

# CNI system design: Secure Remote Access

A Critical National Infrastructure (CNI)–specific look at NCSC guidance on remote access architecture design

Adam B

Remote Access is an essential feature of just about any realistic IT (and OT) estate.

Whether you're separating your compute resources for resiliency and redundancy purposes, or your estate is spread over a wide geographic area, the need for Remote Access is not going away.

But, there's a problem with Remote Access: if you can access resources remotely, potentially, so can attackers!

The NCSC has published guidance and blogs on this subject for more traditional IT scenarios, but in my engagements within the CNI sector, I often hear that this content does not apply, or is not relevant. In this blog post I'd like to dispel that misconception by highlighting several pieces of guidance that can be applied successfully to design remote access systems for the CNI sector.

---

## A few preliminaries

I'm going to walk through the NCSC guidance which will help you to design and improve a remote access architecture in an OT environment. In doing so, I've assumed that the networks under consideration are being built from scratch, but I appreciate that this is not always going to be the case.

If you find yourself retrofitting secure remote access into an existing solution, you should take the key principles I've highlighted and use them to help you identify and prioritise small packages of work that will iteratively build a more secure and robust remote access solution.

---

# The requirements

The first stage of any design process will be to gather requirements for the system. Here is a fairly typical set of user stories:

- As a standard user, I need to access internally provided services
- As a standard user, I need to access internet resources (such as websites and email)
- As a standard user, I need to access services from both corporate sites and at home.
- As a standard engineering user, I need to remotely access operational technology networks, sensors, and actuators
- As an administrator, I need to remotely alter actuator positions to support update/installation/maintenance procedures
- As an operator, I am obliged to provide the means for vendors/suppliers to extract efficiency and management data

## Security Requirements

To the user stories we must add a typical set of security requirements:

- [Make compromise difficult](#) (from the design principles)
- [Make disruption difficult](#) (from the design principles)
- [Make compromise detection easier](#) (from the design principles)
- [Reduce the impact of compromise](#) (from the design principles)
- Ensure integrity of control system data
- Ensure confidentiality of control system data (where appropriate)
- Ensure confidentiality of sensitive data (where appropriate, for example personal data or intellectual property)
- Reduce opportunity for operational impact from the internet

Together, these requirements define what remote access is, and give us a clear idea of the related security requirements. Next, we'll take a look at the

architecture which underlies such a system.

---

## Architectural review

When designing a remote access system, you should begin by assessing how the underlying architecture can meet the stated functionality and security requirements.

In some cases, remote access requirements are considered late in the architectural design process, but this often causes unnecessary re-designs.

By including a remote access solution early in the design plan, you can help both security and functional requirements, without additional delay or cost.

---

## Devices that access resources remotely

A remote access session should originate from a corporately managed device in the possession of a member of staff, to provide a high level of trust. The device could be any suitable phone, tablet, laptop or desktop machine.

Regardless of the device, the security design principles remain the same. [The NCSC's Mobile device guidance](#) details many security deployment practices, for a wide variety of Operating Systems.

---

## A prime target

The device used for remote access is a valuable target for attackers because it easily provides access and credentials for the components or services they wish to target.

Given the threat from compromise via spear phishing and other targeted attacks, you should separate corporate functions such as email, from engineering functions such as remote access.

The purpose of this separation is to prevent the compromise of a single system component in an exposed service impacting critical and sensitive components. This could be either as a result of an attack or misconfiguration.

---

## No to BYOD

A Bring Your Own Device, or BYOD, model could be considered here, but we will rule it out for now, owing to additional administrative complexity and reduced management control over personal devices. [Further reading on the NCSC's stance on BYOD can be found here.](#)

---

## Yes to PAWs

In cases where the impact of a compromise would be critically damaging, you should use a dedicated Privileged Access Workstation, or PAW. This gives you the opportunity to add additional controls and measures, which will reduce the attack surface.

Gaining trust in your management devices, and the value of Privileged Access Workstations, are covered in depth in the [Secure System Administration guidance](#).

---

## Network infrastructure that enables access

Considering the infrastructure that supports remote access is important because it is key to providing access for legitimate users but must be robust enough to ensure attackers cannot gain access.

Remote access infrastructure should give you the opportunity to logically separate different networks by the functions that they support. For example, it's extremely unlikely that your control network needs access to email or other internal services. Separate functions where you can.

You should seek to include the concept of Privilege Access Management (PAM). This has been described in the NCSC blog '[Protecting system administration with PAM](#)'. Specifically, you should avoid the concept of 'Browse-Up,' as detailed in the NCSC's [Security Architecture Anti-Patterns](#) [whitepaper](#).

---

## The devices being accessed

New Operational Technology equipment has features and capabilities that are comparable to those of traditional desktop IT. This means that approved cipher suites and protocols are readily supported, with little configuration – as discussed in the NCSC's recommended [TLS Configuration](#) or [IPSec Configuration](#).


However, not all Operational Technology networks and devices are new. Pre-existing equipment is often dated and should be replaced when possible. Where obsolete technology cannot be replaced, the NCSC has some [obsolete platform guidance](#).

One area that is often not given due consideration is the protection of the management interfaces used by administrators or operators to configure the remote sensors and actuators of the operational devices being managed. This could cover a wide range of protocols and technologies. It could even refer to the physical attributes of the devices themselves. Further reading about management interfaces can be found in the NCSC blog titled [Protection of Management interfaces](#).

---

## Architectural design

In light of the above architectural advice, let's consider an example system. The diagram below lays out the main components of a generic example or use-case.



*This image shows two dataflows. The orange flow is for engineering access from one device. The blue flow shows access to internal corporate resources and the internet.*

## Review of design choices

### Remote access devices

We have decided to give those who require remote access two devices – illustrated by the Orange and Blue laptops in the diagram. All corporate day-to-day tasks can be accomplished on the blue laptop.

The orange device is not capable of accessing the internet or corporate services, such as email. This is intended to [prevent compromise via phishing](#) or watering hole attacks that would result in an attacker gaining access to the remote sensors and actuators.

The dedicated device is locked down and hardened following a principle of least privilege.

Any unnecessary software should be removed as this could provide a legacy issue that introduces vulnerabilities.

---

## Two laptops?

Providing a second device may seem over cautious and unwieldy – especially if you considered using a hypervisor on a single machine with an environment for each function. However, a second device provides a strong physical and procedural defence against attackers targeting control systems and networks.

Two laptops may be undesirable but is undeniably the most robust form of separation. However this option is not always available and you may be forced to explore other options.

Regardless of the approach used, it is important that common methods of compromise for enterprise IT cannot easily lead to compromise of OT systems. For example, you could consider allowing access to enterprise apps such as email and web browsing by running them on a separate machine such as HyperV, or separating the two on the same physical hardware using virtual machines. In the latter case, the more vulnerable enterprise machine should be a guest rather than the host.

The deciding factor here is the impact of compromise within a control system.

Further reading can be found in the NCSC's [network architectures for remote access](#) guidance. This also covers some of these architectural points from a Zero-Trust perspective.

---

## Network infrastructure

Direct access to the field site from a remote office or home machine is not possible. The only way to do this now is via a dedicated, corporately provided, network that has appropriate separation. This is demonstrated by the dedicated and separate flows from the management device to the field site via the corporate offices hosting the network controls.

The 'corporate' device – as indicated by the blue laptop – that is used to access email and the Internet is only routable via a different corporate network.

Network boundary devices (orange hatched boxes) control data flows in and out of each physical location.

---

## Operational Technology Devices

### Legacy devices

Each sensor and actuator in the field sites is accessed by different methods, and has varying degrees of support for modern cryptographic controls that ensure both integrity and confidentiality of the communications.

In cases where the sensor or actuator cannot provide sufficient integrity or confidentiality, routing communications via a Virtual Private Network (VPN) can provide this support.

### Browse down

In this design, we have ruled out accessing remote sensor networks from a standard corporate device.

We should assume that internet facing devices are at a lower level of trust than the target devices and network. Allowing these low-trust devices to connect to the operational devices is known as a browse-up architecture. We are seeking a browse-down approach.

## No jump boxes

We have also avoided using a jump-box or a bastion host, as these add complexity without necessarily improving the security posture of remote access. You should expect that malicious software is capable of laterally moving through an intermediary host in the same way that you can.

---

## Secure system administration

Now that we have explored how we connect to a remote site's sensors or actuators, the next step is to define how this connection should be used.

For this, the NCSC has recently published its [Secure System Administration](#) guidance. This is a comprehensive suite of recommendations on how to administer services and systems securely. It introduces how these concepts should be applied by a CNI organisation.

---

## Good system hygiene

We have explored architectural issues, and how administration should be achieved, but there are some other relevant issues that should be considered to ensure good system hygiene.

The NCSC website includes lots of guidance to help with basic cyber security practices such as our secure design principles and our risk management guidance. We do not intend to repeat these here, but the following aspects are important considerations when securing remote access to Operational Technology networks.

### Multifactor Authentication

The NCSC blog introducing our Two-factor authentication(2FA) guidance highlights the benefits of a second-factor when accessing services. While the guidance can be interpreted for personal use, or for online services, the benefits

have wider application – including critical national services and operational networks.

We can represent a second factor in the diagram below by giving the user a second factor mechanism – in this case a token – in Red. The token could be in a physical form, or generated by a time-based one-time passcode (TOTP).



As previously stated, the credentials used to gain remote access to Operational Technology systems are extremely attractive to attackers. Including MFA in your solution will make it much harder for an attacker to gain remote access to your system should a password for a remote access account be compromised. The NCSC has further guidance on [the use of passwords](#) and [the merits of password managers](#).

---

## Third Party / Guest Access Process

Sometimes it's necessary to provide a hardware vendor with access to the sensor network. While this access is required, it does not mean that we have to give unfettered access at all times.

We could consider a request process to provide such access – this could be as simple as enabling and disabling an account when required, or have the vendor request a time-based, one-time passcode (TOTP).

Guest access and actions should be logged and investigated as appropriate.

It is worth trying to understand how these third parties access your environment – specifically, do they use a common shared host for all of their customers?

Gaining confidence in your supplier's security maturity is as important as your own.

---

## Vulnerability management via asset control

Knowing what assets you have in an IT estate can assist in identifying potential vulnerabilities.

As software and hardware reaches the end of its support life, having an up-to-date record of your physical and digital assets can help future planning for replacement programmes and quickly aid remediation steps, should a critical vulnerability be disclosed. The reduction in time from disclosure to remediation minimises the opportunity for an attacker to use commodity-based attacks against your assets.

For example, an internet facing firewall is a major component of any network security enforcement regime. Understanding the short-term and long-term roadmaps of the products you use is vital to ensuring that your network remains secure. Short-term roadmaps can provide information about expected update schedule windows. Long-term roadmaps can help you understand the vendor's support plans and manage obsolescence.

Vulnerability management is equally applicable to technologies used by provisioned services or devices – such as a leased VPN device or Firewall. The NCSC has recently issued an [alert for vulnerabilities being exploited in VPNs, worldwide](#).

---

## Configuration control management

Configuration files are an often-overlooked digital asset. Version control of running configurations can assist in identifying when changes occur. This provides a known and working baseline and can help to prevent malicious changes to such files.

The [Secure development and deployment](#) guidance introduces the benefits of software asset controls.

---

## IT vs OT

In some cases, Operational Technology devices and networks may be supported and administered by different teams and departments to those handling more traditional network devices and assets. This division is often referred to as OT vs IT – especially when there is friction between teams rising from differing priorities.

For example, security policies may exist for desktop IT so that insecure services such as Telnet are not supported or enabled by default. These policies may not extend to OT, which may lead to components having such services enabled by default for initial configuration but never disabled. In another example, it may be the case that only IT logging events are being monitored and responded to as OT is outside their realm of responsibility.

It is evident that Information Technologies and Operational Technologies are not the same thing. This is reflected in the differing technologies, risks, threats, vulnerabilities, and mitigation techniques.

Notwithstanding these differences, operators should ensure that both OT and IT systems are equally and consistently accounted for in their overall approach to risk management. Not to do so could result in differences and deficiencies in the way cyber security policies are applied and risks are managed across an operator's combined IT and OT estate.

This does not necessarily mean changing one or more team's dynamics, responsibilities, or staff members, or using the same risk management tools, methods or techniques in both IT and OT contexts, but it does mean harmonising the way risks are described and managed across the whole organisation. This will enable good and consistent cyber security risk management and decision making.

The most effective operators are those where any friction between OT and IT teams has been reduced and where the [overall approach to risk management](#) is

applied consistently in both IT and OT environments.

---

## Conclusion

In this blog post I've tried to show that the NCSC's guidance on remote access can be applied to the CNI sector.

The remote access architecture I've described here is one of many approaches that can be used. I have, for example, overlooked cloud based orchestration services. Even so, the key principles and documentation I've indicated, along with the [cloud security principles](#), can assist in your critical evaluation of such services.

### Adam B

Senior Security Architect



**WRITTEN BY**

Adam B

Senior Security Architect

**PUBLISHED**

9 November 2020

**WRITTEN FOR**

[Cyber security professionals](#)

[Public sector](#)

[Large organisations](#)

**PART OF BLOG**

[NCSC publications](#)