INDUSTRIAL CONTROL SYSTEMS
Community of Interest

Guidance on Securing the Border of your Industrial Control System/ Operational Technology environment.

Introduction

In today's rapidly evolving landscape of technological innovation, Industrial Control Systems (ICS) / Operational Technology (OT) has become an increasingly critical component of UK CNI business operations.

To accommodate this, the landscape of both ICS/OT and Information Technology (IT) has undergone significant transformation. A prevailing trend in recent years is the convergence and overlap of technologies used in IT and ICS/OT. Common use cases include:

- Access to third-party systems
- Access to company resources in the cloud
- Remote access for users
- Inter-connecting two or more ICS/OT systems.
- Exporting ICS/OT data for use by wider business systems

However, not all organisations use this technology convergence to drive business level integration; preferring to maintain a distinct separation between IT and ICS/OT due to concerns related to security, safety, and system uptime, considering it a crucial element of their risk-mitigation strategy.

Irrespective of the chosen approach, organisations across industries seek to leverage technology advances to enhance their operations, adapting their unique journey based on industry-specific requirements, regulatory compliance, and their distinct operational needs. This journey towards IT and ICS/OT convergence or separation (or somewhere in between) may involve the adoption of emerging technologies such as Internet of Things (IoT), Artificial Intelligence (AI), and data analytics to improve overall performance and efficiency in their operations.

Regardless of the level of integration of ICS/OT with external networks and systems, the ICS/OT border is the first line of defence in combatting any threats that may seek to exploit the increased connectivity.

ICS/OT environments are facing an increasing and evolving cyber threat landscape, moving beyond simple data breaches to potentially disrupt critical infrastructure and safety. These threats are driven by the convergence of IT and OT systems, the increasing sophistication of attacks, and the diverse motivations of threat actors.

This guidance article has been derived from the real-world experience of experts in the field of ICS/OT cyber security and networking. It aims to provide comprehensive guidance in understanding, designing, maintaining and enhancing the overall security posture of ICS/OT borders, and thus the assets that reside behind them.

This article is intended for IT and ICS/OT network engineers, cyber security teams, and anyone seeking guidance on how to securely implement new ICS/OT network borders, or secure existing ICS/OT network borders.

Other related articles published by the ICS COI can be found here https://ritics.org/ics-coi/#downloads and cover the likes of the use of removable media, and ICS/OT Field devices.

This document provides ICS/OT-specific security principles to assist organisations in understanding, designing, maintaining, and enhancing secure ICS/OT borders. It is intentionally non-prescriptive, allowing organisations to tailor their approach based on their specific risk appetite and level of convergence across their borders.

Scope

The guidance specifically addresses IP-based ICS/OT networks and the interface between the ICS/OT network and any other network, referred to as the boundary.

In the context of this guidance, an ICS/OT boundary is either:

- An ICS/OT Wide Area Network (WAN) interface to another network such as the Internet or a business IT network.
- An ICS/OT Local Area Network (LAN) interface to another network such as an ICS/OT WAN, another ICS/OT LAN, a business IT network, or the Internet.

ICS/OT WAN Definition

An ICS/OT WAN is regarded as a single network or a group of networks that:

- Interconnects multiple ICS/OT LANs.
- Is exclusively for internetworking of ICS/OT traffic.
- Has communication channels, dedicated by physical or logical means to communications between the networks.
- May comprise multiple WAN technologies, which may be from the same or different network providers.

An ICS/OT WAN may be comprised of multiple WAN technologies, which may be from the same or different network providers.

ICS/OT LAN Definition

An ICS/ OT LAN is regarded as a single network or a group of networks that are:

- Exclusively for networking of ICS/OT traffic.
- Located within the same physical site boundary.
- Logically protected under the same authority.
- Interconnected via communication channels comprising solely of equipment and infrastructure owned by the organisation and dedicated to communications between the networks.

Although a LAN is confined to a single physical site, a single site could have multiple LANs.

Security governing the interconnectivity of individual subnetworks, such as a Virtual LAN (VLAN) comprising an ICS/OT LAN is not within the scope of this guidance. However, many of the measures used for border security can be applied to securing the interfaces inter-connecting subnets.

Alignment to Cyber Security Frameworks and Regulations

When designing and implementing ICS/OT border security, it is important to consider the requirements of relevant frameworks and regulations. This will help to ensure that your organisation is meeting its compliance obligations whilst reducing its risk of cyberattacks.

Organisations in scope of the Network and Information Systems Regulations 2018 (NIS) will likely be mandated by their competent authority to meet the security outcomes of the NCSC Cyber Assessment Framework (CAF). Throughout this document, references can be found to relevant CAF indicators of good practice (IGP)s.

Requirements of the international security standard ISA/IEC 62443-3-3 may also need to be considered for your ICS/OT border. Specifically, Foundational Requirement (FR) 5.

Structure of this guidance

This guidance article primarily covers the concepts that should be considered when designing ICS/OT borders, and is split primarily into the 3 following areas:

1. **Routing Basics** – overview of how data can be passed between IP networks to serve as a foundation for understanding ICS/OT border security.
2. **ICS/OT Border Security Controls** – technologies that can be operated at a border to control their use, including solutions for allowing users to securely pass through ICS/OT borders whilst preventing ingress by threats.
3. **ICS/OT Border Architectures** – primarily based upon the concept of De-Militarised Zones (DMZs) showing how a border infrastructure can be arranged to control the level of segregation between networks on either side of the border.

These sections also include “best practice” for the application of the various concepts covered based upon the experience of the contributors to this document.

Finally, there is a summary of guidance providing a concise view of the most important points covered that should be focussed upon when designing an ICS/OT border.

Why Secure Borders?

In many organisations, it was once commonplace to operate ICS/OT across local proprietary networks, with any external connectivity being limited to serial telecommunications. These telecommunications provided two main capabilities:

- **Passing Process Data and Control Signals:** Enabled the transmission of process data and control signals between two geographically dispersed locations, allowing them to operate under a common control scheme.
- **Centralised Data Retrieval:** Allowed for the retrieval of process data back to a centralised telemetry system. In some cases, the telemetry system also offered rudimentary control capabilities.

Whilst serial telecommunications were basic compared to modern day standards, their "point-to-point" nature provided limited exposure to cyber threats from outside of the local network.

Organisations are now increasingly moving towards more modern networking capabilities using IP-based communications; both internally within the local network and externally to other networks. This shift is driving a surge in the use of cloud-based ICS/OT solutions, increased third-party interactions, and greater ICS/OT integration with IT.

Maintaining the integrity and reliability of ICS/OT systems is essential for the safe and efficient operation of industrial processes. Interference or disruptions caused by unauthorised access or changes can lead to operational failures, production losses, and safety hazards. Therefore, it is imperative to establish robust controls to combat the increased threat exposure associated with rising ICS/OT connectivity.

Robust ICS/OT border security may also be required to meet compliance obligations; especially critical nation infrastructure (CNI) organisations like power generation, water treatment, and transportation, who are subject to strict regulations and standards.

Routing Basics

This guidance article focusses on ICS/OT borders that interconnect IP networks. This section covers the basics of routing between IP networks to serve as a foundation for understanding ICS/OT border security. It is not essential to read this section to understand the remainder of the document, however, these concepts are important to understand to enable an ICS/OT border to be securely designed and operated.

The Internet Protocol Suite

The internet protocol (IP) suite is a set of protocols used across the Internet as well as most modern network types including Ethernet and Wi-Fi.

This document focuses primarily on the Internet Protocol, but also references two other important protocols in the application of network borders:

- **Transmission Control Protocol (TCP)** – A connection-orientated protocol, requiring the sender to establish a connection with the receiver before any data can be transmitted. The connection provides a bi-directional channel over which the receiver notifies the sender when it has successfully received a datagram – if the sender does not receive a successful notification the datagram is resent.
- **User Datagram Protocol (UDP)** – A connectionless protocol. Data is sent without first needing to establish a connection. UDP provides no ability to detect if datagrams have been successfully received. UDP is, therefore, uni-directional.

TCP and UDP are both transport layer protocols, often referred to as layer 4 protocols due to their position within the Open Systems Interconnection (OSI) model. They encapsulate application data for transport across a network.

IP is a network layer (layer 3) protocol. It provides the routing capabilities to pass datagrams between network hosts.

Network Router

Hosts on an IP network cannot communicate directly with other hosts outside the local network and must route traffic via a network router. This router is a fundamental component of a network border.

Network routers have multiple interfaces, each connecting to a different network. Each interface can support various network technologies, bridging different network types, such as internal Ethernet networks and external internet-based technologies.

A network host must be configured to route via the network router, which is commonly achieved by configuring the hosts default gateway to the IP address of the network router's interface to the internal network.

Best Practice: Limit border routers to one per network to simplify border security and routing. Dual network routers can be deployed as a high-availability pair to enhance network border resilience.

Routing and Inbound Connectivity

Once the network router is established, a mechanism for inbound connectivity to the ICS/OT network from outside is required. Even if inbound connections to the ICS/OT network are not desired, outbound connections using TCP within IP rely on bidirectional communications.

The primary methods for achieving inbound connectivity include:

- Inbound routes from the external network to the ICS/OT network
- Network Address Translation (NAT)
- Port Forwarding
- Virtual Private Network (VPN)

Inbound Routes

Inbound routes direct data destined for the ICS/OT network from the external network to the ICS/OT network border router. This can be configured in various ways, which are beyond the scope of this guidance but generally involve the external network operator setting up routes to the ICS/OT network.

Network Address Translation

There may be circumstances where it is not possible to configure inbound routes to the ICS/OT network. One example of this is where the ICS/OT border is connected to the internet.

Internet borders will use a public IP address allocated by an Internet Service Provider (ISP). Public addresses are routable across the internet. Any devices within an organisations internal network would typically use private IP addresses which can be managed and allocated independently by the organisation. Private IP addresses are not routable across the Internet.

Another common problem with creating inbound routes to an ICS/OT network is where multiple networks use the same IP address ranges. This can come about when ICS/OT networks have been deployed as isolated environments and using common IP address ranges simplified network management. However, this presents a problem when integrating networks together. Where inbound routes are not possible, NAT can be used, either:

- dynamically on its own.
- dynamically with port forwarding.
- or statically.

Dynamic (Hide) NAT - Dynamic NAT, also known as hide NAT, replaces (or hides) the source IP address in outbound packets, with the IP address of the network router. Packets are then forwarded out to the external network. The network router keeps a record of all translated packets, so that the returned packets received by the network router can have their destination address replaced with the private IP address of the original device, before forwarding back into the internal network for onward routing.

There is a security benefit that comes with dynamic NAT; initiation of inbound communications is not possible. Ergo, threats outside of the ICS/OT network cannot enter directly in via the network border.

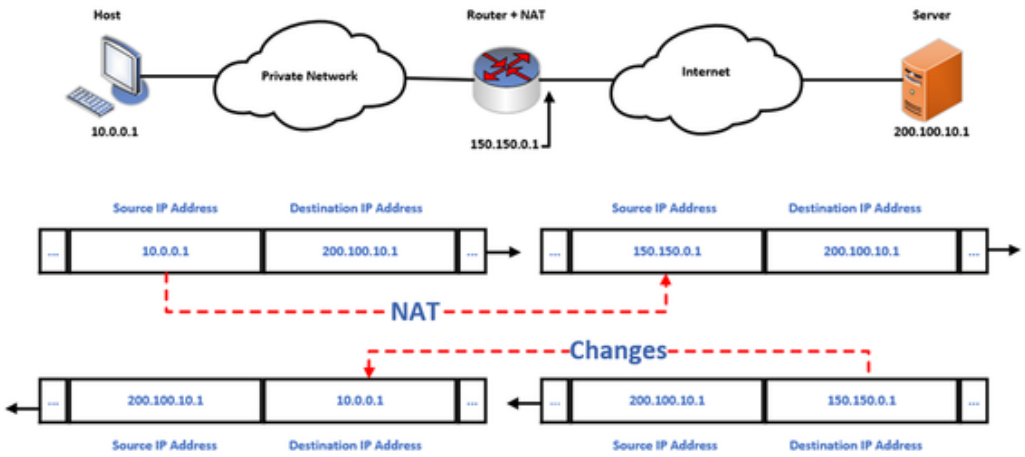

Figure 1 - Network Address Translation

However, this does not mean the connection is secure. Internet access and email can be used as initial access vectors by threats and, therefore, should be blocked. Threats such as malware may also reach the ICS/OT network through hardware additions like portable computers or USB drives. The malware may be reliant upon connecting out to an external command & control (C2) server to cause any harm to the ICS/OT network, so blocking internet access can neutralise this threat.

Port Forwarding - If a connection needs to be established inbound to a device on the internal ICS/OT network, this cannot be done by using the internal devices IP address because it is not routable from the external network. To overcome this, a technique known as port forwarding can be used, whereby the network router is configured to forward packets received for a specific TCP or UDP port to an internal IP address on the ICS/OT network. From the external network, packets are addressed to the network routers public address, and if the port matches one configured for port forwarding, the destination IP address in the IP packet is replaced with the IP address of the device on the internal ICS/OT network and forwarded on.

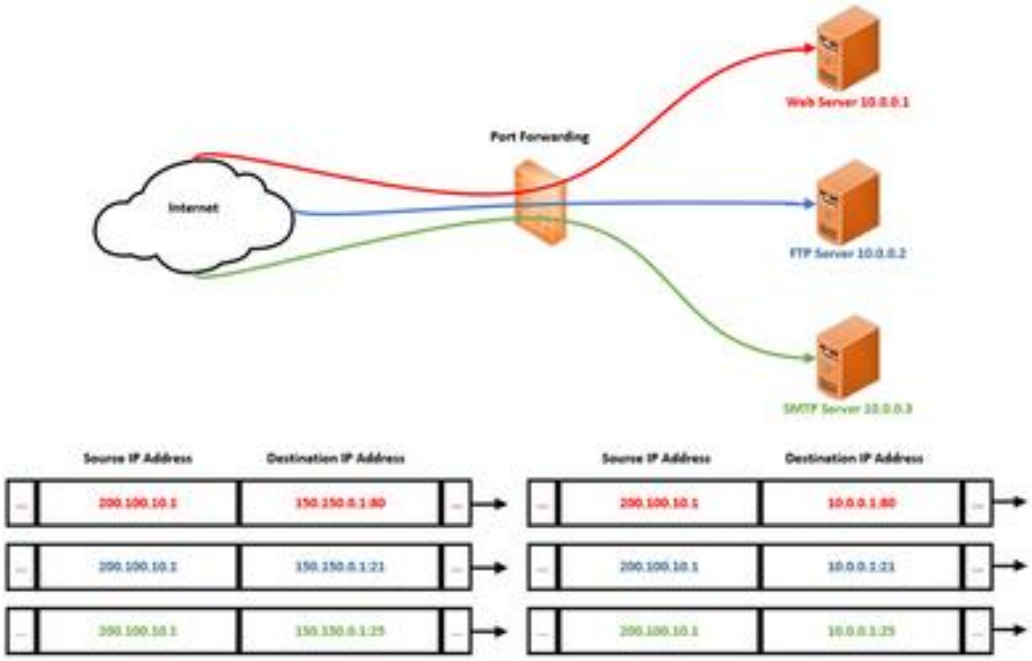

Figure 2 - Port Forwarding

Whilst we have allowed inbound connectivity, the opportunity for exploit by a threat is still limited to the specific TCP/UDP port on a specific host, making compromise more difficult than a fully routed connection to the entire network.

Static NAT - One issue with port forwarding, is that only one internal ICS/OT device can be forwarded packets for any one IP address and network port. E.g. if there are two web servers on the internal network using TCP port 443, only one of those could be reached via port forwarding from the external network via the network routers external IP address. If there are two or more devices on the internal network that need to be reached using the same service port, then static NAT can be used, whereby each internal ICS/OT network IP address is mapped to a different external address assigned to the network router.

Best practice: Carefully consider the use of static NAT and PAT due to the exposure it creates to external threats. Devices reachable by this method would need stringent security controls in operation, as a minimum:

- The latest security updates
- Up-to-date anti-virus/endpoint protection for computers with commercial operating systems such and Windows or Linux.
- Strong authentication with brute-force password protection. Ideally using multi-factor authentication.

Virtual Private Network (VPN)

VPN tunnels can also be used to interconnect private networks across the public internet. Once the VPN is established, hosts on the private networks at each end of the tunnel can freely communicate with one another as if the networks are directly connected.

There are two common modes for VPNs:

- **Site-to-Site VPN** – in this mode a VPN tunnel is established between VPN gateways on different sites. This allows devices on the private networks of each site to communicate between one another.
- **Remote Access VPN** – a VPN can be established directly from an individual device to a VPN gateway.

A VPN gateway can be a standalone device or integrated with a network router. VPNs are also a valuable tool for securing data-in-transit. See the VPN Security section.

ICS/OT Border Security Controls

The introduction of an ICS/OT border via a network router presents potential for unrestricted network communications in-and-out of ICS/OT networks.

In this section the most common security controls for governing ICS/OT border traffic are examined. NCSC have also issued guidance on securing network perimeters that can be found at https://www.ncsc.gov.uk/guidance/network-security-fundamentals#section_6

VPN Security

VPNs not only allow routing between private networks across public networks; they can also secure data in transit to provide comparable integrity and confidentiality protection to a physically secure point-to-point link. This also makes them very useful for protecting data over shared/untrusted networks – be they public or private networks. The main benefits of a VPN are:

- Data is encrypted in transit providing confidentiality, ensuring that the data has not been observed by an unauthorised party.
- Any manipulation of the data during transit should be detected through the integrity checking provided by the VPN tunnel, which will then be discarded.

- Unlike with static NAT and PAT, data cannot simply be sent to host behind the network border. A VPN must first be established, which is only possible if the key for the VPN is known.

When each device is issued with a unique VPN key, and the key is kept secret, then message origin authentication is also achieved. This provides assurance that the data came from the expected entity. I.e. data has come from an authorised source.

Best Practice: Where VPNs are used and there is no requirement to transmit data outside of the VPN, all other mechanisms for routing should be disabled, including inbound routing, NAT and port forwarding.

Further detailed guidance on VPNs can be found at: Virtual Private Networks (VPNs) - NCSC.GOV.UK

Network Firewalls

A firewall analyses inbound and outbound IP traffic traversing the ICS/OT border, comparing it against its configured firewall rules, to determine if the traffic should be accepted or blocked. Each rule typically defines:

- The source and destination IP addresses, which may be network or host addresses.
- The protocol type. E.g. Transmission Control Protocol (TCP) or User Datagram Protocol (UDP)
- The destination port numbers.

Firewalls are also able to analyse the state of each connection, so any packets that are out-of-state, which may be used to defeat the firewall or attack endpoints protected by the firewall, can be dropped.

Network firewalls often have inbuilt network routing capabilities. However, they may also be combined with dedicated network routers where advanced routing capabilities are needed or where connection to access circuits not supported by the network firewall is required. E.g. Digital Subscriber Line (DSL).

Best Practice: Firewall rules should be configured to only allow specific permitted IP addresses for specific permitted ports and protocols. This must be observed for both inbound and outbound connections.

Back-to-back Firewalls is also a consideration, although NCSC has provided some commentary on their use that can be found at https://www.ncsc.gov.uk/whitepaper/security-architecture-anti-patterns#section_5, with an acceptable use case for back-to-back firewalls being where you might have separate teams responsible for the ICS/OT environment and the IT/Internet environment, and therefore one team configures one firewall and the other configures the second firewall.

Network Intrusion Detection and Prevention System (NIDPS)

Traditional firewalls only inspect traffic up to the transport layer (E.g. TCP/UDP) of the IP suite. Firewalls are therefore unable to provide any application layer network protection. To address this gap, we can look to NIDPS.

NIDPS analyses traffic, usually that has been allowed by the firewall, for any potentially malicious communications. Analysis operates, primarily, by three methods:

- **Signature-based** – relies on indicators of compromise (IoC) for known threats by monitoring network traffic for specific threat characteristics, including known byte sequences, file hashes, TLS certificate information, or malicious domains. Signatures must be regularly updated to identify new threats.
- **Anomaly-based** – normal network behaviour is baselined by the system, allowing abnormal or unusual patterns of network traffic to be identified.
- **Policy-based** – bespoke rules specific can be created to identify threats relevant to the specific environment.

As the name implies, NIDPS can be configured to either detect or prevent threats. The action depends on the configuration of the specific analytic. In detection mode, an alert is generated if suspicious traffic is identified, requiring triage and investigation to take any affirmative action.

NIDPS is often combined with firewall technology, as found in Next Generation Firewalls, providing a single product to alert and filter traffic. Systems that only operate in detect mode (NIDS) are also commonly used and, although they can be deployed as dedicated devices, they are now commonly found in Network Discovery and Response (NDR) tools, which can operate completely passively by ingesting mirrored data from the network to detect traffic anomalies and signs of malicious activity.

Best Practice: In prevention mode, suspicious traffic is blocked. This mode should be used cautiously to avoid blocking legitimate traffic, particularly for signature or anomaly analysis. However, policy-based analysis can be used to great effect with an ICS/OT protocol aware NIDPS if applied correctly. For example, to allow data to be read from a Programmable Logic Controller (PLC) but not written.

Unidirectional gateways

Unidirectional gateways, also known as data diodes, only allow communications in one direction. This is achieved by the physical structure of the gateway that only has circuitry to allow data to be sent in one direction. Unidirectional gateways for UDP are simple devices that can effectively be provided by standard gateway technology, just with circuitry removed such that data can only flow in one direction. For TCP, which requires a bi-directional network path, a proxy capability must be incorporated into the unidirectional gateway which can terminate the TCP connection with the sender and re-establish with the receiver.

See The Internet Protocol Suite section for details on UDP and TCP.

Use Cases

Unidirectional gateways can be operated at an ICS/OT border to only allow either inbound or outbound communications.

Outbound: In this configuration, data can be sent out of the network, eliminating any inbound paths for network threats. Typical examples include:

- To pass security event data to a Security Information and Event Management (SIEM) platform for monitoring by a Security Operations Centre (SOC)
- Sending ICS/OT backups off site to they are protected from an extreme event.

Inbound: When there is only a need for a border to provide inbound communications, such as importing software updates or backups, a data diode provides the ability to do this without opening up any outbound connectivity. As previously discussed, outbound communications can be exploited by threats such as malware to connect out to a C2 server.

Reversable Unidirectional Gateways

These are useful for when both outbound and inbound communications are needed, but not simultaneously. Combining the previous examples, a reversable unidirectional gateway can be used as follows:

- Under normal operation, the gateway allows outbound communications for passing security event data and backups.
- When there is a need to import data, the gateway direction is reversed. When the import is complete, the gateway is returned to the outbound direction.

There is a trade-off here which must be considered. While this approach increases threat protection, there would need to be an acceptance of potential gaps in security event data received into SIEM during the import period.

Best Practice: Use unidirectional gateways where a network needs to achieve the security protection from inbound threats of an air-gapped network, whilst still allowing the centralised monitoring of security events.

CAF IGPs:

- **CAF - B5.c – Achieved IGP** "Your comprehensive, automatic and tested technical and procedural backups are secured at centrally accessible or secondary sites to recover from an extreme event."
- **CAF – C1.a – Achieved IGP** "You have extensive monitoring coverage that includes host-based monitoring and network gateways."

Zero Trust Network Access (ZTNA)

A disadvantage from a security perspective of traditional VPNs is that they typically give users access all resources accessible via the VPN. Whilst a firewall can be used to limit what resources are accessible, they are not limited to individual users. This can unnecessarily expose resources.

Building upon the principles of zero trust architectures, ZTNA is technology that provides secure access to supported applications, such as:

- Web Services
- Secure Shell (SSH)

- Remote Desktop Protocol (RDP)
- Databases

Access to each instance of an application can be assigned to individual users or groups of users.

ZTNA can also be used to gain assurance over the security posture of the remote access device. This can be helping when the connection is being used to perform privileged actions requiring a 'browse-down' approach to be achieved, whereby the device being used to perform administration is of a higher-trust than the system being administered.

Some of the security aspects that can be validated to gain assurance over the device include:

- Up to date operating system and anti-virus protection
- Host firewall
- Drive encryption
- Screen lock
- https://www.ncsc.gov.uk/whitepaper/security-architecture-anti-patterns

At the heart of ZTNA is a trust broker, which is comprised of:

- **A control plane** – this hosts the Policy Decision Points (PDP) where the policies are defined, establishing the conditions that must be met by users and devices before access is granted.
- **A data plane** – this hosts the Policy Enforcement Points (PEP) restricting access only to applications allowed by the policy.

These can both be provided by the same resource or by separate resources. The resources can be hosted in the cloud or on-prem.

The process from the user's perspective is as follows:

- Users initiate access by establishment of a secure web connection to a trust broker. This can be hosted either on prem or in a cloud.
- The user provides the required authentication data.
- Optionally, a device posture check can be enforced to validate the cyber hygiene of the user's device.

- After successful user authentication and device validation the applications the individual is specifically authorised to access and presented to the user.
- The user selects the desired application, and they are routed to it via the trust broker.

Best Practice: Use zero trust network access to provide more granular control over access to resources.

CAF IGPs

- **CAF – B2.a – Achieved IGP** "The number of authorised users and systems that have access to all your network and information systems supporting the essential function(s) is limited to the minimum necessary."

Bastion hosts

Bastion hosts are used at ICS/OT borders to broker connectivity between internal and external networks. Bastion hosts can facilitate more secure border configurations by negating the need for direct connectivity to the ICS/OT network.

Common use cases for Bastion hosts are:

- **Jump boxes (aka terminal servers)** to give remote users access to systems and data without providing a direct connection from the user's computer. When users connect to the bastion host, they are presented with a virtual desktop installed with the applications required to interface with the system.
- **Data import –** Data may need to be brought into an ICS/OT network from external systems. Bastion hosts can be used as the initial landing platform for imported data and security screened for threats. The bastion host can also be used to transform complex file types like Excel, Word or PDF which can be embedded with malicious into simple file types such as CSV of TXT before onward transmission to the internal network.
- **Data Export –** Providing a mechanism for exporting data increases the risk of unauthorised data egress. Within ICS/OT, data integrity and availability are likely to be more important to protect than confidentiality. However, information about the system, such as PLC configuration, could be extremely useful to an attacker. Bastion hosts can be used as export control servers to allow an authorising party to control data exports before onward forwarding to the intended recipient.

The Bastion host also provides the benefit of creating connection breaks and protocol change:

- **Connection break** – the lack of direct TCP/IP connectivity between external systems and the ICS/OT LAN reduce the ability for threats to traverse directly to the ICS/OT LAN.
- **Protocol change –** the service used to access the bastion host, say RDP, would be different than the protocols used to access a host on the trusted network, e.g., a back-end Supervisory Control and Data Acquisition (SCADA) database. Because protocols used by the trusted network do not need to be exposed remotely, this provides a degree of protection over any vulnerabilities or security weaknesses in those protocols. This is particularly useful when needing to communicate to end devices via ICS/OT protocols which often lack message authentication capabilities.

Bastion hosts should be configured with more stringent security protection than devices within the internal network due to their increased exposure to external threats. In particular, if using a Windows operating system on the bastion host, drive and clipboard sharing should be blocked. This prevents data being copied over the connection from the client to the server, or vice-versa.

Bastion hosts that provide protocol breaks, including encryption breaks, are excellent points for implementing logging and monitoring. NCSC provides some commentary on the use of Bastion hosts that can be found at https://www.ncsc.gov.uk/whitepaper/security-architecture-anti-patterns#section_3

Best Practice: Use bastion hosts to broker connections in and out of the ICS/OT network to provide greater control over data connections and reduce exposure to threats.

CAF IGPs:

CAF – B4.a – Achieved IGP "Content-based attacks are mitigated for all inputs to network and information systems that affect the essential function(s) (e.g. via transformation and inspection)."

Privileged Access Workstation (PAW)

Privileged access workstations are devices dedicated for management of the trusted infrastructure. They should have no access to browse the Internet, access email, or connect to the organisation's business systems. Where PAWs are used, all privileged access should be via a PAW, including from third parties. NCSC have issued guidance on the development and use of PAWS that can be found at https://www.ncsc.gov.uk/collection/principles-for-secure-paws

PAWs must provide a 'browse-down' approach and, therefore, achieve a comparable or higher level of security protection and trust than the systems they are managing. The level of security will be dependent upon the type of workstation, but broadly they fall into one of the following categories:

- **PAWs that only connect to the trusted infrastructure**. These may only require a level of security protection and trust of a comparable level to the trusted infrastructure. Although, if possible, greater security protection should be applied if feasible to do so.
- **PAWs used for accessing the trusted infrastructure from external environments**. the environment the PAW is being operated from is likely to be less secure than the one being accessed, greater security controls are likely to be needed to achieve a comparable or higher level of trust.

It will not always be feasible to enforce the use of PAWs. Where access needs to be permitted from other types of workstations, then minimum security requirements should be mandated to minimise the risk of these devices coming under the control of a threat. One method of enforcing this is via the compliance check capabilities often found in ZTNA products. This checks the device for things like the presence of up-to-date antivirus (AV) software, or that the operating system (OS) is fully patched with security updates. Whilst not a substitute for PAWs, this is preferable than allowing access from devices with an unknown security posture.

It should be noted that bastion hosts cannot be considered a substitute for PAWs. The device used to input user credentials must be of a higher trust than the system being accessed, otherwise it cannot be guaranteed to have a higher level of security protection, which increases the chance of the device and the user credentials being compromised.

Best Practice: Use PAWs coupled with bastion hosts to broker communications at the network border.

CAF IGPs:

- **CAF - B2.b – Achieved IGP** "All privileged operations are performed from highly trusted devices, such as Privileged Access Workstations, dedicated solely to those operations."
- **CAF - B2.b – Achieved** IGP "You either obtain independent and professional assurance of the security of third-party devices or networks before they connect to your systems, or you only allow third-party devices or networks dedicated to supporting your systems to connect."

Temporary / Timebound Access

Keeping remote access permanently enabled increases the opportunity for it to be exploited by a threat as an access vector to the ICS/OT network. To minimise this risk, remote access can be established on a temporary of timebound basis:

- Temporary access is useful for ad hoc requirements where access is enabled to perform a particular activity, which is revoked when the activity is complete.
- Timebound access is used for regular activities that happen periodically at specific times. The access is configured such that access can only be achieved during the pre-defined time periods.

Temporary and timebound access can be achieved via:

- **Remote access accounts** – Each account can be kept disabled and only enabled when required. Some remote access services allow time schedules to be created for users so they can only access systems at specific times.
- **Firewall rules** – Firewall rules permitting connectivity can also be kept disabled and enabled when a connection is required. Firewalls often support scheduling so that firewalls only allow connections at specific times.
- **Switch port** – Keeping switch ports disabled stops all network connectivity through the interface. The interface used would need careful consideration, so that the remote access capability can be disabled without impact other communication links that may not be suited to temporary/timebound access.
- **Cloud-based bastion host** – Where a bastion host is operated as a cloud virtual machine (VM), a fresh bastion host can be provisioned for each new session from a

preconfigured bastion host VM image. Once the remote session is complete, the VM is archived to provide auditing of remote activity to reconstruct events if required.

NCSC provides additional Privileged Access Management guidance in its principle based article - https://www.ncsc.gov.uk/collection/secure-system-administration/use-privileged-access-management

Best Practice: Keep all remote access disabled and only enable when specifically required.

CAF IGPs:

- **CAF - B2.c – Achieved IGP** "The issuing of temporary, time-bound rights for privileged user access and / or external third-party support access is in place."

OT Border Architectures

Introducing multiple network layers at the network border can further enhance the protection offered. Connections can be brokered via intermediary devices hosted within these network layers, limiting direct pathways between external networks and critical internal assets. In some arrangements, direct routing paths can be eliminated so that data control is not reliant purely on firewall rules that are prone to configuration errors.

The network layers operated at a network border providing additional segregation between internal and external networks are commonly known as a de-militarised zone (DMZ).

Dual Network Firewall DMZ

Two layers of network firewalls can provide a more secure architecture by limiting data routing directly between the internal and external networks.

Typically, there are two methods for deploying two layers of network firewalls, shown in figures 3 & 4 below.

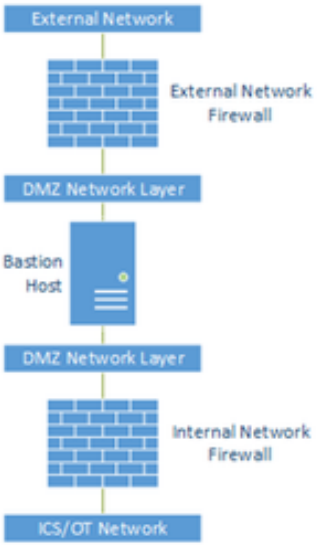

Figure 3 - Dual homed Bastion host DMZ

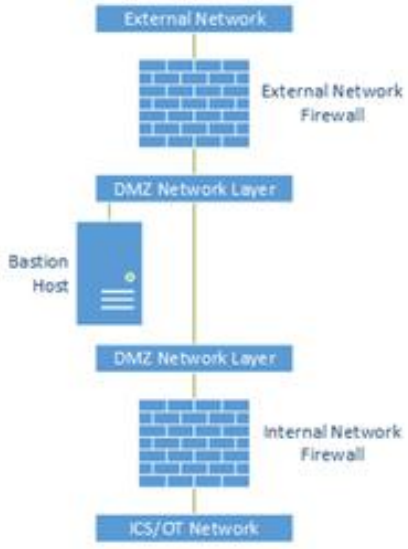

Figure 4 - Single homed Bastion Host DMZ

Where dual-homed bastion hosts are used there is no physical path between the external and internal networks at the network border and all data must be brokered via the bastion host.

Even with the single-homed bastion host, it won't be possible for data to be routed from external networks to the internal network, unless the external gateway is configured with specific routes for the internal network.

There may be use cases where there is a need to pass data between internal and external networks directly without brokering on bastion hosts. Two firewall layers can still be beneficial to provide governance separation over the management of the firewall layers, by:

- Preventing a single party being able to open unauthorised connectivity across the border, either due to accidental errors or malicious intent.
- Requiring an external threat to gain access to the management interface of two separate firewall layers to breach the border.

Single Network Firewall DMZ

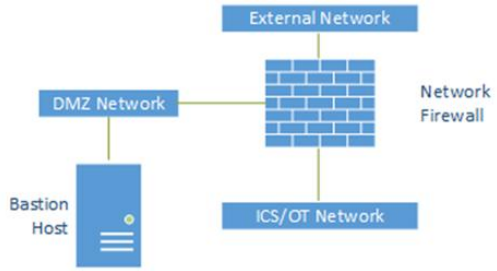

Figure 5 - Single Network Firewall DMZ

A DMZ can be established with single security gateway by provisioning it on a dedicated network interface. Firewall rules are relied upon to limit direct connectivity between the internal and extern all networks. Whilst a much simpler architecture routing separation cannot be achieved, and segregation is solely dependent upon a single set of firewall rules.

Segregating Multiple Networks

The three border architectures that have been depicted provide the basic building blocks upon which network segregation can be achieved. These can be adapted as needed to meet specific applications. For example:

- Multiple DMZ VLANs can be used allowing separation of multiple bastion hosts. Firewall rules are used to block connectivity between bastion hosts.
- Multiple ICS/OT VLANs can be created within the internal network. Again, using firewall rules to limit connectivity between subnets. This can be used, for example, to segregate networks operating at different levels of the Purdue Model.

The figure below shows an example of this:

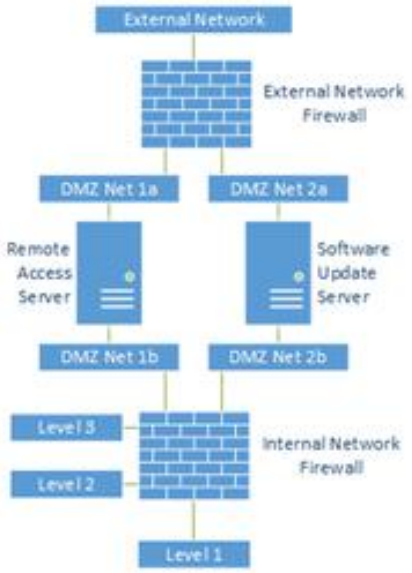

Figure 6 - Multiple VLANs example.

CAF IGPs:

- **CAF – B5.b – Achieved IGP** "Network and information systems that support the operation of the essential function(s) are segregated from other business and external systems by appropriate technical and physical means (e.g. separate network and system infrastructure with independent user administration). Internet services are not accessible from network and information systems."

Putting the Guidance into Practice

This paper has discussed a wide range of topics related to ICS/OT border security. There is no one single approach that fits all; each organisation must determine how to achieve the necessary level of connectivity while protecting against potential security threats. The following are key areas that should be considered when making those decisions:

- Always use a network firewall at the ICS/OT border to restrict connections to the minimum necessary and only from specific authorised sources.
- Use network firewall per network to simplify border security and routing. Dual network firewalls can be deployed as a combined high-availability pair to create a more resilient network border.
- On network firewalls that connect to the Internet, disable NAT if there are no devices on the ICS/OT network that need to connect to hosts on the internet with public IP addresses.
- If inbound access to an ICS/OT network is required from a non-routable network use VPNs rather than static NAT or PAT.
- Consider using VPNs for all communications outside of the ICS/OT network to restrict connections to authorised sources and to protect data-in-transit.
- NIDPS provides a deeper level of traffic inspection than firewalls which can help to detect and/or block threat activity.
- For remote user access, consider using ZTNA instead of VPNs to provide greater control over what each user can access.
- Only allow remote user access from PAWs or trusted third-party devices with validated security controls. Use ZTNA to assess the security posture of these devices.
- Keep all remote access disabled and only enable when specifically required.
- Use bastion hosts to broker connections in and out of the ICS/OT network to provide greater control over data connections and reduce exposure to threats.
- Locate bastion hosts in a DMZ network.
- Use unidirectional gateways where a network needs to achieve the security protection from inbound threats of an air-gapped network, whilst still allowing the centralised monitoring of security events.

CAF Indicators of Good Practice Summary

This case study discusses measures that contribute to the following CAF V3.1 Indicators of Good Practice (IGP)s:

- **CAF – B2.a – Achieved IGP** "The number of authorised users and systems that have access to all your network and information systems supporting the essential function(s) is limited to the minimum necessary."
- **CAF - B2.b – Achieved IGP** "All privileged operations are performed from highly trusted devices, such as Privileged Access Workstations, dedicated solely to those operations."
- **CAF - B2.b – Achieved** IGP "You either obtain independent and professional assurance of the security of third-party devices or networks before they connect to your systems, or you only allow third-party devices or networks dedicated to supporting your systems to connect."
- **CAF - B2.c – Achieved IGP** "The issuing of temporary, time-bound rights for privileged user access and / or external third-party support access is in place."
- **CAF – B4.a – Achieved IGP** "Content-based attacks are mitigated for all inputs to network and information systems that affect the essential function(s) (e.g. via transformation and inspection)."
- **CAF – B5.b – Achieved IGP** "Network and information systems that support the operation of the essential function(s) are segregated from other business and external systems by appropriate technical and physical means (e.g. separate network and system infrastructure with independent user administration). Internet services are not accessible from network and information systems."
- **CAF - B5.c – Achieved IGP** "Your comprehensive, automatic and tested technical and procedural backups are secured at centrally accessible or secondary sites to recover from an extreme event."
- **CAF – C1.a – Achieved IGP** "You have extensive monitoring coverage that includes host-based monitoring and network gateways."

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, ICS-COI members have used all reasonable care in verifying the guidance contained within using the data sources available to it, they provide no warranty as to its accuracy or completeness. To the fullest extent permitted by law, the ICS-COI accept no liability whatsoever for any expense, liability, loss, damage, claim or proceedings incurred or arising as a result of any error or omission in the report or arising from any person acting, refraining from acting, relying upon or otherwise using this document. You should make your own judgment as regards use of this document and seek independent professional advice on your particular circumstances.

Reference to any specific commercial product, process or service by trade name, trademark, manufacturer or otherwise, does not constitute or imply its endorsement, recommendation or favour by the ICS-COI. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.