



INDUSTRIAL CONTROL SYSTEMS

Community of Interest

GUIDANCE: RESOLVING ANTI-PATTERNS IN INDUSTRIAL CONTROL SYSTEM / OPERATIONAL TECHNOLOGY ENVIRONMENTS.

Introduction

Organisations such as [RITICS](#) have researched implementations of the [NIS Directive](#) and the [Cyber Assessment Framework \(CAF\)](#) across several critical sectors including transport, water, and energy.

During the NCSC's reviews of these implementations, they have identified some patterns often seen in system designs in CNI organisations that you should **avoid**. The term 'anti-pattern' has been developed to refer to these repeated (but ineffective) solutions to a common problems.

This guidance is for operators, consultants, suppliers and regulators working with organisations who own or operate Industrial Control Systems (ICS) / Operational Technology (OT) to assist in identifying and resolving anti-patterns found in these environments. The guidance:

- unpicks the thinking behind the anti-patterns
- explains why the patterns are bad and
- proposes better alternatives

It supplements NCSC's existing [anti-pattern advice](#) by examining anti-patterns commonly seen in the ICS/OT environment, although it will also help cyber security professionals working in traditional enterprise IT environments.

In this guidance

- [Anti-pattern 1: Flat, unsegmented/unsegregated networks](#)
 - [Anti-pattern 2: Uncontrolled access to ICS/OT networks](#)
 - [Anti-pattern 3: Lack of authentication and data security](#)
 - [Anti-pattern 4: Inaccurate asset inventory](#)
 - [Anti-pattern 5: Unchecked backups](#)
-

Anti-pattern 1: Flat, unsegmented/unsegregated architectures

Flat, unsegmented/unsegregated networks are characterised by devices and hosts being able to communicate across to other devices and hosts on a network unhindered and where they have no legitimate need to do so. Flat unsegmented networks are commonly built using a switch (or several switches) to connect all the devices on the network, without VLAN technology or routers to enforce segregation. The same effect can be seen where firewalls are being used without restrictive rules. Thus all hosts are routable to all other hosts




Figure 1 Flat, unsegregated architecture

What's wrong with this pattern?

This anti-pattern potentially exposes critical assets to unauthorised access and possible compromise. In many cases, the flat structure means there are numerous connections with corporate IT infrastructures, resulting in multiple points of entry. A compromised standard business system can give the attacker network access to the majority of critical ICS/OT systems.

Firewalls can help to restrict to both IT and ICS/OT networks by limiting traffic permitted across the interface enforcing segregation. However, if the firewalls simply restrict communications between two different IP addresses (that is, without also specifying the protocols allowed) then the system is still vulnerable to attack. The firewalls may also need to be configured so that the communication between resources can only flow one way.

An alternative to using firewalls is to divide the network up, segmentation, into subnets that are not directly reachable from each other, and use a service that straddles the two networks to provide gated access. Such services can include remote access servers (RAS), hardened jump boxes, bastion hosts and reverse proxies. However, such a design can be an [anti-pattern itself](#), as users can access high trust parts of the network from a low trust position through these services.

A better approach

A better approach is to introduce structure into flat and legacy architectures while preserving safety and reliability. This can be achieved by:

- Implementing a zoned architecture model (such as that described in ISA/[IEC62443](#)). For legacy installations, you can start by addressing the key assets within the architecture, and provide segmentation and segregation between them (and the rest of the network).
- Ensuring management interfaces of devices are only accessible to trusted management devices and not to networks more broadly.
- Use configuration management tools to formalise system deployments so that it is easier to track, update and re-deploy systems over time.
- Validating your control measures (for example by conducting tests where you try to reach segregated/segmented areas, or by commissioning a red teaming/adversary simulation event, with specific scope to focus on reaching ICS/OT assets).

For more detailed information please refer to the [NCSC's secure design principles within an OT environment guidance](#).

Purdue levels

A common approach for ICS network segmentation is to use the [Purdue enterprise reference architecture](#). This involves creating a network segment for corporate users with a lower security level than the more critical operations zone, as shown in Figure 2 below.




Figure 2 Purdue control hierarchy

While this segmentation offers some improvement over a flat network, any compromised internet-facing application could still provide a pivot point for access to internal network segments through exploiting a vulnerability or misconfiguration. The problem grows as more devices are connected. This can potentially lead to exposure of Purdue levels 0, 1, 2 where sensing and control equipment

is connected to physical infrastructure (and where there is most likely no authentication of devices).

To meet operational needs with minimal communication pathways, identify the data that needs to move between segments/zones. It may be important to keep functions self-sufficient within a segment/zone (for example, to minimise the need for real-time data passing between different trust levels). Define the trust relationship between hosts and between different zones/trust levels with authorisation required from both sides of the connection, and only the required content passing between different trust levels.

Implementing a segregated architecture can improve visibility of the applications, users, devices and content on the network, making it easier to detect anomalous or malicious activity. Knowing and establishing this internal communications structure will make it possible to isolate a compromised element or network segment more rapidly, reduce the impact of an intrusion, or limit the spread and impact of malware (including [ransomware](#)).

Identify and protect key assets

Rather than using a flat unsegregated/unsegmented network (or trying to compartment the whole network), decide what needs protecting, identify the critical and important assets, applications, services, and the protocols and control commands that are legitimate for an ICS/OT network. Understand the interdependencies and interactions that need to happen between assets. It may be important to segment critical assets from the rest of a process network. There could be non-process related components to consider (such as maintenance workstations, or network and protection monitoring).

As well as restrictions at the network layer, restrict needed communications down to port and protocol level to define how a host can communicate with other hosts. Tighter controls can also be introduced at the application layer to restrict the content of communications. The proposed segregation needs to be examined carefully in relation to common services (such as authentication requests that are passed between domain controllers) to ensure trust relationships are applied appropriately.

Understanding the risk

Segregation can be a complex exercise, so to be cost effective you may need to focus on high-risk functions and applications. By understanding the risks posed to your most critical assets (which can be done using [Crown Jewels Analysis](#)), you can then prioritise the implementation steps needed.

Some vendors can provide 'worst case' estimates of delays introduced by their firewalls. Where these delays would be harmful to the organisation, segregation/segmentation should be designed such that both parts of a time-critical process are placed in the same zone.

Segmentation/segregation can reduce the exposure of legacy systems to attack and assist security monitoring. Consider the compatibility of the segmentation/segregation solution with legacy equipment that has less capability to (for example) identify and authorise access and users.

Implementing network segregation

Begin with a baseline of user and device behaviour to define the segregations. From there, segregation is a dynamic and iterative process. For example, using the baseline data, implement VLANs but with unrestricted routing between them. By monitoring traffic content and flow, you can then develop the required controls and restrictions. Logging and monitoring can help to find improvements and maintain appropriate segregation.

'Normal' activity presents differently in separate segments/zones, having different functions and trust levels. Normal activity in one segment could even be malicious activity in another segment. Planning the expected communications allows anomalous traffic to be identified (and dealt with) more quickly. It makes it much harder to move around a network undetected to steal information, or seek access to more critical assets.

Recheck the segmentation/segregation if new threats or system changes necessitate a review of the network architecture; ongoing effort is required to maintain effective segmentation/segregation.

Limitations of network segregation/segmentation

In addition to segregation/segmentation, you should also consider other control measures using a defence-in-depth approach (such as access control, security monitoring and authentication where possible). Increase the effectiveness of the segregation allowing known and expected traffic rather than attempting to deny potential malicious traffic. Define the granular detail of *who* can access with *what* application, *when* access is needed and for *how long*.

Using multiple vendors across multiple domains introduces further complexity and requires consistency in deploying segregation policies. Preventative measures applied locally should also seek to limit connections with partner networks (or prevent such connections entirely) in order to prevent compromise from a partner leading to a breach of key systems.

Anti-pattern 2: Uncontrolled access to ICS/OT networks

As interconnectivity between ICS/OT assets and systems continues to grow, access control also grows, this can be done in an uncontrolled manner. For instance giving access to IT administrators to administer access control to the ICS/OT environment, without knowledge of how the IT administration works, can provide uncontrolled privileged access. The goal is for access to be controlled to ICS/OT assets so that only users who need to carry out an action on (or interface with) a system may do so. These controls are based around user authentication to enforce privileges and authority (but could also be physical, such as the use of keys to access a cabinet).

Given the operational lifetimes of many ICS/OT systems, access control may not initially have been a design concern. Where these systems have not been modified, implicit trust is granted to those working within the ICS/OT environment.

What's wrong with this pattern?

Where there is no controlled access to operational networks, either through physical or technical means, it is possible for an attacker to have unchallenged, unrestricted access to operational systems. Access control may not have been implemented due to implicit trust in the environment (and those working in it). This can result in unchallenged access and potential modification of these systems. Where access control had been implemented, it has typically provided no granularity (so users either have full or zero control of the system).

There are a number of ways in which this anti-pattern can manifest itself. Some are described below.

- **Unrestricted connections to ICS/OT networks.** Having direct, unrestricted access means engineers may carry out break-fixes or maintenance by plugging in third party systems into the same network as the ICS/OT systems. Without any effective monitoring to detect new devices entering the network, it is not possible to guarantee the secure and predictable state of assets in the ICS/OT environment, where arbitrary configuration changes (or manipulation) may occur without the asset owner knowing.
- **Engineering stations with no monitoring or controls.** Engineering stations on a production ICS/OT network typically have a higher degree of privilege over the system, in particular, having the ability to make configuration changes or control state within the ICS/OT system. Without granular controls on access and monitoring to ensure that all actions are logged, it is not possible to find out what changes may take place (and more importantly, if the change was carried out by a competent, authorised user). This is of broader concern where engineering station capability is also on devices on the corporate network, especially where generic PCs with internet access are used.
- **Uncontrolled use of removable storage media.** In many environments, there is an 'airgap' between systems, meaning software and data may be transferred using removable media. Unless the media is known to be trustworthy, it may contain malware which can alter the state, configuration and safe operation of the system.
- **Use of shared accounts.** Knowing exactly who is authorised to access a specific part of the system enables asset owners to be confident that the person making a change (or accessing a privileged part of a system) is competent and authorised to make that change. If a shared set of credentials is used within a system (which is very common with legacy ICS/OT elements), it is not possible to trace changes, and whether the user was authorised and competent.

A better approach

It is critical to ensure that access to the ICS/OT environment is authorised, controlled and challenged when appropriate. This ensures that the state of the system continues to be predictable and well-understood. Solutions and controls which should be followed include:

Network controls. Many network systems have provision to authenticate devices to a network, either through certificates, or simply by verifying the MAC address of the device. In some instances, switches can be configured to deny network access until the device has been authorised. This ensures that any 'new' devices are knowingly authorised to access the network. In some environments, a controlled gateway may also limit the number of entry points to specific ones provided for maintenance purposes. Simpler forms of network access include:

- tying used ports to known MAC addresses and disabling unused ports (or placing them in a 'disabled' VLAN)
- using physical port locks to deter the connection of devices to unused ports
- using captive connections to prevent devices from being disconnected

Control policies for removable media. Removable media is used for many different purposes within an ICS/OT Environment, including:

- installing software patches on ICS/OT systems
- importing anti-virus signatures onto an ICS/OT system
- installing new PLC/SCADA programs on ICS/OT systems
- exporting reports and other data
- provision of off-site backups

A key aspect of the design of any ICS/OT system is to understand all such scenarios where media is to be used, ensuring that all maintenance activities are considered, such that appropriate management of these interfaces can be implemented.

Removable media within an ICS/OT system can include USB drives, CD, DVD, SD Cards, floppy disks, and portable HDDs. You should ensure that:

- media is checked to confirm that it does not contain any known malware
- the integrity of all files being transferred into the ICS/OT environment is checked to confirm files only contain the expected, authorised, content
- a register is kept of all media used to import/export data, and that only authorised media is permitted for use within the ICS/OT environment
- records of all files/data transferred into, or out of, the ICS/OT environment are kept
- any sensitive data is appropriately secured/encrypted to ensure its confidentiality
- all media is appropriately sanitised either between use or when disposed, [as per NCSC's Secure sanitisation of storage media guidance](#)

Network monitoring. Logging and monitoring of network traffic within the ICS/OT network has many benefits, from defining a baseline to determining anomalous activity, to identifying changes to the network (such as newly-connected devices) and logging what data was exchanged between parties.

Trusted devices for maintenance. Where third party maintainers introduce their own devices into the ICS/OT environment, this can present new threats. For critical or commonly serviced systems,

running a fleet of maintenance devices ([as per NCSC's Privileged Access Workstation guidance](#)) with accounts allocated to third party maintainers) ensures that a known, good, configuration is maintained, and that trusted devices can be introduced.

Segregated accounts for access. Using individual, named accounts (with appropriate authority and permissions) ensures actions can be traced. These accounts should be configured using [NCSC's Secure system administration guidance](#). There should be an effective process for managing accounts, especially for staff joining/leaving/moving, and a clearly defined joint process of approval and validation by both the operator of the ICS/OT environment and the third party. Due to legacy ICS/OT equipment not being able to support individual named accounts, a 'defence in depth' approach may be required which could include:

- auditing who has access to shared accounts
- changing credentials at regular intervals
- using change management control and physical access control mechanisms

Anti-pattern 3: Lack of authentication and data security

Authentication solutions which are commonplace in corporate networks, such as Microsoft's Active Directory, that define what users have access to, and what type of access they have, may not be deployable for some aspects of ICS/OT environments. Likewise integrity solutions such as encryption may not be available within ICS/OT environments. In some cases, especially where safe operation of the environment is a strict requirement (such as relaying trusted data from a sensor to a controller) it is critical to ensure that the received data is authentic, trustworthy and reliable, to ensure that the safe operation of the system is not affected. If the integrity of this data can't be guaranteed, actions taken by the system are no longer predictable, which could allow an attacker to carry out uncontrolled, arbitrary actions.

What's wrong with this pattern?

Leaving ICS/OT and safety-critical systems with no authentication or integrity measures, means that whilst the system can be proven to operate safely within normal parameters, if an adversary were able to forge messages, the safe and predictable operation of that operation is compromised. If there is a reason why authentication and integrity solutions cannot be adopted, it should be raised as a risk within the risk assessment process and appropriate mitigations put in place.

Ideally, if the products you're using don't allow for authentication, you should be raising it with your vendor and/or considering whether the risk of using such products is still acceptable to the business.

There are a number of ways in which this anti-pattern can manifest itself:

- **Data sent without integrity protection.** When data is received from another source, if there is no variable hash that proves integrity (i.e. that the data has not been modified), there is no way of telling that it came from a trustworthy source.
- **Data sent without any form of authentication.** Whilst integrity checks may be in place to ensure that data arrives without any corruption, this is no guarantee that it came from the original sender. Both the sending device and user need to be authenticated (the NCSC has produced guidance on [device authentication](#) and [user authentication](#)).
- **Untrusted data is turned into an action.** If the integrity and authenticity of input data (such as sensor data) has not been verified, but that data is then turned into a corresponding action (such as actuator movement or authority to carry out a task), then it's impossible to find out whether the input data was trustworthy (and if the corresponding action should have been taken).
- **Revoked, invalid and outdated certificates.** Authentication of a peer user or device is often achieved through a PKI certificate. These certificates must be checked to ensure they've not been revoked by the issuer, or past their expiration date, and have valid signatures that can be traced back to a trusted root authority. If invalid certificates are presented and warnings are 'clicked through' or ignored, then there is a risk that an attacker could impersonate the true certificate owner.
- **Use of ICS/OT protocols without additional communications security.** Many historic ICS/OT protocols were designed with little or no cryptographic security. Where communications channels are not physically secure, they should be used in conjunction with security protocols that provide appropriate authentication (such as Secure Authentication additions to the DNP3 protocol, standards such as IEC62351 for security of 60870-5 and 61850 series protocols, or sending the data through a TLS tunnel).

A better approach

Introducing authentication and integrity checks on messages within a system (in addition to having the [data between devices encrypted using TLS](#)) ensures messages exchanged between them are known to be complete, untampered, and from a known source. Where legacy systems may not be able to carry out the necessary cryptographic operations, then a defence in depth approach needs to be implemented to secure the network (ideally via physical security, so it's not possible to connect untrusted devices to the network).

However, authentication is not a panacea. Whilst it guarantees that the messages in peer systems are genuine and have not been tampered with, a legitimate system may have been compromised by an attacker. This means a maliciously created, modified or suppressed message could be sent by a compromised system that presents the correct authentication credentials. Therefore, the full range of defence in depth control measures must be considered to ensure that the system is trustworthy (including the likes of host-based and network-based intrusion detection systems).

It also requires selecting the right encryption to be used for the right application, based on risk appetite and owner concerns, as covered in [the NCSC's guidance on using VPNs](#).

Anti-pattern 4: Inaccurate asset inventory

ICS/OT environments present their own unique set of challenges for asset discovery and knowledge. Asset discovery and maintenance of asset registers within an ICS/OT environment can be very difficult. ICS/OT systems are often designed to last a minimum of 25 years, and many systems are in service much longer than that. Over time, parts are replaced, systems are modified for new requirements and new systems are added. If accurate records of these changes are not maintained, then it is easy to lose track of what is installed. In addition, the data required about each asset will naturally change over time, maturing with the ICS/OT security discipline of the Asset Owner. ICS/OT network security was not a mature discipline when older systems were designed, so information about network configurations, MAC addresses, firmware versions etc. would not have been recorded in a centralised location.

What's wrong with this pattern?

The increasing quantity of devices connecting to ICS/OT networks is expanding the potential attack surface. Without a full appreciation of all devices, components and services that exist in a system (and a full understanding of their function and purpose) it is not possible to ascertain whether the environment is secure and what potential risks exist. By not having an accurate understanding of the assets within the ICS/OT environment, if you're trying to undertake improvement work then you will be doing this inaccurately and will miss elements you may need to secure, which in turn means wasted effort, investment and incorrect solutions and controls employed.

Documenting every change that is made is equally critical, as it ensures that any deviation is known, its effects communicated and why it was carried out in the first place. When an incident occurs, the first steps of incident response and forensic root cause investigation will typically involve a review of the baseline, where efforts may be hampered if a deviation exists that is not documented or fully understood.

When an asset develops a fault and is taken to for repair, software configuration changes may be made which are well within the remit of a maintenance operator, but cause a deviation from the baseline, 'good' configuration. This again, can hamper investigation and recovery efforts if the changes are not fully documented and communicated to the asset owner. Some problems include:

- varying configurations between similar/like units in an estate of devices
- lack of a standard configuration which is rolled out across an estate of devices
- allowing configuration changes to be made (either in software or on the device) without communication and logging change in the asset register
- lack of communication of configuration changes to stakeholders

A better approach

Using an asset inventory is about visibility and having a single view of what is being managed. It provides details of known assets which can be used to develop a configuration baseline, from which a change control process can be used to manage changes.

It is important to recognise that the asset inventory is essential to an organisation's cyber security program; it provides the map on which the protect, detect, respond and recover processes are

built. In an ICS/OT environment, the approach to asset inventory needs to incorporate the context of ICS/OT, especially the need for asset management at scale, identifying priority assets for updates, compensating controls that must remain in place where patching is not possible, and holding the security and backup status for each device.

Understanding and identifying deviations from a 'known good' configuration is critical, where regular baselining of assets allows asset owners to identify where configuration changes have been made from the baseline. Document any local changes with appropriate approvals to ensure that the business understands:

- the reasons for that change
- the impact of that change
- what risks are associated with that change in the long-term

This ensures that active, conscious decisions are made, their safety and security impacts are well-understood, and that any changes are universally made (and if not, the reasoning is accountable and traceable).

A fully detailed asset inventory that is continuously updated is essential for managing an operational environment. It can also assist with investigating the impact of a vulnerability or the extent of a compromise. An asset inventory for ICS/OT holds the configuration baseline including the software, hardware and firmware. A full and complete asset inventory should contain the full list of devices within an environment, and for each device key information such as the model and manufacturer, version, IP address, location etc.

Where possible, a Software Bill of Materials (SBOM) for each application provides a transparent view of the software components and their origin. When vulnerabilities become known in a piece of software, a SBOM helps the asset owner to assess the impact of the vulnerability on their ICS/OT or IT environment.

Anti-pattern 5: Unchecked backups

If there is not a clear and prescriptive secure back up strategy for the ICS/OT environment, then one of the problems that can be encountered is unchecked backups are created. Backups are taken in an adhoc manner with no checks undertaken to verify the integrity and accuracy of the data being backed up, if the correct data has been backed up, and if the backup storage solution is secure.

What's wrong with this pattern?

If these backups are not verified, or they are the incorrect data, or have been tampered with, then there is no guarantee that, when a system fails, it can be recovered in a reasonable period, and that the safe return of operation cannot be guaranteed. Some example ways in which this anti-pattern manifests include:

- lack of tested, checked and documented backups for assets in the ICS/OT environment
- lack of disaster recovery processes and procedures (including testing)
- lack of creation of 'milestone' backups, created on a regular basis, or when significant changes are made to the environment which invalidate the previous backup
- lack of secure storage of backups including [offline storage](#)
- no firmware backup held for assets in an ICS/OT environment

It is essential to be able to restore systems to a clean state following a compromise. Regular checking and testing of backups is therefore an important part of recovery preparations, so the operational environment can be restored to its pre-compromised state.

In the IT domain, backups have proven invaluable in recent years, in particular during the response of a ransomware attack. In the ICS/OT domain, having validated backups is probably more important, as many ICS/OT environments are part of critical national infrastructure. These backups provide a trusted recovery point in which an environment can be safely reset to in the event of an incident. At the same time, when a component has to be replaced, the known 'good' state enables rapid return to operation as a 'plug and play' repair.

A better approach

Within an ICS/OT environment, configuration backups (including control programmes) for critical assets used within an ICS/OT system should be made on a regular basis and stored in a secure location. [These backups should be tested](#) to ensure that, in the wake of an incident, the backups are reliable and can be trusted to restore the function of a system.

For a real-time environment where operations are managed from real-time data, frequent automated backups are needed to ensure recovery data is near to real time (whereas an unchanging data set would require less frequent backups).

More than one storage location and backup are preferable. Ransomware can target backups and encrypt those as well (if they are connected to the network), [so storing backups offline or offsite is also important](#). There should always be a backup available that is **not** connected to the network. If there cannot be any downtime for essential functions, then a backup site with clean hardware to allow a full and rapid restore may be necessary.

The recovery of operations and services should be planned in advance by considering the risks involved. This requires an impact analysis on operational systems (including potential safety consequences) to plan the process of restoration and decide the priorities. This can be informed by the likes of [Crown Jewels Analysis](#) and other risk assessment processes. The importance of minimising downtime is likely to require a fast recovery of functions. If the rollout of backups and restoration needs to be done at speed, it must become a familiar process that is practiced regularly.

The frequency of backups required depends on the nature of the organisation and the volume of data it can create. An operation which follows an unchanging set of protocols may require a less frequent backup than an operation that relies on a feed of time-critical process data to inform operations. Backup systems can perform with near real-time snapshot backups to ensure all data is completely up-to-date, and therefore enables recovery as near to the incident as possible. Note that simply ignoring the ransom demand whilst restoring your systems from backups may no longer be an option, as operators now have to consider the possibility of sensitive information being made public (as we explain in our ['The Rise of Ransomware' blog](#)).

Accurate inventory information is emphasised during restoration efforts. Recovering a clean version of data to a huge quantity of devices will require a clear mapping of digital assets to physical assets.

References

Anti-pattern 1: Flat, unsegregated/unsegmented networks

- [NIST](#) recommends network segmentation and segregation as “one of the most effective architectural concepts that an organization can implement to protect its ICS”.
- [Standard IEC62443](#) recommends partitioning into zones and conduits and restricting data flows, creating separate zones based on different security levels, with communication between zones going through conduits.
- [ISO 27001 control A.13.1.3](#) also recommends segregation in networks.
- The [NCSC Cyber Assessment Framework \(CAF\)](#) recommends segregation in designing for security and resilience by specifying the following outcomes:
 - CAF B4a Secure by Design. “Operational systems for the essential function are segregated in a highly trusted, more secure zone”.
 - CAF B5b Design for Resilience. Operational systems “are segregated from other business and external systems by appropriate technical and physical means.” “Internet services are not accessible from operational systems”.
 - CAF B3c on protection of Stored Data. “You have suitable, secured backups of data to allow the operation of the essential function to continue should the original data not be available. This may include offline or segregated backups, or appropriate alternative forms such as paper copies”.
 - CAF B4c Secure Management. “Systems and devices supporting the operation of the essential function are only administered or maintained by authorised privileged users from dedicated devices that are technically segregated and secured to the same level as the networks and systems being maintained”.
 - [NCSC Design principles and Operational Technology guidance](#).

Anti-pattern 2: Uncontrolled access to ICS/O networks

The [NCSC Cyber Assessment Framework \(CAF\)](#) recommends that systems provide technical controls to prevent compromise and also event detection, specifying the following outcomes:

- CAF B2b.Device Management “Dedicated devices are used for privileged access” and “you fully know and have trust in the devices that are used to access your networks”
- CAF B4b. Secure Configuration “You closely and effectively manage changes”.
- CAF B4c.Secure Management “Systems and Devices supporting the operation of the essential function are only administered or maintained by authorised privileged users from dedicated devices”.
- CAF C1a. Monitoring Coverage “You monitor traffic crossing your network boundary”
- CAF C2b. Proactive Attack Discovery “Routine searching for system abnormalities indicative of malicious activity”.

Anti-pattern 3: Lack of authentication and data security

The [NCSC Cyber Assessment Framework \(CAF\)](#) recommends that systems provide technical controls to enable authentication and data security, specifying the following outcomes:

- CAF B.4. – System Security and B.3. Data Security - Integrity of messages is recommended
- CAF A.4 – Supply Chain Risk - Vulnerability management is covered.

<https://csrc.nist.gov/publications/detail/sp/800-131a/rev-2/final> details which algorithms are deemed acceptable.

Anti-pattern 4: Inaccurate asset inventory

The [NCSC Cyber Assessment Framework \(CAF\)](#) recommends that systems provide technical controls to enable authentication and data security, specifying the following outcomes:

Deviations from configuration baselines, and a lack of understanding of the *what* and *why* can affect compliance with the NIS Directive (CAF B.4. and is highlighted in [NCSC Cyber Assessment Framework \(CAF\)](#)

- CAF B.4. System Security
- CAF B.6. – Staff awareness and training

For any safety-critical or automated process, it is vital to know whether the asset is carrying out the functionality it states it does, something which the RITICS/NCSC '[Effective Solutions to Comply with the NIS Directive – Supply Chain Requirements](#)' project has assessed and identified a number of recommendations to ensure security and traceability in the [supply chain](#).

NCSC has also issued [generic asset management guidance](#) and specific guidance related to Asset Management within ICS/OT Environments.

Anti-pattern 5: Unchecked backups

The [NCSC CAF](#) sections B.5 (Resilient Networks and Systems) and D.1. highlight how it is critical to have reliable working backups, the value they provide towards resilience, and having effective recovery plans in place.

NCSC has issued various items of guidance around backups that include:

- [Offline backups in an online world - How to protect your backups that are stored in the public cloud.](#)
- [Cloud backup options for mitigating the threat of ransomware](#)
- [10 Steps to Cyber Security - this includes a step around data backup](#)
- [Mitigating malware and ransomware attacks](#)

A key aspect of this guidance is to ensure that backups are not left accessible on the network, to ensure that an attacker can not compromise the backup.

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), with support from the National Cyber Security Centre (NCSC). This guidance is not intended to replace formal NCSC guidance where already available, and care has been taken to reference such existing guidance where applicable.

This document is provided on an information basis only, and whilst ICS-COI members and NCSC 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 NCSC and 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 or NCSC. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.