



National Cyber
Security Centre
a part of GCHQ



Australian Government
Australian Signals Directorate

ASD
ACSC

AUSTRALIAN
SIGNALS
DIRECTORATE
Australian
Cyber Security
Centre



Communications Security
Establishment Canada
Canadian Centre
for Cyber Security

Centre de la sécurité des
télécommunications Canada
Centre canadien
pour la cybersécurité



Federal Office
for Information Security



National Cyber Security Centre
Ministry of Justice and Security



National Cyber
Security Centre
NEW ZEALAND

Secure connectivity principles for Operational Technology (OT)

How organisations should design, secure, and
manage connectivity in OT.



Table of Contents

- [Introduction](#)
- [Principle 1: Balance the risks and opportunities](#)
- [Principle 2: Limit the exposure of your connectivity](#)
- [Principle 3: Centralise and standardise network connections](#)
- [Principle 4: Use standardised and secure protocols](#)
- [Principle 5: Harden your OT boundary](#)
- [Principle 6: Limit the impact of compromise](#)
- [Principle 7: Ensure all connectivity is logged and monitored](#)
- [Principle 8: Establish an isolation plan](#)

This guidance has been developed with contributions from partnering agencies and is part of a series of publications aiming to highlighting the importance of cyber security in operational technology.

It is produced by the UK National Cyber Security Centre (NCSC) in partnership with the Australian Signals Directorate's Australian Cyber Security Centre (ASD's ACSC), the Canadian Centre for Cyber Security (Cyber Centre), the US Cybersecurity and Infrastructure Security Agency (CISA), the US Federal Bureau of Investigation (FBI), Germany's Federal Office for Information Security (BSI), Netherlands' National Cyber Security Centre (NCSC-NL), and New Zealand's National Cyber Security Centre (NCSC-NZ).

Introduction

Operational technology (OT) environments - which have long been centred on safety, uptime, and operational continuity - are now more interconnected than ever. Driven by the need for increased efficiency, agility, and integration, these advancements offer significant operational benefits (such as real-time analytics, predictive maintenance and remote monitoring & administration), but they also introduce risks.

Organisations deploying or operating OT systems often face challenges in prioritising cyber security due to operational constraints, such as dependence on legacy technologies that were never designed for modern connectivity or security requirements. These challenges are compounded by the increasing use of third-party vendors, remote access solutions and supply chain integrations, all of which expand the potential attack surface. In an OT environment, risks are elevated since a cyber intrusion can lead to physical harm, environmental impact, or potentially the disruption of an [operator of essential service \(OES\)](#).

Exposed and insecure OT connectivity is known to be targeted by both opportunistic and highly capable actors. This activity includes [state-sponsored actors actively targeting critical national infrastructure \(CNI\) networks](#). The threat is not just limited to state-sponsored actors with recent incidents [showing how exposed OT infrastructure is opportunistically targeted by hacktivists](#). Strengthening the cyber security of CNI, including securing OT connections, can challenge attackers' efforts and raise the threshold necessary to cause physical harm, environmental impact, and disruption.

Prioritising Systems

Due to potentially limited resources, organisations may not be able to complete all mitigating steps at once. When prioritising systems within your organisation, some topics that should be considered are:

- The role and impact of the device or process to your operations, including the ability to control and/or monitor key functions.
- The presence of fail-safe systems or redundant systems that maintain availability and reduce the risk of unsafe operating conditions or service outages.
- The time it would take to implement the change, including currently available funds and complexity. Keep in mind that the cheapest option may not be the most impactful option to securing connectivity.
- Active threat activity from attackers ranging in sophistication, including the consideration for current geo-political events and the potential national security significance of your organisation and/or your customers' organisations.

Principles-based guidance

This guidance outlines the desirable end-states that organisations should look to achieve when designing connectivity into OT environments. They are intended as goals rather than minimum requirements.

System owners should use these principles as a framework to design, implement, and manage secure OT connectivity, for both new and existing OT systems. These principles are particularly critical for operators of essential services.

Integrators and device manufacturers are encouraged to make these principles easier for organisations to achieve, through providing products that are secure by design, easy to implement and maintain. Integrators and manufacturers should ensure they are providing documentation to allow organisations assess connectivity risks. It is especially important that this documentation is available for 'turnkey' solutions, allowing operators to understand the design and implement appropriate security controls throughout the system's lifecycle.

Principle 1: Balance the risks and opportunities

Before you undertake any design work for new or existing connections to your OT environment, you must ensure you are equipped to **make risk-informed decisions** about when, how, and where connectivity is permitted within OT systems and to external/third party systems. Ensure these decisions are thoroughly documented to allow the reasoning to be auditable.

The first step for all OT connectivity should be the documentation of a formal business case to support decision-making. This should be stored centrally and referred to regularly during the design process. At a minimum the business case should document the following:

- **Requirement:** is the connection required and what does it aim to achieve?
- **Business Benefit:** what benefits arise from the added connectivity?
- **Risk Tolerance:** what cyber and operational risks are acceptable?
- **Potential Impacts:** what are the potential impacts of a compromise to this connectivity?
- **Introduced Dependencies:** will the connection make the system reliant on external services, making isolation harder in an incident?
- **Senior Accountability:** who is the senior risk owner for the new connectivity?

Tip: Define risk thresholds in the business case so future design and review decisions can be measured against agreed limits.

Note: To be able to effectively assess introduced dependencies and impacts you will require a definitive view of your OT architecture. To achieve this, use the [NCSC guidance on creating and maintaining a definitive view of your OT architecture](#).

For OT environments it is critical to give additional consideration within the business case for both risks to obsolete products and operational risks that could arise from increased connectivity:

Obsolete products

OT networks frequently contain obsolete products, commonly referred to as legacy products, attributable to extended system lifecycles. In this guidance, obsolete products refer to both software and physical assets. However, this does not include the use of legacy protocols by these

devices for device-to-device communications. This aspect will be addressed in [Principle 3 – Use standardised and secure protocols](#).

Using obsolete products compounds several related problems:

- **The product no longer receives security updates:** If developers are no longer providing [security updates](#), this increases the likelihood that vulnerabilities will be exploited by attackers.
- **The latest security mitigations are not present:** Older products may lack the latest security measures, increasing the impact of vulnerabilities, making exploitation more likely to succeed, and detection of any exploitation more difficult. This may include insufficient support for modern human-to-machine authentication mechanisms or modern cryptographic standards.
- **Unmanageable requirement for compensating controls, old knowledge and skills:** When manufacturers no longer support a product, organisations must invest in maintaining specialised knowledge and technical skills, either through in-house expertise or external contractors. This ongoing investment can divert resources from modernisation efforts and increase security costs due to the need for additional compensating controls. Additionally, the lack of readily available expertise for obsolete technologies can complicate incident response and recovery, making it harder to resolve issues quickly and effectively.

These combined issues heighten the risk of severe security incidents from low-skilled and low-capability attackers. [Obsolete products](#) should be treated as untrusted and should **not** be used to implement security controls.

Where obsolete products are present, segmentation and network controls may help manage associated risks. However, organisations should view these as temporary measures and assess if these measures are sufficient while establishing a timeline for asset replacement.

Note: If up-to-date software depends on an obsolete operating system, then the software must be treated as obsolete as the underlying platform remains insecure.

It is also possible to have 'self-established obsolescence', for example when (due to operational constraints) you choose **not** to update, leaving a device in a known vulnerable state.

Further reading

There is a range of guidance on managing obsolete systems. This includes:

- The NCSC's Device security guidance on [managing obsolete products](#).
- ASD guidance on [Managing the risks of legacy IT](#).
- BSI guidance on [Dealing with "End of Support" in industrial control and automation systems](#)
- The Five Eyes intelligence and security alliance, joint guidance [Secure by Demand: Priority Considerations for Operational Technology Owners and Operators when Selecting Digital Products](#).

Operational risk

OT connectivity should be designed with operational resilience in mind, and should not compromise the safety, reliability, or availability of OT systems. This means understanding how systems behave under failure conditions and ensuring that critical functions are not overly reliant on fragile or non-resilient links. When assessing connectivity, consider the following operational risk factors:

- **Unintended impacts on process or safety:** Could the connectivity introduce risks that affect the system's core functions or compromise safety mechanisms?
- **Loss of connectivity:** What would be the operational consequences if the communications links were to fail or become unavailable? Would critical processes be disrupted?
- **New interdependencies and single points of failure:** Are you introducing dependencies on systems or services that are not designed for resilience? For example, could a business system outage cause an OT process to shut down?
- **Manual fallback capability:** In the event of degraded or lost connectivity, can operations continue manually? Is there a clear and tested procedure for manual intervention?

At each stage of the design process you should assess if the connectivity is able to meet the risk-thresholds defined in your business case and that it aligns with your organisational threat context. In order to do this effectively you will need to verify your organisation has existing knowledge and processes to support this.

A comprehensive risk management framework

A risk management framework helps decision-makers identify, assess, and prioritise potential threats, enabling more informed and resilient systems. It also ensures consistency and accountability by providing structured processes for evaluating risks and implementing mitigation strategies. Use of a comprehensive risk management process ensures that connectivity decisions:

- are evaluated against **organisational risk appetite and threat context**
- consistently factor in **cyber, safety, and physical risks**

Your risk management framework should enable a holistic approach to risk that looks beyond the directly involved assets and data flows and considers the whole system architecture; its dependencies, existing security measures, and the policies and processes that govern it. To facilitate this [creating and maintaining a definitive record of your OT architecture](#) is critical. Building whole-system understanding will aid the selection of appropriate security controls that address the unique risks and constraints of your environment.

A whole-system understanding is essential for effective threat modelling. This approach enables you to evaluate how technical controls can mitigate connectivity-related risks and supports a proactive cyber security posture.

Further reading

- NCSC's [risk management framework](#) aligns with international standards such as the [ISO/IEC 27000](#) series. This includes sections on [system driven risk management](#) and [threat modelling](#).
- Additionally standards such as [ISA/IEC 62443-3-2](#) offers a tailored risk assessment methodology for Industrial Automation and Control Systems (IACS).
- The [IEC 62443 Zones and Conduits model](#) can help your organisation design secure OT architectures by grouping assets ('zones') with similar security needs, and defining controlled communication paths ('conduits') between them. This approach supports secure connectivity by enabling you to draw clear trust boundaries within your environment.

Effective control and oversight of your supply chain

The supply chain plays a critical role in OT security, with a wide range of third parties often involved in the design, integration, and ongoing maintenance of systems. The NCSC has established [supply chain principles](#) for managing this risk.

It is particularly important to manage supply chain risk when procuring new products. Ensuring that devices are [secure by design](#) and developed following a secure product development lifecycle. This helps reduce the risk of introducing vulnerabilities through third-party components or insecure design practices. Supply chain factors that may affect your ability to implement effective secure connectivity could include:

- **Ability to influence:** Do you have the ability to influence the security controls architected into the supplier's solution?
- **Contractual controls:** Does your contract allow you to define and enforce minimum product security requirements, including capabilities for updating, access control, [digital forensics and protective monitoring](#)?
- **Component visibility:** Do you know all technologies in the supplied product? Hidden or undocumented devices may alter your connectivity model. This sometimes is referred to as a bill of materials.
- **Supplier trustworthiness:** Consider the supplier's policies, how they protect development and maintenance environments, and how they handle your designs/configurations.
- **Supplier track-record:** Has the supplier previously demonstrated that they respond to security issues and incidents in a mature manner, and do they positively engage with vulnerability researchers.

Further reading

- CISA's [Secure by Demand](#) publication outlines priority considerations for OT owners and operators when selecting digital products, emphasising the importance of devices being secure by design.
- ACSC's [secure by design publication on choosing secure and verifiable technologies](#) can further assist procuring organisations to make informed, risk-based decisions within their own operational context.
- Additionally, [IEC 62443-4-1](#) provides vendor-specific requirements for secure development practices in Industrial Automation and Control Systems (IACS), offering a robust framework for evaluating supplier security maturity.

Principle 2: Limit the exposure of your connectivity

Exposure refers to where an asset sits within the wider system architecture, and how accessible it is to external or adjacent networks, taking into account defence in depth controls. Ultimately, the more assets exposed at the network edge, the broader your attack surface becomes.

To manage this risk, organisations should adopt an **exposure management approach**. This involves proactively identifying, assessing, and mitigating risks associated with digital assets and their connectivity. This includes evaluating the asset's placement in the network, the type of connectivity implemented, and the strength of cyber security controls. The NCSC Netherlands [exposure management guidance](#) defines categories such as 'internet or external facing' and 'adjacent network-facing' assets, helping organisations assess exposure levels systematically.

It is especially critical to ensure that you are managing the exposure of your admin interfaces. Where possible limit administration of devices or systems to only be achievable through [privileged access workstations \(PAW\)](#), which provide secure and trusted endpoints for system management. When administering critical security controls or obsolete systems via a PAW it may be appropriate to limit administration to only be achievable via local physical access to further reduce the exposure of these interfaces.

Note: Network edges in OT can be hard to identify. For example, an unsecured radio link inside an OT network may not appear on network diagrams but still forms part of the network edge. Ensure you identify all data flows within the OT network and the components that facilitate this connectivity.

Reduce time of exposure

Not all connections need to be continuously active. Where possible, use just-in-time access, enabling connectivity only when required and disabling it otherwise. This significantly reduces the window of opportunity for attackers.

Remove inbound port exposure

All connections with the OT environment should be initiated as outbound connections from within the OT environment. This principle helps avoid exposing inbound ports on the OT network perimeter or between internal zones, which can significantly increase security risk.

In scenarios where systems outside the OT environment require access to OT assets (for example remote vendor support), use brokered connections through a secure gateway located in a separate, security-controlled segment such as a demilitarised zone (DMZ). A brokered connection is a method where the external party connects to an intermediary system (the broker), which then securely relays the connection to the OT asset. This ensures that the OT system is never directly exposed to the internet or external networks, and that all access is mediated, monitored, and controlled.

Note: The security of the brokered connection is critical. It must be actively updated and use modern authentication methods.

Manage obsolescence risks

Obsolete devices pose a known and increasing security risk, making them unsuitable for direct external connectivity beyond the OT network boundary.

However, operational constraints often mean that migrating away from obsolete devices takes time. During this transition period, these devices may still need to communicate with other systems or receive vendor support.

To manage the associated risks, organisations should enable indirect access to external networks through compensating controls, including:

- **Network segmentation:** isolate the obsolete device from the wider OT network using logical or physical segmentation to limit lateral movement and reduce exposure.
- **Trusted boundary controls:** place up-to-date, security-hardened components between the obsolete device and external systems. Examples include:
 - a **protocol gateway** to translate and inspect traffic before it reaches the device
 - a **hardened jump host** to facilitate vendor support, ensuring access is controlled and monitored

- **Access restrictions:** limit connectivity to only what is operationally necessary.
- **Monitoring and Logging:** ensure all interactions with the obsolete device are logged and monitored for anomalous behaviour.

Manage unique connectivity bearer risks

The risk associated with a connection depends on the network type and transmission medium. For example:

- public internet links carry higher exposure than private fibre
- wireless technologies, even within private networks, may introduce risks of unauthorised access

Security controls should be tailored to the connectivity type to ensure adequate protection.

External attack surface management

Your public visibility can be monitored actively by using [external attack surface management \(EASM\)](#) tools or other internet-facing asset discovery tools, to identify accidental or unmanaged exposure before attackers do. These discovery tools index internet-connected assets and protocols, allowing anyone to find exposed web servers, remote access portals, or industrial devices. EASM tools can be used to reduce your attack surface and the likelihood of being targeted.

If your systems are visible to these scanning services then they are highly likely to be found and targeted by malicious actors, significantly increasing the risk to these systems. There are several things you should consider before using tools to manage your exposure:

- **Static public IP address space:** you must maintain a thorough and current list of all public IP addresses alongside exposed network ports in use. This task can be particularly challenging, especially for organisations that operate across vast geographical areas or use multiple ISPs, or have several infrastructure teams.
- **Dynamic public IP addressing:** particularly where your organisation uses cloud services, you may use dynamic addressing. This can increase the challenge of automated scanning. It is critical that you establish a process to maintain an accurate record of these addresses.
- **Continual monitoring:** you should establish ongoing processes to automate exposure management within your organisation. This should include scanning the full public IPv4 and IPv6 ranges belonging to your organisation, not just the Ips currently in use.

- **Third parties:** if connectivity to your environment is provided by a third party, then it may be proportionate to include these endpoints in your exposure management program. This would require establishing processes with the third party for sharing IP addresses of endpoints within or directly supporting your OT systems.

Any device found through a EASM, or other internet-facing asset discovery tool should be deemed at risk and promptly investigated. This should include establishing:

- Was the system originally designed to be connected for direct internet access?
- Was this system designed to be exposed temporarily or permanently, and how long has it been exposed for?
- Are there any additional services/ports/protocols exposed by the system to those expected?
- Does the system have security controls implemented that are appropriate to its exposure?
- Is the system is updated and hardened?
- Does this systems exposure cause additional risk to other connected OT systems?
- Does the system need to be exposed directly to the public internet, or can connectivity be restricted to only those who need it?
- Is the asset and/or its internet connectivity critical to the delivery of your OT process?

Where these factors highlight that an asset is not designed or lacks sufficient security controls to operate in the threat context of the public internet, remediation actions should be immediately taken. Depending on your operational constraints this could include:

- disconnecting the asset
- reconfiguring/Updating the asset
- adding compensating controls to manage the exposure risk

Further reading

International partners have published a range of resources on this topic:

- [CISA Internet Exposure Reduction Guidance](#)
- [NCSC-NZ Cyber Security Guidance: Preventing unintentional operational technology device exposure](#)

Managing wireless networks exposure

It is easy to think of exposure only at a device level, but the communication medium also needs to be considered. Devices that use wireless communications or send traffic using wireless signals can be exposed to a greater level of risk.

A wireless network is not constrained by the boundary of your site. This means physical controls can't be relied on to protect this network. If a wireless network is not configured with security in mind, it is trivial for an attacker to intercept and capture signals, as well as to potentially inject traffic into these networks.

Understanding the factors that impact security of your signals is critical to enabling you to [build appropriate and proportionate controls](#).

Principle 3: Centralise and standardise network connections

The connectivity models of OT systems can be complicated, involving various stakeholders such as business systems, billing platforms, and external vendors responsible for ongoing maintenance. As organisations evolve, these connectivity models often become more complex, adapting to new business requirements or integrating modernised processes. This increasing complexity can significantly expand the organisation's attack surface, making it harder to monitor, control, and secure communications across the OT environment. Each additional connection, especially if implemented in an ad hoc or bespoke manner, introduces potential vulnerabilities that attackers can exploit.

Centralising and standardising connectivity helps address this challenge by consolidating access points and enforcing uniform security controls across the OT estate. A centralised architecture enables consistent monitoring, logging, and enforcement of security policies, making the management overhead of cyber security easier. Standardisation ensures that all connections follow a repeatable and well-understood pattern, reducing the risk of misconfigurations and simplifying the deployment of protective measures such as encryption, authentication, and segmentation.

To effectively manage your attack surface, it is essential to ensure that your OT remote connectivity should be **flexible**, **repeatable** and **categorised**.

Flexible

Controls must be regularly assessed and refined to keep pace with emerging threats. A threat-informed approach should include routine reviews of threat advisories and an understanding of how adversaries exploit connectivity to gain access or disrupt functionality.

Organisations should maintain robust change management processes and the agility to transition to new solutions when existing ones become outdated. Flexibility also means selecting products that offer ongoing support for new security controls, enabling the organisation to adapt as threat models and regulatory requirements change.

Where third parties require access through remote connectivity, flexibility should be embedded within contractual agreements to accommodate evolving security requirements.

Repeatable

Connectivity should be robust and reusable, minimising the need for bespoke solutions for each use case. New connectivity implementations should avoid duplicating existing routes into the network, thereby reducing the overhead associated with deployment, updates, and maintenance.

For example, rather than deploying separate virtual private network (VPN) endpoints within the OT network for each third party, centralise remote access through a secure solution hosted in the DMZ. This allows for consistent enforcement of access controls, session monitoring, and reuse of a single hardened access path across multiple vendors.

Organisations may have legacy products or brownfield deployments that do not align with new connectivity patterns. A clear process should be established to manage interim risks and migrate these systems within a sensible timeframe.

Categorised

Security controls should be tailored to the nature of the data flow. Categorising connectivity types helps identify the most appropriate and proportionate controls. For example, distinctions should be made between human-to-human, human-to-machine, and machine-to-machine interactions.

Categorisation supports the application of targeted protections and ensures that connectivity is aligned with operational and security requirements.

Principle 4: Use standardised and secure protocols

In addition to securing the networks and devices used to establish communications, your organisation must also consider the security of the protocols employed.

As outlined in [Creating and maintaining a definitive view of OT architecture](#), it is common for industrial environments to prioritise availability over the confidentiality and integrity of communications. It is essential that all components of the confidentiality, integrity & availability (CIA) triad are considered. However, you may prioritise different aspects of CIA depending on the connection. For instance, in field networks, authentication and integrity are essential to limit an attackers' ability to send malicious traffic. Conversely, in north-south traffic at network boundary points, encryption becomes critical to prevent attackers from discerning information on how to impact the system.

Protocol validation

Protocols used within and between your network environments should feature data formats to enable simple validation of both the protocol and its data. To ensure that traffic is expected and within acceptable bounds, schemas should be applied to these data flows. Data simplicity is key as it reduces the attack surface and makes it more difficult for adversaries to inject malicious data.

Schemas should be used to inspect and verify protocols and data payloads at key trust boundaries. These boundaries may exist between networks (for example the OT/IT boundary) or between services (for example in front of SCADA control software or a programmable logic controller). Ideally, verification should be schema-based and follow a 'known good' model, only allowing traffic that conforms to expected structures and values.

When validating data, consider any nested or encoded content. For example, a field might contain a base64-encoded value. A basic check might verify the length or format of the base64 string, but a more robust approach would decode the value and validate the underlying data against expected patterns or constraints.

Industrial protocols

When evaluating industrial protocols within your OT environment, you should:

- Default to the latest secure versions of industrial protocols (e.g. DNP3 to DNP3-SAv5, CIP to CIP Security, Modbus to Modbus Security, OPC DA to OPC UA).
- Ensure that protocols support cryptographic protections for authenticity and integrity, such as digital signatures. Your protocol use should be flexible, enabling you to update to new cryptographic algorithms. In particular ensure protocols support 'crypto agility', the ability to switch and update cryptographic algorithms, to ensure the lifetime of the product is matched by the lifetime of the cryptographic algorithm. Where there are no hardware constraints, crypto agility is likely to enable you to migrate to [post-quantum cryptography](#) algorithms when and where required.
- Prefer protocols that support open standards and interoperability to facilitate vendor-agnostic solutions. This approach can help reduce the number of bespoke data flows between environments, thereby decreasing complexity within your architecture. Pay special attention to security functionality as interoperability issues can break the entire communication stack rather than a specific proprietary function.
- Require a business case for the use of insecure protocols within your environment, making their use the exception rather than the norm. For instance, in time-critical safety applications, encryption may not be feasible. However, compensating controls should be implemented to manage associated risks, and it is critical these controls are documented as part of your risk management framework.

Where your organisation has insecure industrial protocols in use, you should establish a roadmap for migration to secure industrial protocol variants. This will enable you to make considerations to enable this in asset uplifts and system maintenance.

Tip: Use resources published by your manufacturer to support your development of a migration plan. This could include:

- direct support from engineers that understand your current deployment
- use of public materials on migration or new solutions from vendors
- evaluation of the manuals and specification products you already own to identify supported protocols

Industrial control protocols (Modbus, OPC DA, EtherNet/IP, etc.) should be restricted to isolated OT network segments. External connections for data exchange between OT and IT should be brokered through a DMZ and use secure, standardised protocols designed for interoperability (such as OPC UA over TLS, MQTT over TLS, HTTPS). Where operational data needs to be shared, replicate the OT historian to a historian instance in the DMZ via a unidirectional, secure transfer mechanism, ensuring no inbound connectivity from IT to OT. IT systems should query the DMZ historian via a [secure HTTP-based API](#) with strong authentication, rather than directly accessing OT systems.

Principle 5: Harden your OT boundary

The prevalence of obsolete assets and weak security controls within the OT environment makes hardening the OT boundary critical. Network segmentation and segregation remain key controls to reduce exposure, where built-in protections at the device or protocol level may be limited. Although these measures provide a robust first layer of defence, they are even more effective when combined with native security capabilities within OT systems. This could include secure by design devices and authenticated communication protocols.

Note: In certain sectors, the security of connectivity patterns may be set by regulations. In some instances, these patterns can restrict your ability to choose the security controls implemented.

Regulated organisations should ensure they adopt new secure patterns, introduced by the relevant industry body or regulator, as quickly as feasible. If the security controls established by the regulator or industry body fall short of the standards anticipated by your organisation, you should implement compensating controls, such as the hardware flow controls outlined in this section, to mitigate residual risk.

Regulators and supporting industry bodies should be routinely updating these patterns to improve their security as the threat landscape evolves, and new controls become available. Old connectivity patterns should be deprecated and all users moved to the new secure pattern.

Because many OT systems are difficult to update or replace, the boundary becomes the primary defence against external threats. Organisations should therefore invest in modern, modular, and easily replaceable boundary assets. This could include deploying a firewall with application-layer (Layer 7) inspection capabilities, often referred to as a next-generation firewall. These assets offer greater flexibility for patching, upgrading, and reconfiguring security controls. Importantly, they can be maintained without disrupting core OT operations. This makes them essential for adapting to evolving threats and maintaining long-term resilience.

Your OT boundary security controls need to be flexible to enable your boundary to evolve to meet changing threats; if a device or system is directly exposed to an external network it has no additional lines of defence. This means that the failure or compromise of a single control or measure would result in an attacker gaining immediate and complete access to the system.

When designing security across all connections, a layered approach, commonly referred to as [defence in depth](#), should be prioritised.

Note: It is critical that devices facilitating connectivity with external services and networks are designed for this threat context and have suitable technical controls. It is critical that OT boundary assets are:

- not obsolete
- routinely [updated by default](#) with the latest firmware
- replaced before they reach their end-of-life (EOL)
- **developed following a secure development lifecycle**, ensuring they meet modern security standards

Given their location at the edge of your network, these devices are particularly susceptible to exploitation by attackers if known vulnerabilities are present.

The [NCSC's vulnerability management guidance](#) provides detailed advice on managing risks throughout an asset's lifecycle.

The [DSIT Software Security Code of Practice](#) can help assess whether vendors develop software securely.

To harden your OT boundary you should consider the following:

- **Remove unused services and ports:** Devices should be configured to only expose required services within the network to reduce their attack surface.
- **Implement phishing resistant multi-factor authentication (MFA) for external services:** Access to human-to-machine connectivity should require MFA in order to protect sensitive information and prevent unauthorised control actions. The NCSC's guidance on [Multi-factor authentication for your corporate online services](#) can help you to implement strong methods of MFA.
- **Change default passwords:** No devices deployed in your network should have default passwords, as using default credentials makes the compromise of assets trivial. This is particularly important for devices exposed to the public internet, or your external networks. Your organisation should ensure it has a [password policy](#) in place that sets out ways to generate and store secure credentials.
- **Enforce the principle of least privilege:** Both human-to-machine and machine-to-machine connectivity should follow the concepts of least privilege. Only the permissions required for the action should be granted. Where this is human-to-machine the connectivity should be user aware so actions produce an employee specific audit trail. Human-to-machine permissions should also be integrated into joiners, movers, leavers (JML) processes to ensure access rights are revoked if the user's role changes or they leave the organisation. The NCSC has guidance on [privileged access management](#) that covers this in more depth.

- **Use context aware access:** Where available you should implement context aware controls on external connectivity. These controls make decisions on connectivity based on a number of factors such as device location, device type and configuration and user pattern of life.
- **Enforce security requirements on third parties:** When a third party designs and implements connectivity into your OT environment, it is crucial to ensure that the solution aligns with your organisation's security expectations. For more information, please refer to [principle 5 in the creating and maintaining a definitive view of your OT architecture guidance](#).
- **Enforce uni-directional traffic flows:** Where possible, establish outbound only uni-directional connectivity, to minimise risks from external connections. This approach reduces the potential impact of external systems on your OT systems and ensures their independence from outside influences. At its simplest form this could be using uni-directional protocols such as UDP and network policy enforcement. For high-threat or high-risk environments you may wish to consider further hardware-based security controls.

Hardware-based controls

- **Cross Domain solutions:** Cross Domain is a methodology and architectural approach to enabling risk-managed bi-directional data flows across trust boundaries. A Cross Domain solution represents a collection of security controls to enable a specific data flow, providing security across the layers of the network stack, backed up by products with hardware security controls at key boundaries. The NCSCs [Cross Domain principles](#) can be used to guide design, development, assessment and deployments of these as a security control.
- **Data diodes:** Data diodes aid in establishing assurance of uni-directional data flows, through physically enforced directionality designed into the hardware. This feature ensures that bi-directional communications cannot be accidentally re-enabled due to misconfigurations at either the device or network layer. It is important to understand that:
 - a diode solely ensures data directionality; it does not encompass the broader security measures typically present in a Cross Domain solution.
 - an architecture where a diode ingests untrusted content, followed by software parsers on the more trusted side to 'make the content safe' is **not** considered Cross Domain; to be Cross Domain, we would expect at the minimum for the untrusted content to be structurally verified in hardware to reduce the risk of software parser vulnerabilities.
 - a diode in each direction (to enable bi-directional communications such as APIs) with software on the more trusted side to process untrusted content and orchestrate bi-directional flows is considered an anti-pattern.

Tip: Data diodes can be a useful control when ingesting hard-to-inspect data into isolated networks e.g. ingesting logs/network captures into an isolated security monitoring environment.

Principle 6: Limit the impact of compromise

Modern OT networks should be designed with controls that extend beyond the OT boundary. These should implement layered controls that reduce the impact from insider threats, third parties and external compromise. For OT connectivity these layered defences should focus on two main risks:

Contamination

This refers to the unintended or unauthorised introduction of malicious code, compromised data, or insecure configurations into a trusted environment. Contamination can occur through infected devices, vulnerable software updates, or poor operational hygiene.

For example, if a malware-infected laptop is connected to a production network, the malware may propagate across systems, undermining their integrity and security. Contamination can also result from misconfigured devices or outdated firmware, which attackers exploit to maintain persistence and evade detection.

Lateral movement

[Lateral movement](#) is the process attackers use to expand their reach after gaining initial access. It involves internal system mapping, compromising additional hosts, and escalating privileges to control critical systems. Attackers often use stolen credentials to access more systems, enabling data exfiltration or sabotage.

For example, in a flat OT network with VPN access for vendors, a contractor connecting their designated equipment could also gain access to all other devices on the network. Without layered controls, a motivated insider could compromise multiple systems and disrupt OT processes.

It is important to consider that although lateral movement is often thought of in the context of external attackers, lateral movement techniques can be used by insider and third party threats that already have a foothold on the network. [Identifying and mitigating 'living off the land' techniques](#), where attackers use legitimate tools and processes, is critical to preventing lateral movement risks.

Note: OT gateway devices, serial gateways, and network switches aggregate multiple assets, making them high-value targets for machine-in-the-middle attacks. An insecure gateway can also offer attackers a persistent point of presence on your network to enable further attacks.

These devices form a critical role in the security of the OT network. To maintain cyber resilience and reduce exposure to known vulnerabilities, such devices should be subject to regular updating, robust configuration management, and timely replacement before reaching end-of-life or becoming unsupported by vendors.

Further reading:

- The NCSC has published guidance on [preventing lateral movement](#).
- The [ICS MITRE ATT&CK® Framework](#) provides detailed examples and mitigations to lateral movement techniques specifically for OT systems.

Segmentation

One of the most effective strategies to [reduce the impact of compromise](#) is to implement a **zoned or segmented network architecture**. By dividing the network into smaller, functionally isolated segments, organisations can contain threats within the zone where they originate.

Micro-segmentation

In OT environments, micro-segmentation offers a more granular approach to network segmentation by dividing zones into smaller units based on specific workloads, applications, or device functions. Unlike traditional segmentation, which typically separates large network zones (for example IT vs. OT, or control vs. monitoring), micro-segmentation applies controls at a finer level, often down to individual devices, services or protocols. This approach allows for highly targeted traffic policies, enabling organisations to restrict communication paths to only what is strictly necessary.

For example, a sensor may only be permitted to communicate with its associated controller, but not with other devices in the same zone. This significantly reduces the attack surface and limits the potential for lateral movement within zones.

Micro-segmentation is particularly valuable in environments with mixed trust levels, legacy systems, or varying security requirements. It supports [zero trust principles](#) by enforcing least privilege access and ensuring that even within a zone, traffic is subject to inspection and control.

Separation of duties:

Separation of duties ensures no single system, role, or individual has complete control over all aspects of a critical function. In OT environments, this means dividing responsibilities and access across systems and users to reduce the risk of accidental or malicious actions: and to limit the impact of compromise.

Applying separation of duties in OT helps reduce exposure and limits risk propagation. Organisations can contain potential compromise and improve resilience by **functionally separating systems** especially those involved in control, monitoring, and business operations. It also supports auditing and accountability by clearly defining which systems and roles are responsible for specific actions.

For example, you should ensure that monitoring, analytics, and business systems do **not** have direct control capabilities over OT assets. These systems should be designed to observe and analyse, not command or alter operations.

Browse Down

The browse down principle dictates that you should trust your administration device as much as (or more than) the system you are managing. Implementing the browse down pattern correctly is vital to ensure that a connectivity compromise does not enable an adversary to alter systems or security controls and policies. The NCSC has further guidance on [gaining trust in your management devices](#) and [privileged access workstations \(PAWs\)](#).

Boundary Controls

Organisations should implement strong controls at the boundary of their OT environment. These controls should reside in a separate network segment, often called a DMZ. A DMZ acts as a buffer zone that isolates external-facing systems (such as remote access gateways, update servers, or vendor portals) from the core OT network.

This architectural separation ensures that any compromise of externally connected systems does not directly expose OT assets. It also enables tighter control over traffic entering and leaving the OT environment, supporting inspection, logging, and policy enforcement.

You should strictly regulate traffic between OT zones. This can be achieved using techniques such as:

Host-based controls

Host-based firewalls operate directly on individual devices. They enforce rules based on:

- source and destination IP addresses or MAC addresses
- ports and protocols (e.g. TCP/UDP)
- connection flags (e.g. SYN, ACK)
- directionality (inbound vs. outbound)

These filters are essential for enforcing local security policies and should default to a 'deny all' posture, with only explicitly authorised traffic permitted. Suppliers should assist in defining appropriate rulesets tailored to the device's operational role. These should be viewed as your last line of defence for limiting network flow to a device as part of a layered defence model.

Static network controls

Route filtering and access control lists (ACLs) play an essential role in enforcing zone boundaries and limiting unnecessary connectivity in segmented OT networks.

- **Route filtering** operates at the routing layer, controlling which network paths are advertised, accepted, or propagated between zones. This prevents unintended or insecure routing of traffic across segments and helps maintain the integrity of the network architecture by ensuring only authorised networks are reachable.
- **Access Control Lists (ACLs)**, typically implemented on routers and switches, function at the packet forwarding level. They provide a straightforward mechanism to permit or deny traffic based on criteria such as IP addresses, ports, and protocols. While ACLs are less granular than Deep packet inspection (DPI) and do not support **stateful inspection** (that is, they do **not** track the state of connections), they are effective for enforcing basic perimeter rules such as blocking external access to sensitive OT devices, or restricting inter-zone communication to specific services.

Static network filtering and routing controls help reduce the attack surface and enforce segmentation policies, particularly in environments where simplicity, performance, and predictability are critical.

This is the minimum level of control that your OT environment should implement between network zones.

Dynamic network controls

Dynamic network control mechanisms provide intelligent, context-aware enforcement of traffic policies across segmented OT environments. These mechanisms evaluate traffic based not only on fixed attributes, but also on connection state, protocol behaviour, and command-level content.

- **Stateful filtering** enhances traffic control by tracking the state of network connections over time. Unlike stateless controls that inspect each packet in isolation, stateful inspection understands whether a packet is part of a legitimate, established session. This allows for dynamic rule enforcement, for example, permitting return traffic for an authorised request without requiring an explicit rule for the response. By maintaining context, stateful filtering reduces false positives, improves security, and supports more adaptive segmentation policies.
- **Deep packet inspection (DPI)** adds further granularity by analysing the full payload of network packets, enabling interpretation of protocol-specific commands. This is particularly valuable in OT environments where control over specific operations (such as read versus write commands) is critical. DPI can be integrated directly into security devices like layer 7 application firewalls to actively block traffic based on its content. Alternatively, it can function as a passive control, alerting operators to unexpected or suspicious commands when embedded in monitoring tools or intrusion detection systems. In encrypted channels, such as those secured with TLS, effective DPI requires interception and decryption, which means considering where cryptographic protections apply and ensuring the inspecting device is fully trusted. Any security control that handles decrypted, plaintext traffic is in a position of significant trust: if that device is compromised, the confidentiality and integrity of the protected communications can be undermined.

Together, stateful filtering and DPI form a dynamic layer of traffic control that complements static mechanisms. They enable more precise enforcement of segmentation boundaries, particularly in OT environments where protocol-specific control and connection awareness are essential to maintaining operational integrity and resilience.

Threat detection and response:

While static and dynamic controls enforce segmentation boundaries, threat detection and response systems, such as intrusion detection systems (IDS) and intrusion prevention systems (IPS), offer an additional layer of defence by identifying and addressing threats that may bypass or exploit those controls. Unlike DPI and stateful filtering, which enforce traffic control, IDS and IPS focus on identifying and responding to threats that may exploit weaknesses in those controls.

The use of IDS/IPS can enhance segmentation by providing an additional layer of protection for OT networks:

- **IDS** is a passive system for monitoring traffic for known attack signatures, protocol violations, or behavioural anomalies, alerting operators to suspicious activity
- **IPS** goes further by actively blocking or quarantining malicious traffic in real time

Both systems often leverage DPI to gain deep visibility into traffic content, but they also use other techniques such as signature matching, anomaly detection, and protocol validation. When deployed at critical network boundaries, IDS/IPS reinforce segmentation by ensuring that only legitimate and safe traffic is permitted, and that any attempts to exploit vulnerabilities or bypass controls are swiftly detected.

Principle 7: Ensure all connectivity is logged and monitored

Even with all possible precautions in place, there remains a risk that your system could be compromised. Monitoring is your last line of defence when designing secure connectivity.

It is critical that your organisation [makes compromise detection easier](#) by implementing comprehensive logging and monitoring throughout your OT environment. These logs will help your organisation establish a baseline of 'normal' activity, allowing operators or detection systems to identify abnormalities faster.

The end-goal of logging should not just be to collect logs. Instead, you should understand how attackers may seek to exploit your systems through identifying weak points. Then design monitoring and alerting to help identify potential attacks. This can help guide what logging or packet captures you need to support these monitoring and alerting rules. Within OT environments, specific considerations should be made regarding how logging addresses:

Unauthorised activity:

Changes to OT environments are managed through strict controls, including detailed planning, change logs, and advance notifications. During maintenance, monitoring rules may be temporarily disabled to reduce false positives. Keeping your SOC informed ensures alerts for maintenance data flows are re-enabled outside planned windows. Work management tools can support this visibility. However, business processes should verify maintenance actions are legitimate to prevent attackers from exploiting this system to hide malicious activities.

Anomaly detection:

This refers to the process of identifying patterns or activities that deviate significantly from normal or expected behaviour within a system or network. This approach can be especially advantageous in OT systems, where there are relatively static, repetitive processes characterised by consistent command structures. However, it is critical that anomaly detection does not replace the implementation of technical controls designed to disable or block unused command sets, services, or ports.

Break-glass:

Break-glass access, where typical security controls are bypassed for safety incidents, is intended solely for emergency situations and should not be used as a standard remote access method. It is critical that any attempt to use a break-glass account triggers the highest criticality alarm within your Security Operations Centre (SOC).

Data flow monitoring:

Continuous monitoring of data flows within and between network segments is crucial to enable you to validate segmentation policies, and identify early signs of compromise or misconfiguration in your controls.

Further reading:

This is an area that there is already extensive [guidance on logging produced by the NCSC](#) and the wider topic of [building a security operations centre \(SOC\)](#).

For monitoring of external facing systems, our [External Attack Surface Management \(EASM\) buyers guide](#) outlines the features you should look for in these tools.

Manufacturers have a role to play in ensuring standard logging and forensic features that are robust and 'secure by default', so that network defenders can more easily detect malicious activity and investigate following an intrusion. The NCSC has produced [guidance on digital forensics and protective monitoring specifications for producers of network devices and appliances](#) to aid in assessing these characteristics.

Principle 8: Establish an isolation plan

In certain circumstances, it may be necessary to isolate OT environments from external influences. This need can arise from various factors, including increased threats or confirmed compromises within connected systems.

The isolation process for the system should be considering any potential impacts to wider business or any national interdependencies. This plan should be linked to and part of your wider [business continuity plans](#). It should be regularly tested to ensure that the system works as intended, and does not impact your organisation's services.

Note: Isolation plans should include an understanding of your contractual arrangement with third parties and suppliers. This could include requirements such as the ability to switch from remote support to having to physically attend the site.

OT systems that provide critical functions should, where possible, be designed to facilitate isolation, allowing them to function independently of external dependencies. It is essential to incorporate isolation planning into the system design process to prevent any unintended consequences that may arise from isolation measures.

For organisations managing multiple sites, it is important to develop not only site-specific isolation plans but also comprehensive strategies that address large-scale isolation needs. These strategies should consider scenarios where a crucial infrastructure component is compromised, which could lead to potential lateral movement across all sites.

When planning large-scale isolation, it's important to identify critical data flows that must remain operational. Some data flows from sites may need exemptions from isolation measures, especially where losing these data flows could cause a national-level impact or create unsafe operating conditions. Technical controls, such as data diodes, should be in place to allow these flows to safely operate during a compromise.

There are three primary isolation strategies:

Site isolation

This strategy is applicable when you are managing a site built on a flat network, or one that has restricted security measures. In this approach, the options for isolation are primarily confined to removing all external network connections, either through physical disconnections or by modifying network rules. If using modification of network rules for isolation, then the network appliance must be a secure and up to date device.

Application/service-specific isolation

This strategy is applicable when you have implemented secure connectivity and network controls, as outlined in this guidance. It enables you to isolate specific affected services and network routes, thereby minimising the potential impact of security incidents. For example, if you become aware that a third party with remote access to your environment has been compromised, a well-architected 'just in time' access model allows you to temporarily revoke their access. You can maintain connectivity with other third parties while re-enabling the disabled connection once the risk has been mitigated.

Site isolation with hardware-enforced trusted communications

If your site architecture involves hardware-enforced data flow security controls, you may be able to safely maintain these data flows while isolating other non-hardware-enforced data flows. For instance, if you use data diodes to transfer telemetry and logging from the environment, you may be able to keep this link up to maintain visibility while isolating the rest of the site. Such an approach can enable the continuation of business and/or national functions while effectively isolating external influences from the environment. However, you may need to implement additional monitoring processes temporarily to ensure that the incoming data aligns with the expected value ranges.

Investing in cyber security controls to effectively manage the risk of connectivity can help minimise the operational risks associated with isolation. If you can trust that services are properly isolated, you can take a more targeted approach to OT security and reduce business impacts.

© Crown copyright 2026. Photographs and infographics may include material under licence from third parties and are not available for re-use. Text content is licensed for re-use under the Open Government Licence v3.0.

(<https://www.nationalarchives.gov.uk/doc/open-government-licence/version/3/>)



NCSC.GOV.UK



@NCSC



@CYBERHQ



@CYBERHQ



National Cyber Security