance should be read in conjunction with related guidance the NCSC has published on logging and monitoring.

***Important Note:** While the ICS COI is supported by the NCSC, and NCSC staff are involved in a range of its activities, no formal review of this guidance article has been undertaken by the NCSC. The ICS COI and its members strive to produce relevant ICS/OT specific cyber security guidance to supplement principle based cyber security guidance published by NCSC and have taken care to reference this guidance where applicable. This guidance article will be reviewed every 18 months to ensure that it has not been superseded by guidance published by NCSC, relevance and that any references are still accurate. The ICS COI and its activities are purely voluntary, with guidance articles produced that are deemed needed by UK Operators and their supportive industry partners. The fact that this guidance article has been published by the ICS COI has no relevance to the priority and focus of guidance published by NCSC.*

Architecture

Regardless of the type of systems within the ICS/OT environment or the type of logs obtained (network, host, application, controller, and others) from the systems, there is a need to store the logs securely. This article will focus on the general considerations for a typical ICS/OT environment to guide ICS/OT operators to a position of collating the logs securely. Other guidance articles will address the skills required to perform ICS/OT log analysis and who should perform the analysis. Therefore, this article primarily focuses on the architectural principles for ICS/OT log storage.

What drives the decision?

Deciding where to store log files is equally as important as what to store, who will have access to them and how long they should be stored for.

There is no defined template or operational architecture which outlines where is the best place to store log files, the most crucial decision is to agree to collect these in the first instance. The responsibility of collecting and storing logs rests with the Cyber Security team, although the OT operations team are always involved and consulted in establishing the process.

It may be effective for a small site to store data locally on its network (so long as it is well connected), equally, it may make more sense to send these direct to a cloud storage solution for other reasons such as time, efficiency, access in real time, lack of local resources etc.

Log files should always be stored securely, whatever location decision is chosen.

Another consideration for log storage will be the expected volume of log data. For larger systems or sites, the volume of data collected can scale rapidly, so it is crucial that the advice from the other guidance papers (Why and What to Log) is taken into account to ensure a manageable amount of logs are collected and stored. The key point is that only relevant data is captured depending on the identified risk mitigations, Security Operation Centre (SOC) use cases, or data forensics requirements, that logging/monitoring can be useful for.

What may influence the storage location is the requirement on how the data is accessed and what happens to data if the site is placed into island mode. While there are many more factors to consider, these are just a few to note.

Centralised location

This approach is common for environments which started life as a disconnected system and have expanded data collection into the enterprise. Industries which require assurance of data sovereignty may often take an approach to retain data within their own boundaries. Principles to consider include:

- The most important architectural principle for ICS/OT log storage is to collect and store the logs centrally within the environment.
- Data should be stored in a secure manner with appropriate access controls in place, in a separate zone, accessible through a Demilitarised Zone (DMZ) with appropriate tracking and auditing. Consideration should be given to make the data stored read only to minimise the risk of logs being tampered with. Appropriate Identity and Access Management, and Role Based Access Control processes should be implemented.
- Established processes to be in place to control access and ensure nonrepudiation, processes implemented for onboarding and offboarding access.
- In larger environments logs may be forwarded to a Security Information and Event Management **(**SIEM) system or log aggregator (this should also be considered if environments are being expanded).
- In smaller environments logs may be forwarded to a secured engineering workstation.
- Policy, process, and technology to provide local log store and forward and possibly processing if the connection or host is lost.

Local collection

This approach is common for environments which operate disconnected systems and have not or cannot integrate data collection easily or safely into the enterprise, or to a dedicated local centralised log collection point. This example may also be common with disparate systems or those which operate autonomously without wider data connections. This approach is more subject to human error than other collection methods and requires administrative controls to be established to support this operation. Principles to consider include:

- The most important principle is to have an established process for log collection using safe and clean removable media.
- Process for collection may be resource intensive and complex; this may require engaging multiple teams for collection and processing.
- Data collected manually will be snapshot in time and therefore stale when reviewed in some cases.
- In larger environments it should be considered whether an approach to collect logs centrally should be adopted for critical assets (hybrid).
- In smaller environments manual collection and review is likely to be suitable approach, however this is true so long as it is practical to do so, supported with a well-documented process and consistently applied.

Cloud location

Cloud storage is common for established estates where there is a satisfactory level of maturity and well-established security processes to control traffic traversing security boundaries. Likewise, this collection method may also suit smaller sites where there is limited ability to self-host, such as a lack of skills locally, or remote locations with lack of resources. This approach scales well to include data from disparate sources of interest including endpoint, edge devices, controllers, and other operational assets where logging is possible or relevant.

It may be appropriate to have a local collection point that then feeds data to a cloud tenant to enforce control measures, as well as providing some level of reporting should there be a need to operate in island mode. This may be a site SIEM or log collector that has local reporting capability which offers a store and forward capability. Cloud storage may also be the operational approach of third parties who offer logging and event management as a managed service. Principles to consider include:

- The most important principle is to have child and parent collectors to ensure that devices are not communicating directly with a cloud collector, strategically placed in a DMZ (ideally there is also one-way flow control implemented out of the ICS/OT environment).
- Data should be proportionate and appropriate as far as sensitive data is not inadvertently sent off site.
- Access to the data may be to a wider audience than with other collection methods, therefore this should be routinely reviewed, and established access controls implement.
- Data retention may also be a consideration with cloud collection due to the elasticity of storage, also considerations for Service Level Agreements (SLA) with a Software-as-a-Service (SAAS) model, as well as data sovereignty.
- Design and capacity for local store and forwarding in the event of loss of connection, this may be undertaken at the local site collector or on device.
- Managed services through a 3rd party may collect data locally through an established Virtual Private Network (VPN) or through a local site collector, this may be for Network Operations Centre (NOC) as well as SOC.
- Considerations on the use of public and private cloud for the storage of data, design decision on this may differ if this is a pure SOC or SAAS which is an addition to an Original Equipment Manufacturer (OEM) Supervisory Control and Data Acquisition (SCADA) which may have cloud reporting capability for asset management.
- There may be a compelling reason to split Information Technology (IT)/OT logging into separate cloud tenancies to ensure alerts go to the correct resource and are not lost in enterprise noise, there may also be a single reporting portal that aggregates this data.

Hybrid location

Hybrid collection is an accumulation of local storage and cloud, while there may be elements of cross over between the cloud and local storage, as well as manual collection in some instances.

This approach has the most flexibility to deliver local access as well as providing greater scales of access of data to a wider audience. Data which is cloud hosted may not always be accessible during wider provider outages which needs to be considered, likewise, benefits which may include the ability to access log files during a cyber incident where local log access is simply not possible. Principles to consider include:

- The important principle with a hybrid model is ensuring that the correct resources are defined for storing or processing the data.
- Opportunities to create data pools which can be mined or processed with high levels of computer resources, including machine learning models.
- Considerations for where data is stored and processed, including data sovereignty.
- Offers the widest flexibility of all models that can be scaled as required to capture all collections models.
- Data may be duplicated or only stream a smaller subset thereof; a comprehensive approach needs to be taken into account where consumers only have access to a single source of data.
- Ability to locally buffer data when external connection is limited.
- Opportunities to scale solutions, especially where data is stored e.g. 1-week local storage and 1 years on cloud services.

Zone and conduit model

Zones and conduits are essential for securing ICS/OT environments by segmenting networks and controlling communication between different areas. Data flows may be inline or out of band of control traffic, each methods have risks and rewards. Traffic that is inline may be simpler; however, this should be reviewed for additional load and appropriate isolation.

Out of band collection may be appropriate where vendors do not warrant modifications to equipment or service provision. Principles to consider include:

- Using a zones and conduits model, widely referenced in standards such as IEC 62443, is considered good practice for log collection and storage
- This model allows logs to be obtained and forwarded from operational systems up and out to a dedicated zone where a log store can reside.
- This also allows the logs to be forwarded further or pushed up from an ICS/OT DMZ to an external location for further analysis and/or archiving.

Caution needs to be taken to ensure data collection does not inadvertently create unintentional communication paths for lateral movement.

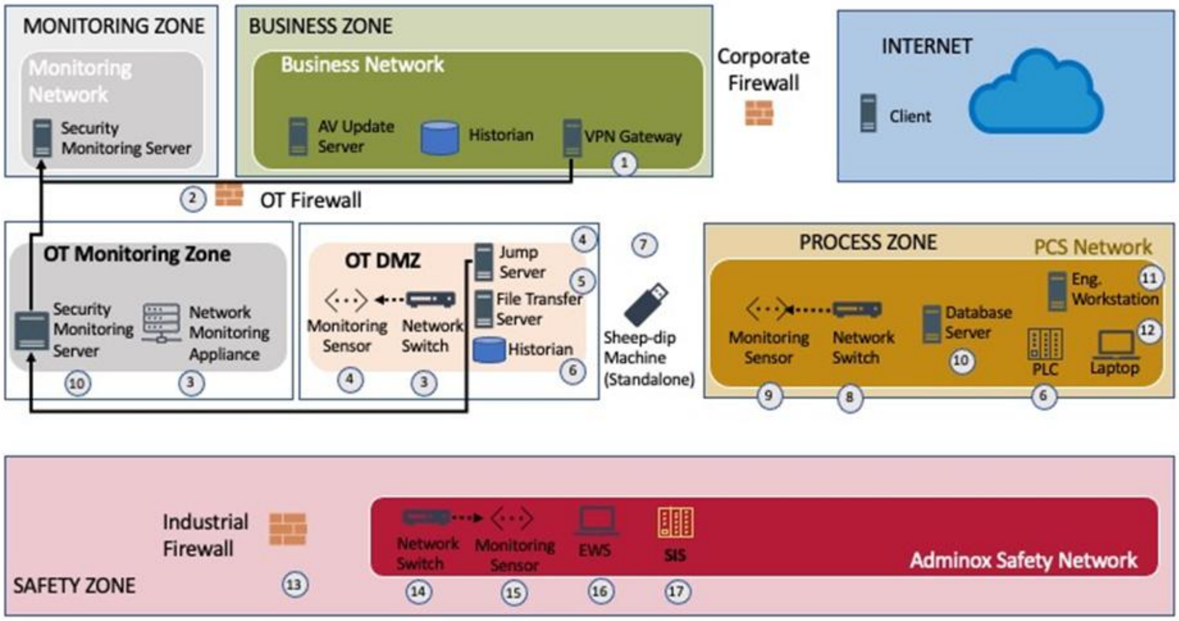

Figure 1- Example architecture model diagram showing zones/conduits.

Out of band

Collecting monitoring data via an out-of-band (OOB) network can significantly reduce the overhead associated with transporting this traffic over a production network. While there are notable advantages to this approach, it is also important to weigh the associated trade-offs.

Why Use Out-of-Band Networks?

Where possible, it is recommended to deploy OOB networks for the implementation of security monitoring tools and the forwarding of logs. The key benefits include:

- Monitoring traffic is kept separate, avoiding interference with operational processes.
- Providing the ability to capture malicious activity from the systems requiring monitoring.
- Enables structured, documented processes for log collection from air-gapped or standalone systems.
- OEM and vendor tools, such as syslog, may support out-of-band data collection methods.
- Facilitates controlled, secure log forwarding to MSSPs and other external providers.

Key Considerations and Trade-Offs:

However, deploying OOB networks also introduces other practical and security-related considerations:

- Additional planning, design, and ongoing maintenance is required, including consideration of local site capabilities and technical skills.
- There is a risk of unintentionally creating a ‘common’ network that undermines existing ICS/OT security boundaries.
- Running a secondary OOB network increases expenditure on hardware, site capacity, and security enforcement tools.
- OOB networks often fall outside the remit of standard operational teams, requiring clearly defined management responsibilities.
- Monitoring data may be collected in clear text and is not always natively encrypted, requiring additional security controls for safe transmission.

Tools available for use

Most solutions will have some method of data and log collection which is built into the service or hardware. This may have some limiting factors on how this data can be collected and consumed, to help address this there are open standards which are available to send data in a common format to collection agents or using application API frameworks. These may include:

- Windows Event Forwarding & Windows Event Collector.
- Dedicated ICS/OT monitoring solutions.
- Other Operating system built-in tools.

- OEM and native tools, including scripting functions.
- Open collection methods through syslog or similar.

Challenges

Collecting log data and processing is not a one size fits all, there are often challenges which may differ from operating system and hardware vendor. Where there is a standardisation on hardware and software this lends itself to a common approach, however more often than not, there is a mix of devices from different solution sets that provide access in diverse ways. Add into the mix the variety of ways to collect, store and process and this can often lend itself to agreeing on levels of compromise, especially where cost and safety is a factor. These may include:

- Pressure to migrate logging storage to cloud platforms.
- Networking architectures which do not allow low level collection of data from assets to be presented to a higher level.
- Hardware limitations on interfaces and data collection and aggregation methods.
- Dependency (or perceived dependency) on the SOC operating model:
  - In-house vs outsourced vs hybrid
  - Skills between IT/OT for threat analysis
  - Central SOC vs dedicated OT SOC or tenancy
- Collection of logs from different environments such as Programmable Logic Controllers, Firewalls, Windows, and Linux operating systems.
- Data Sovereignty for data stored outside of the ICS/OT environment on cloud.
- Access control to data and the storage of sensitive data on hosted platforms consumed as SAAS.

Further reading

There is lot to consider where logging and monitoring data can and should be stored, this document outlines some of the high-level principles that should be considered. This is the first in a short series of document which will explore where the best place is to store event data based on reference sites with differing security needs and goals. Other related reading includes NIST SP 800-92 - Guide to Computer Security Log Management and NIST SP 800-82 Rev. 3 - Guide to Operational Technology (OT) Security.

CAF Indicators of Good Practice Summary

This article discusses measures that contribute to the following CAF Indicators of Good Practice (IGPs):

- C1.b Securing Logs - You hold log data securely and grant appropriate access only to accounts with business a need. No system or user should ever need to modify or delete master copies of log data within an agreed retention period, after which it should be deleted.

Statement of Support

This guidance has been produced with support from members of the Industrial Control System Community of Interest (ICS-COI) for publication via the Research Institute for Trustworthy Inter-connected Cyber-Physical Systems (RITICS). This guidance is not intended to replace formal NCSC guidance where already available, and care has been taken to reference such existing guidance where applicable.

This document is provided on an information basis only, and whilst ICS-COI members have exercised reasonable care in compiling the guidance, they provide no warranty as to its accuracy, completeness, or suitability for any particular purpose.

To the fullest extent permitted by law, neither the ICS-COI or its members accept any liability for any loss, damage, cost, or expense arising directly or indirectly from the use of and / or reliance on, this document. Users of this guidance are advised to exercise their own judgement and consider taking independent professional advice.

Any reference to commercial products, services, or entities by name or otherwise, does not constitute or imply endorsement, recommendation or preference by the ICS-COI. The views and opinions expressed in this document shall not be used for advertising or product endorsement purposes.

Document Details

This document is version 1.0 and was published on 23/09/2025. It will be reviewed every 18 months.