waterfall

Rethinking Secure Remote Access to Industrial and OT Networks

Contents

Rethinking Secure Remote Access to Industrial and OT Networks

Summary

Demand for remote work continues to increase materially but concerns about the security of software-based remote access solutions are increasing as well. The latest guidance from CISA, CCCS and other authorities is to stop using VPNs and jump hosts and instead use stronger solutions for both IT and OT networks.

But compounding understanding of the problem and the alternatives is that "secure" remote access vendors are notoriously opaque as to the design of their solutions, and security limitations inherent in those designs.

In this book we examine the structure and security characteristics of remote access technologies including:

- Classic remote access clients
- Firewalls
- VPNs
- Two-factor authentication (2FA)
- Jump hosts
- Cloud-rendezvous systems
- Demilitarized zones
- Protocol breaks
- Next-gen firewalls
- Deep packet inspection
- Unidirectional gateways
- Unidirectional remote screen view
- Timed Switches
- Hardware-enforced remote access

We then look at a number of common attack scenarios, evaluate the most common combinations of the above technologies against these attacks, and observe what kinds of attacks are credible threats to which kinds of defensive technologies.

Introduction: Remote Everything

"Secure" remote access (SRA) has become essential to industrial and operational technology (OT) networks in applications from manufacturing and building automation to heavy industry and critical industrial infrastructures. SRA enables:

- Routine IT helpdesk for OT systems and networks,
- OT experts calling in while on the road,
- Equipment vendors carrying out routine adjustments to minimize maintenance costs and outages,
- Industrial automation vendors carrying out remote management and adjustments,
- Enterprise-security and cloud-based security operations center (SOC) analysts logging in to adjust security alarm levels and investigate alarms,

And much more. The word "secure" in the widely used SRA name is, however, a misnomer – security is a spectrum not a pair of discrete states. Nothing is or ever can be "secure," only "more secure" or "less secure." That said, marketing teams pretty much universally call their remote access solutions "secure" in hopes of convincing us to buy their products. For consistency with this common marketing usage, this book uses the phrase "secure" remote access as well.

All remote access deployments come with risks. Depending on the solution deployed, risks may include:

- Stolen credentials risk attackers simply logging in as if they were legitimate users,
- Software vulnerabilities in software remote access systems risk attackers taking over those systems and attacking what should have been protected OT assets,
- Weak access controls mean even well-meaning users might mistakenly access and manipulate sensitive equipment that should never be operated remotely, and allow attackers such access as well,

And many more. To manage these risks, we must understand the spectrum of remote access solutions that are available and how these solutions perform in the face of different kinds of cyber attacks. We select our remote access and other OT cybersecurity solutions based on these risks and the consequences they may cause.

All this can be very difficult because:

a. Remote access solution providers are notoriously opaque in their description of what their technologies do and how they do it, and
b. Remote access providers are equally opaque as to what are residual risks that remain when their "secure" solutions are deployed

Hence this book, where we describe a wide range of remote access solutions, evaluate their security in terms of cyber attacks those solutions do and do not defeat with a high degree of confidence, and explore the latest recommendations for OT remote access from national authorities in light of these insights.

Terminology: To simplify this book, we use the terms:

- OT assets: HMI workstations, turbine vibration monitoring stations, PLCs, RTUs, Flow Computers, Safety-Instrumented Systems, Protected Relays, embedded circuit boards and any other cyber asset accessible in an OT / industrial automation network.
- Remote laptop: for simplicity, we describe the remote asset accessing the OT network as a "remote laptop," imagining a remote user at a conference, in a hotel room. In practice, this may be a permanent workstation in a central IT office, a cloud-based SOC, or any other remote service provider's permanent location.
- Internet: for simplicity, we describe the network(s) between the remote laptop and the OT assets as "the Internet." In practice, this may be the "real" Internet, one or more intervening IT networks and/or demilitarized zone (DMZ) networks.

The first part of this book looks at a wide variety of remote access technologies and how they work, while the second part evaluates the security of OT-focused technologies and scenarios.

Classic Software Remote Access

Firewalls

Firewalls are Internet Protocol (IP) routers – they forward network traffic from one network to another. Firewalls are routers with filters that examine every message and decide whether the message is allowed to be forwarded between networks or should instead be rejected or more commonly dropped. The filter function in firewalls can be very complex, including features such as built-in anti-virus inspections, intrusion detection functions, intrusion prevention functions, and deep packet inspection for greater flexibility in deciding which messages are and are not allowed.

The widely used "Purdue Model" or "Defense in Depth" (DiD) designs use layers of firewalls with networks between them. Best practice in these designs is that no connection from one asset to another passes through more than one or two layers of firewalls. This practice means, for example, that no PLC should be able to open a connection to any Internet-based host, with that connection "punching through" many DiD layers of firewalls.

Caution: Simply exposing OT assets such as web-based HMIs, PLCs and RTUs to Internet connectivity through one or more layers of firewall is strongly discouraged by most authorities. Usernames and passwords for these exposed OT systems can be guessed or phished (if the passwords exist at all) and OT systems are often difficult to patch promptly and consistently, meaning Internet-based attackers may be able to exploit known vulnerabilities to attack exposed OT systems, without even bothering to phish passwords.

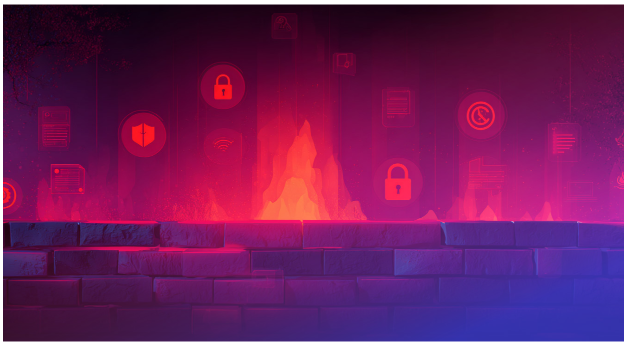

Remote Desktop, VNC, SSH, Et Al

Remote Desktop (RDP), VNC, Secure Shell (SSH) and equivalent technologies are designed for local networks, not remote access. They were designed for a simpler time. They connect to arbitrary IP address/port combinations, encrypt those connections, and show remote users the screens of accessed OT assets. Users move their mice and type on their keyboards and the RDP/VNC class solutions transmit the KMM information to the controlled OT computers.

In principle, we could simply open ports on our IT/OT firewall or worse an OT/Internet firewall, thus exposing RDP, VNC, SSH and other of these remote access tools to remote users on the Internet. In practice this would be high risk. RDP and VNC are difficult to configure so that they use consistently strong encryption, making them vulnerable to Internet-based man-in-the-middle (MiM) attacks. Worse, OT assets are most often difficult to patch consistently, so a subset of our Internet-exposed ports connect Internet-based attackers to poorly patched OT assets, leading to exploits of known, Internet-exposed OS, RDP, and VNC vulnerabilities.

Caution: OT assets with RDP / VNC / SSH exposed to Internet connectivity through a firewall are not recommended by any modern authority or guidance as sufficiently secure for OT remote access.

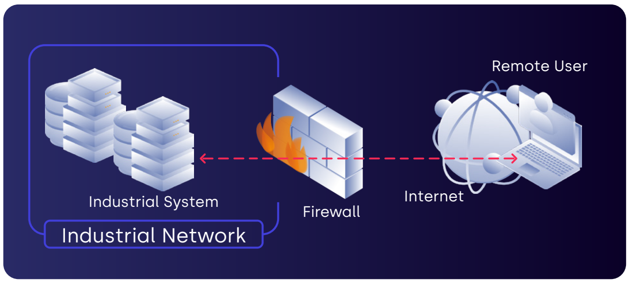

Virtual Private Network (VPN)

A VPN is a way to give a remote computer & user the "look and feel" of plugging a remote OT network switch. The computer has an IP address in the OT network and can send any Ethernet message to the remote switch, as if it were physically connected to the switch. Literally, this is a virtual connection to a private OT network. VPN connections are always encrypted and generally protected by additional layers of passwords and other authentication. Modern VPN systems may include additional features such as checking that the remote laptop is fully up to date with patches and anti-virus signatures as a prerequisite for connecting to the OT assets.

Caution: RDP/VNC/SSH + VPN together are not recommended by any modern authority or guidance as sufficiently secure for OT remote access, in large part because usernames and passwords are routinely stolen through phishing attacks.

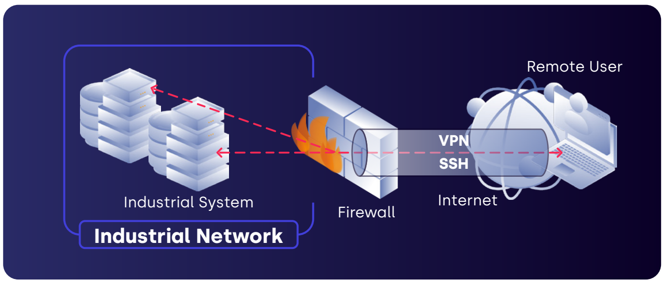

Two-Factor Authentication (2FA)

Two-factor authentication, or more generally multi-factor authentication (MFA), is technology that permits a user to log in and access a system only after presenting at least two kinds of authentication. The most common "kinds" of authentication include:

- Something you know, such as a username and password,
- Something you have, such as a cell phone or key fob with a rotating password, and
- Something you are, such as your fingerprint, iris print or facial recognition.

Note: VPN + 2FA are the bare minimum software remote access recommended by some authorities for OT networks. Communications for RDP/VNC/SSH style remote access solutions can be transmitted to OT assets by the VPN solution for the look and feel of using these access tools as if the remote laptop were connected to the OT network.

Note: It is widely recommended that the VPN and 2FA systems be kept fully and preferably automatically patched to reduce the risk of Internet-based attackers exploiting known vulnerabilities.

Jump Hosts / Intermediate Hosts

A jump host is most commonly a Windows computer or virtual machine set up as the only host that remote laptops can connect to in the OT network. For example, a physical Windows workstation in an electric substation might be set up with RDP, VNC or other remote access software. Each jump host workstation permits one user at a time to take over the machine with a remote laptop and use the jump host to access OT assets. Those assets might be accessed with a second RDP, VNC or SSH session launched in the jump host, or might be accessed by OT tools installed on the jump host, such as local HMI software.

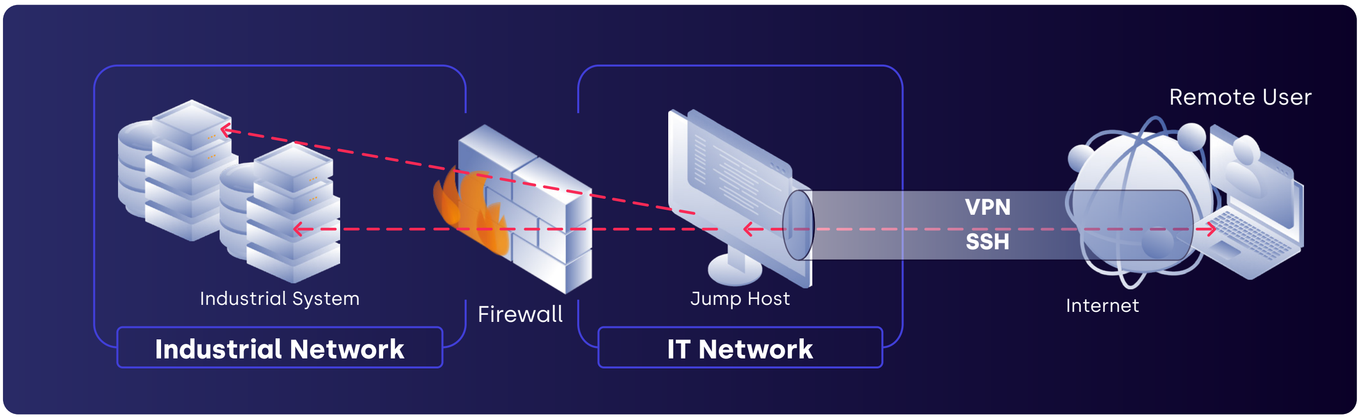

All jump hosts are also intermediate hosts, but intermediate hosts may be either more or less than a jump host. Like a jump host, an intermediate host is a computer that all remote users must log in to as part of their process of connecting with an OT asset. Unlike jump hosts, intermediate host may, for example, also be either of:

- Zero-trust access point: Many users can simultaneously log into a web-based user interface and are given a choice of the OT assets that they are permitted to log into – typically a subset of all OT assets/computers. Users must still have credentials that let them log into the OT assets. The zero-trust access point serves as an extra step in the login process that allows connectivity to only "allowed" hosts for that user, even if the user knows the login credentials for other OT hosts.
- Password vault: This is a zero-trust access point that relieves the user of needing to know the credentials for the OT host the user is allowed to log into. Such vaults are used on IT networks to share local administrator credentials on Windows and Linux workstations. It is used less frequently on OT networks for the same purpose. When a user asks the vault to log into one of the assets the vault has been set up to permit the user to connect to, the vault reaches out and uses its knowledge of the admin password to change the password. The vault then logs the user in with the changed password or gives the user the temporary password. When the user session is complete, the vault once more changes the admin password on the host / device that was accessed, so the user no longer knows the credential or has access to the device.

Note: 2FA, VPNs and jump / intermediate hosts are recommended by many (but not all), authorities as the minimum configuration sufficient to secure OT remote access.

Cloud Rendezvous

There is no standard term for the kinds of connections that Teams, Zoom, GotoMeeting and other conferencing services establish, and so we invent the phrase "cloud rendezvous." In a cloud rendezvous, each end of the connection connects to the same cloud service and the service connects them – forwards screen images from one to the other. Most of the remote access solutions that claim to be designed for OT do this – they have an intermediate host in the OT network and that host connects to the remote access provider's cloud service. Remote users connect to the same cloud service. Depending on credentials, permissions and user requests, the cloud service mediates a "rendezvous" – sending screen images from an OT asset through the cloud service to the remote user / laptop, with keystrokes and mouse movements going the other way.

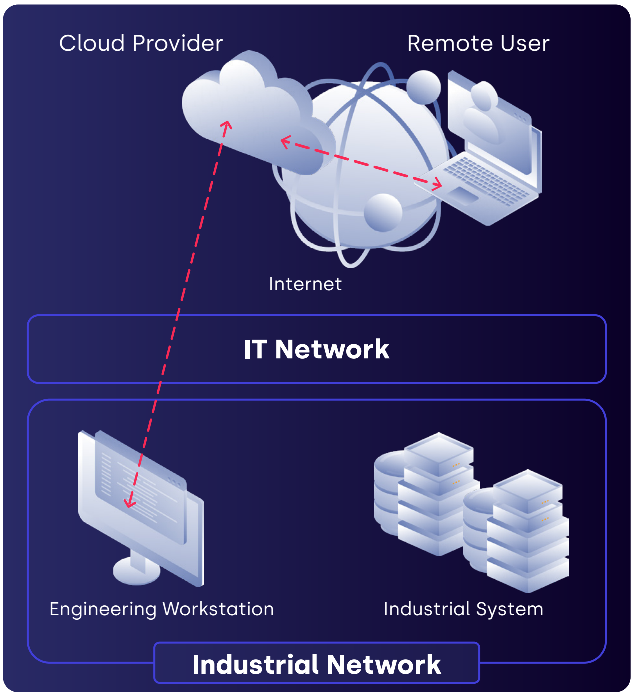

Caution: providers of this style of OT remote access frequently claim "complete" security. This because "no changes are needed to the IT/OT firewall" to allow the providers' OT intermediate host to create a connection out to the cloud. The assumption is that IT/OT firewalls are configured to permit outbound connections but not inbound connections.

In fact, IT/OT firewall best practice is a "deny by default" rule for OT to IT connections, with a specific "allow" rule for every connection to every IT or Internet destination to which an OT asset needs to connect. In this kind of configuration, we do need to add a rule to the firewall to permit the outbound connection. Worse, the cloud service is universally accessible throughout the Internet so that remote users, no matter where they are in the world, can log in to the cloud to request a rendezvous. This means attackers can test the cloud-based service constantly with known vulnerabilities, zero days, phished credentials and so on. Nothing is or ever can be "completely secure." Caveat emptor.

Note: Despite the cloud attack exposure, authorities increasingly recommend cloud-rendezvous style OT remote access for at least low to medium-consequence OT installations, provided these systems are coupled with 2FA to log into the cloud and with Purdue Model style layers of firewalls.

Demilitarized Zone (DMZ)

There is no standard term for the kinds of connections that Teams, Zoom, GotoMeeting and other conferencing services establish, and so we invent the phrase "cloud rendezvous." In a cloud rendezvous, each end of the connection connects to the same cloud service and the service connects them – forwards screen images from one to the other. Most of the remote access solutions that claim to be designed for OT do this – they have an intermediate host in the OT network and that host connects to the remote access provider's cloud service. Remote users connect to the same cloud service. Depending on credentials, permissions and user requests, the cloud service mediates a "rendezvous" – sending screen images from an OT asset through the cloud service to the remote user / laptop, with keystrokes and mouse movements going the other way.

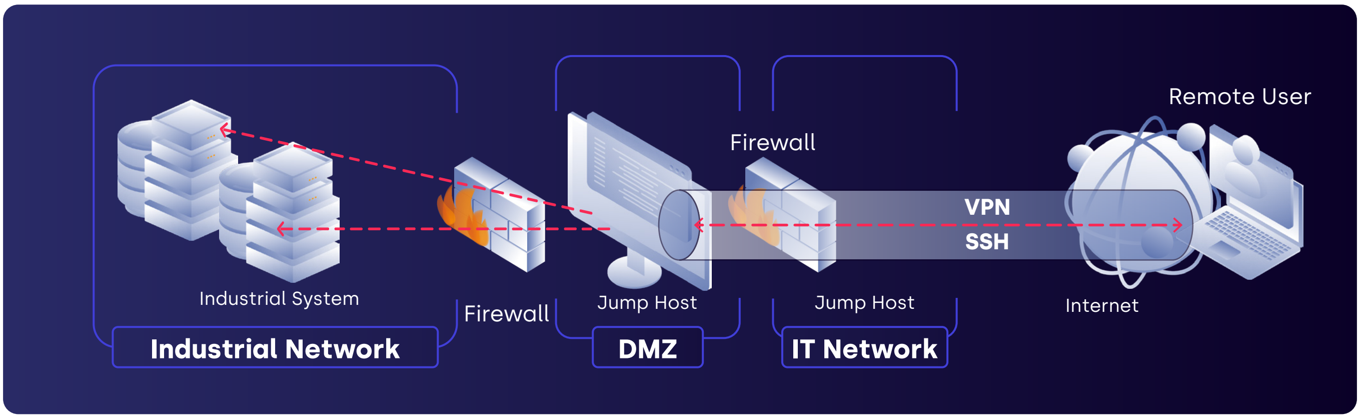

Caution: providers of this style of OT remote access frequently claim "complete" security. This because "no changes are needed to the IT/OT firewall" to allow the providers' OT intermediate host to create a connection out to the cloud. The assumption is that IT/OT firewalls are configured to permit outbound connections but not inbound connections.

In fact, IT/OT firewall best practice is a "deny by default" rule for OT to IT connections, with a specific "allow" rule for every connection to every IT or Internet destination to which an OT asset needs to connect. In this kind of configuration, we do need to add a rule to the firewall to permit the outbound connection. Worse, the cloud service is universally accessible throughout the Internet so that remote users, no matter where they are in the world, can log in to the cloud to request a rendezvous. This means attackers can test the cloud-based service constantly with known vulnerabilities, zero days, phished credentials and so on. Nothing is or ever can be "completely secure." Caveat emptor.

Note: Despite the cloud attack exposure, authorities increasingly recommend cloud-rendezvous style OT remote access for at least low to medium-consequence OT installations, provided these systems are coupled with 2FA to log into the cloud and with Purdue Model style layers of firewalls.

Protocol Break

A protocol break is a term used to describe the Purdue Model rule of not permitting TCP connections to pass through more than one or two firewalls in a DiD stack of firewalls between the most sensitive OT systems and the Internet. A slightly more sophisticated protocol break is where an intervening network, such as a DMZ, has different TCP protocols passing through its "south" or "bottom" firewall than through its "north" or "top" firewall. Eg: plant historians are often deployed in an IT/OT DMZ. In this case, device protocols such as OPC-DA might pass through the "bottom" firewall of the IT/OT DMZ, and the historian client/server protocol might pass through the "top" firewall.

Protocol breaks are widely seen to have security benefits similar to DMZ's, namely that they should slow down cyber attacks somewhat. In principle, to break into our example historian from the IT network takes a phishing attack, or exploitable vulnerability in the historian – one the firewall does not recognize and passes through to the historian from a malicious client on the IT network. Once the historian is compromised, to pivot through to the OT network (use the compromised historian to attack the OT network), an adversary must now phish credentials or find a vulnerability in the OPC protocol.

Caution: in practice protocol breaks tend to have some security value, but they are not "silver bullets." Slowing down an attack may increase the likelihood of being detected before losses are suffered, but there is no guarantee.

Next-Generation Firewalls and Deep Packet Inspection

"Next generation" firewalls (NGFW) are firewalls that carry out "deep packet inspection." Both are marketing terms for firewalls that look deeper into messages and communications than did the previous generation of firewalls when NGFW were introduced over a decade ago. In OT terms, for example, a NGFW rule might say "the DMZ historian is allowed to connect to the OT OPC server, but is allowed to issue only "read" requests to that server, not "write" requests. The previous generation of firewalls over a decade ago could not enforce such fine distinctions in their rule sets.

Caution: Next-gen firewalls are no panaceas. For example, the firewalls themselves have occasional vulnerabilities and zero days that can be exploited. More universally, buyers must be very careful that today's firewalls that claim to "understand OT protocols" do in fact have support for the protocols that will be passing through the firewalls and can make the distinctions that we hope to encode in our rule sets intended for those firewalls.

Remote Access Advice

MODERN APPROACHES TO NETWORK ACCESS SECURITY

Publication: June 18, 2024

U.S. Cybersecurity and Infrastructure Security Agency
U.S. Federal Bureau of Investigation
New Zealand's Government Communications Security Bureau
New Zealand's Computer Emergency Response Team
Canadian Centre for Cyber Security

This document is distributed as TLP:CLEAR. Disclosure is not limited. Sources may use TLP:CLEAR when information carries minimal or no foreseeable risk of misuse, in accordance with applicable rules and procedures for public release. Subject to standard copyright rules, TLP:CLEAR information may be distributed without restrictions. For more information on the Traffic Light Protocol, see cisa.gov/tlp.

US, Canadian and New Zealand authorities recently issued a joint report: "Modern Approaches to Network Access Security". The name is a bit misleading – it is guidance for remote access. Most of the guidance is focused on remote access to OT networks, and most of the recommendations relate to cloud-based remote access and related tools.

The rationale in the report is that many businesses already have most of their employees working remotely "in the cloud" as it were. It therefore makes good sense to use cloud-based services to secure their working environment, rather than try to route all security services through a distant home office. IT remote access advice included:

- Stop using VPN technology – too many VPNs give too much access to IT networks.
- Instead, use a zero-trust approach – deploy technologies that give specific users remote access to only the specific IT assets those users need to do their jobs, not the entire IT network.
- Use Secure Service Edge (SSE) – a collection of cloud services that include:
  - Zero Trust Network Access (ZTNA) – a "security broker," such as a cloud-based service sometimes an IT intermediate host, provides enforces permissions as to which IT services remote users can access.
  - Secure Cloud Web Gateway (SWG) – not specifically for remote access, an SWG is similar to an old-style web proxy server. Remote laptops are configured to route all web requests through the SWG, and the SWG enforces rules about which websites are safe or allowed for the user to access.
  - Cloud Access Security Broker (CASB) – not specifically for remote access, a CASB is technology that sits between remote or local users and cloud services. CASB's can enforce policy regarding which cloud services users are allowed to use, and can have a role in data loss prevention, for example by forbidding and raising alarms when users try to upload a sensitive document to an unauthorized file sharing site.
  - Firewall-as-a-Service (FWaaS) – again is not specifically for remote access. A FwaaS is a cloud-based firewall. Some vendors FWaaS requires software agents to be installed on each protected IT host and remove the IT/Internet firewall. Others still require an IT/Internet firewall, but that firewall simply routes all traffic to/from the cloud-based FWaaS. In both cases, the cloud-based system provides company-wide control over who connects to what, without the need to manage rule sets any longer in individual sites' firewall appliances.

- Secure Access Service Edge (SASE) – a collection of capabilities, including NGFW, SWG, CASB and:

  - Software-Defined Wide Area Networking (SD-WAN) – systems that let us define an abstract IT / cloud-based network architecture, with preferred (eg: cheaper) service providers, alternate providers, and the ability to switch between them as reliability and throughput demands. SD-WAN's can also include the ability to reconfigure firewalls and routing in response to remote access connection requests, creating routes and permissions through firewalls for the duration of the session, and deleting these again after the session completes.

Caution: Automatic updates as well as essentially universal encryption and authentication are assumed by the government guidance for IT remote access. Many OT assets cannot support this kind of management discipline, and so are poor fits for this class of remote access.

Note: The multi-government "modern network access" guidance document recognizes that IT networks differ from high-consequence OT networks, and for the latter networks recommends remote access in the form of unidirectional remote screen view (URSV), timed A/B switches, and hardware-enforced remote access. We discuss these options in the following sections.

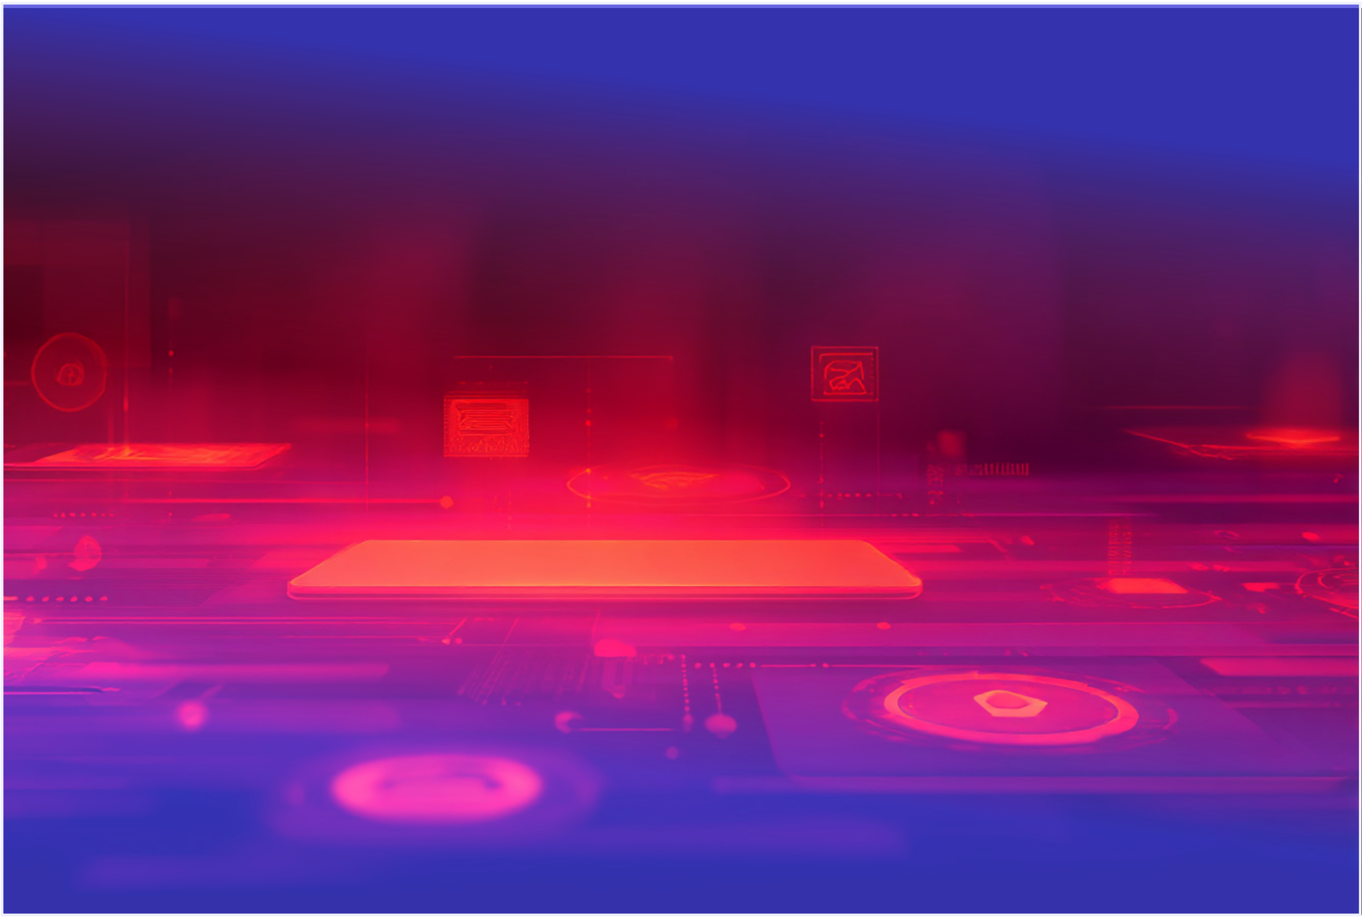

Hardware-Enforced Unidirectional Remote Access

Many authorities recommend hardware-enforced unidirectional remote access for high-consequence OT & industrial sites. In this section, we describe a number of such stronger-than-software technologies.

Unidirectional Gateways

Unidirectional gateways are not of themselves remote access solutions but are essential components of URSV and hardware-enforced remote access. To this end, the NIST 800-82 standard defines a unidirectional gateway as:

"Unidirectional gateways are a combination of hardware and software. The hardware permits data to flow from one network to another but is physically unable to send any information at all back to the source network. The software replicates databases and emulates protocol servers and devices."

All cyber-sabotage attacks are information. The only way a control system can change from a normal state to a compromised state is if attack information enters the system, somehow. The gateways are significant because, while a fully patched and correctly configured firewall SHOULD NOT allow attack information into a protected OT system, the gateways CAN NOT allow attack information back into the OT system, even if they are misconfigured, even if they are incompletely patched, no matter how sophisticated the cyber attack on the gateways.

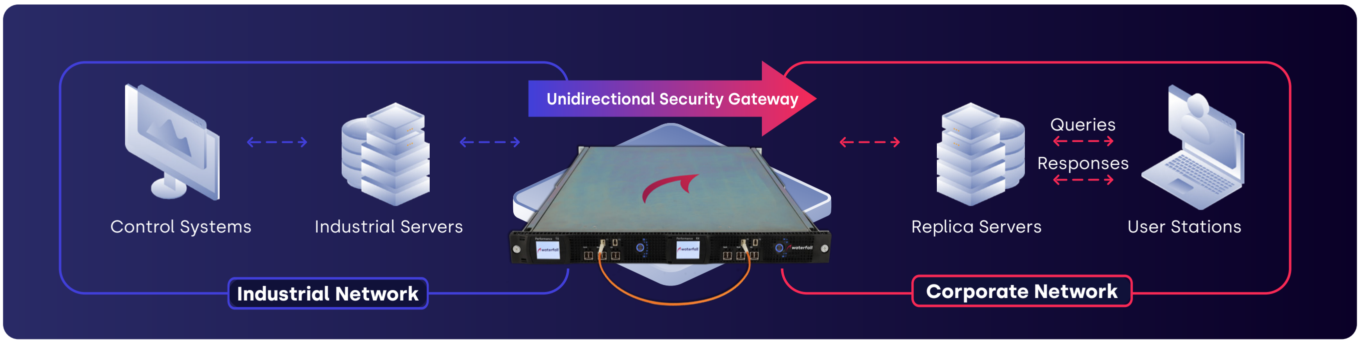

For example, Waterfall Security's Unidirectional Gateway products routinely replicate historians, all the variations of OPC servers, publish/subscribe protocols and more out to IT networks. IT applications, such as the SAP enterprise resource planning system, predictive maintenance applications and others, as well as IT users interact with the IT replicas normally. The gateway hardware means that (potentially malicious) queries for data can no longer be sent into OT systems. The gateway software, replicating servers, means such queries no longer NEED to be sent into OT systems, because all of the data that is allowed to be shared with IT and external users is in the IT network already in the replica systems.

Caution: The unidirectional gateway product space can be confusing. Some vendors claim to have "all software" unidirectional gateways, or "unidirectional firewalls" and other similar-sounding products. Buyers looking for true unidirectional gateways may wish to require purchasers to provide binding attestations that the equipment being purchased complies with the NIST 800-82 definition.

Unidirectional Remote Screen View

Remote screen view is when screen images of OT assets, such as engineering workstations, are sent through unidirectional gateway hardware to service providers on external networks. With URSV, however, there is no way for remote service providers to send keystroke, mouse or other information back through the gateway to remote-control the OT assets. Instead, the service providers are generally on the phone with an OT person sitting at the workstation or other OT system, moving the mouse in response to advice from the remote expert. This is why URSV is sometimes referred to as "attended remote access." To operate equipment remotely, the remote service provider must have the cooperation of, and be in contact with, someone in the OT network with the ability to operate the OT systems.

Note: URSV is used and recommended routinely to enable remote access to high-consequence sites. Waterfall Security's Remote Screen View product for example, is used routinely for remote access to offshore platforms and to enable remote adjustments of large and expensive assets such as steam and gas turbines.

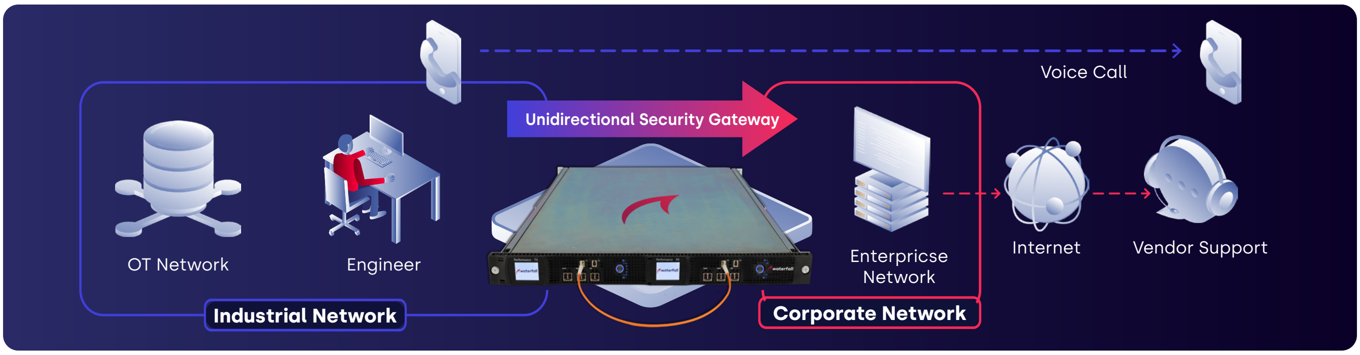

Time-Limited Hardware Switches

Time-limited hardware switches enable temporary connectivity between two computers or networks. These switches are often deployed in parallel with a unidirectional gateway. The gateway replicates servers and devices to the IT network continuously, and the hardware switches control access to a conventional software remote access system, such as a VPN & jump host.

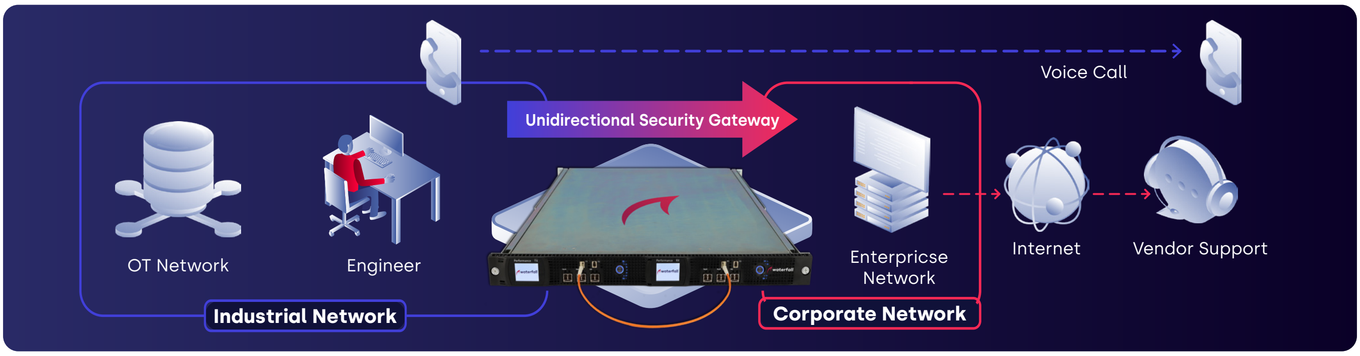

The switches are typically enabled with 1Gbps copper connectivity and are normally in a physically disconnected state, where no electrons can pass through the switch. When a remote service provider needs access to the OT network, they typically call their contact point in the OT network with their request. That person generally hangs up and calls back to the service provider, to verify that the request was legitimate, and not for example an AI-generated forgery. With the request verified, the OT person retrieves a physical key, inserts it into the hardware switch, turns the key and so enables bi-directional copper connectivity into the software remote access system.

The remote expert then logs in normally with username, password, VPN, 2FA and other credentials. After a pre-programmed interval, the switch physically disengages, reverting to its normal disconnected state. The security benefit here is that the software remote access system is connected to external networks and so exposed to cyber attacks only during remote access sessions, and the OT site has physical control over when and for how long remote access is permitted.

Note: This kind of switch is recommended as an extra layer of protection for high-consequence OT sites and is most effective for sites that will enable remote access only infrequently, and for comparatively short periods of time, thus minimizing the amount of time the software remote access solution is exposed to attack from the Internet.

Time-Limited Hardware Switches

This class of technology uses two cooperating unidirectional gateways, one oriented into the OT network, and the other oriented out of the network. Waterfall Security's HERA, for example, sends encrypted screen images through the outbound gateway to a remote client, while sending encrypted keystroke and mouse information into a built-in intermediate host through the inbound gateway. In addition, the hardware of the inbound gateway has a built-in filter that discards all IP traffic, permitting only encrypted keystroke, mouse and other remote access information to enter the OT side of the device.

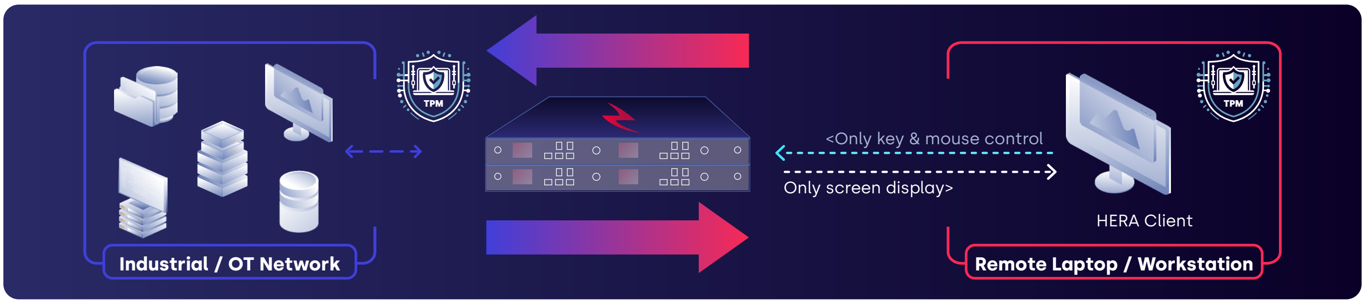

This hardware filtering dramatically improves security because even if an adversary somehow exploits a zero-day vulnerability or otherwise takes over the Internet-facing parts of the hardware-enforced solution, pivoting the attack through the compromised hardware-enforced device is extremely difficult. No attack information at all can reach back into the OT network through the outbound "screen image" gateway. And only encrypted keystroke and mouse movements can enter through the inbound gateway – conventional TCP SYN floods, buffer overflows and other exploits are discarded by the Waterfall hardware.

Even better for defenders, if the Internet-facing parts of the hardware-enforced solution are compromised, that hardware does not contain the key information or other credentials needed to forge keystroke and mouse movements and so take over a session in progress. Those credentials exist only in the remote laptop or workstation, and in the CPUs on the unidirectionally-protected OT sides of the hardware-enforced solution. In contrast, if Internet-facing parts of conventional software remote access solutions such as VPNs, jump host and cloud rendezvous systems are compromised, pivoting through to the next stage of the solution, or directly into the OT network is much more straightforward.

Caution: Not all "inbound / outbound" unidirectional designs are equally secure. Purchasers should ask for a description of the kinds of attacks that candidate solutions defeat more reliably than do software solutions, and how they defeat these attacks, to evaluate security benefits of potential solutions.

Note: True hardware-enforced remote access is recommended by CISA, CCCS and other authorities for remote access to high-consequence sites.

Part 2: Attack Scenarios

To compare the security of two or more technologies or designs, we look at various attack types and which attack types the potential defense technologies can defeat with a high degree of confidence.

In this section, we look at four classes of attacks and evaluate those attacks against remote access solutions. To further simplify the analysis, we combine a number of the simpler SRA solutions into Defense-in-depth layered solution. The result is the following chart. Each row, column and element of the chart is evaluated in the sections that follow.

	High Conn VPNs	Pivoting Attacks	Session Hijack	Client Pivot	Type
2FA, DNZ, VPN, Jhost, NGFW	✗	✗	✓	✗	Unattended
OT SRA	✓	✗	✗	✗	Unattended
Timed Switch	✗✓	✗	✗✓	✗	Minimally Attended
HERA	✓	✓	✓	✗	Unattended
Unidirectional Rem Scr View	✓	✓	✓	✓	Attended

✗ does not defeat

✓ does defeat

✗✓ dramatically reduces attack opportunities

High Connectivity VPNs

The Modern Approaches to Network Access Security report cited earlier describes this vulnerability / attack scenario as "Although some VPNs can be configured to enforce granular firewall policies to provide limited levels of access to company resources, not all VPN providers offer this as an option." This is a most serious vulnerability when any sort of malware has infected the VPN client system, such as a remote laptop that ten minutes earlier was browsing the Internet and clicking on links in email. When a VPN connects a compromised remote system to an OT network with minimal additional segmentation, even the simplest of malware can jump through the VPN connection to infect unpatched or otherwise vulnerable OT assets.

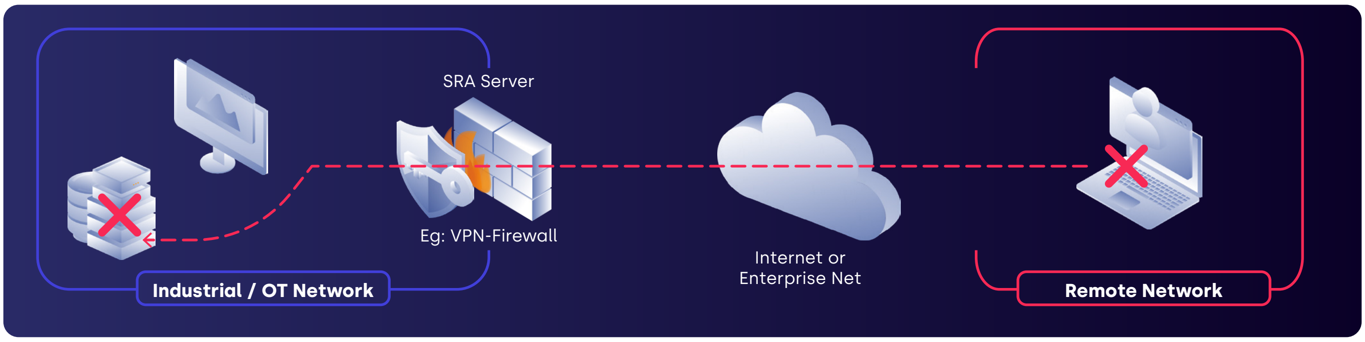

Pivoting Attack on Zero-Trust Server

Today, pivoting is by far the most widely used of sophisticated cyber attacks. Pivoting is defined as "using a compromised host to attack other hosts." Attackers routinely pivot from their attack machines on the Internet through IT/Internet firewalls into IT networks using a host of techniques ranging from malicious attachments, malicious links/downloads, remote access credential theft, and exploiting Internet-exposed vulnerabilities in IT assets. Attackers then use the compromised IT assets to "move laterally," taking over additional IT assets to establish persistence, search for credentials, and eventually pivot again through additional layers of firewalls into higher-value targets, such as customer relationship management systems or OT systems.

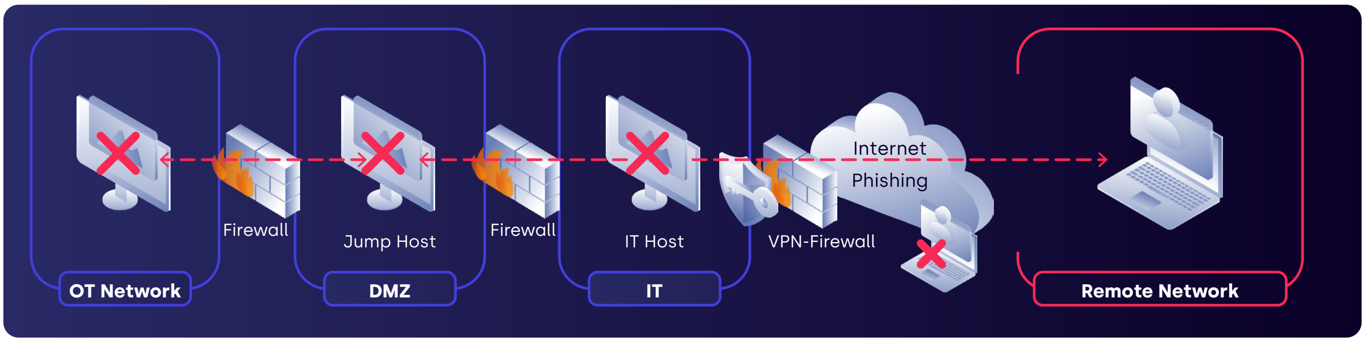

The fundamental problem with software-based remote access systems, both on-prem and cloud-based, is that they are complex software artifacts. All such software has defects, and some of those defects are vulnerabilities. Some of those vulnerabilities we the defenders have discovered and are desperately trying to patch before they are exploited, and some our attackers have discovered (zero-days) and are using against us without our knowledge. The compromised software assets may be IT assets, OT assets, remote access servers, two-factor servers, cloud-based servers or even remote laptops / software-remote-access clients.

Session Hijacking

- Session hijacking is attackers taking over and mis-operating existing remote access sessions. Strong encryption is the standard defense against session hijacking, but despite clever marketing, not all remote access solutions have this capability. In particular, browser-based software remote access solutions are more vulnerable than average to session hijacking. A typical session-hijacking attack follows:

- The attacker compromises the remote laptop with some very simple malware in a malicious attachment or download,
- The malware waits until the machine has opened a web-based remote access session, for example through a cloud-based OT rendezvous,
- The malware steals the session cookies from the browser – a simple operation – and then sends the cookies to the attacker's command and control center and causes a malfunction of the laptop or browser, for example a "blue screen" or hard reboot.
- The attacker promptly collects the cookies from the C2 and reconnects to the OT SRA server as if they were the original user

Features of web-based solutions that make this kind of attack particularly straightforward include:

- HTTPS, the workhorse of browser encryption, is designed to prevent theft of information and credentials by machines on the Internet that the HTTPS connection is passing through – so-called "man in the middle" (MiM) attacks. The protocol has few protections against credential theft within the client laptop.
- Web browsers and HTTPS are designed to work across an unreliable Internet, which means connections to remote web servers can be "broken" and re-established many times, each time negotiating a new encrypted connection. Breaking such connections and re-establishing them from another host in another part of the Internet is straightforward.
- Web-based applications routinely identify which new connections belong to which user sessions using "cookies" – small random-looking chunks of data that are sent with every request to a web server identifying the request. Steal the cookies and you steal the session.

Even though many browser-based SRA solutions try to compensate for this vulnerability with other elements of design, it is difficult to reliably defeat this kind of attack in a web browser / HTTPS / web server system that is intrinsically designed to support intermittent connectivity across an unreliable Internet.

Client Pivot

A client pivot involves much more sophisticated malware and attack techniques on the remote laptop. Instead of stealing credentials (cookies) and resuming the session from another machine, the client pivot takes over the remote laptop and uses the compromised laptop itself to establish remote connections into the OT target.

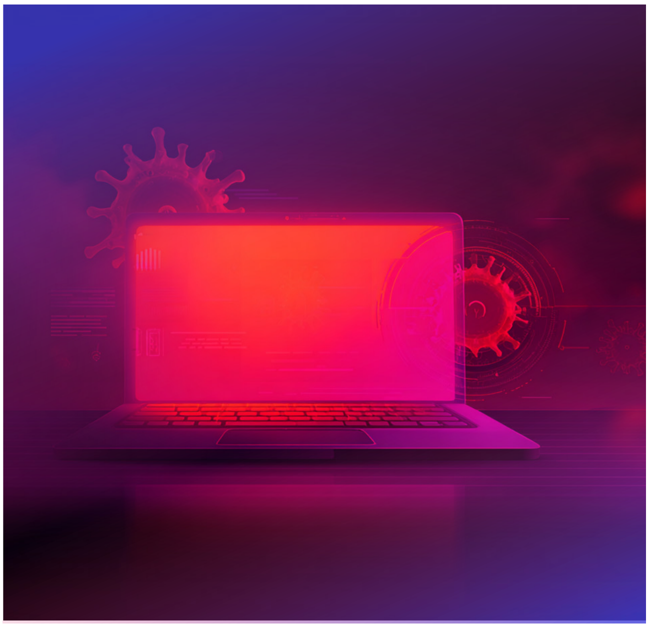

For example, the malware might wait until a remote session is in progress and then create an invisible virtual display, using the same kind of virtual display technology that is used routinely for USB-based external monitors. The difference is that the virtual display is not visible on any monitor. The malware then moves the SRA client window into the virtual display, transmits the contents of the display to the attacker's own machine/display and operates the remote access window maliciously. To distract the legitimate user, the malware then posts confusing error messages encouraging the user to "submit details for analysis" or "press here to diagnose the problem" or some such, to buy the attacker tens of minutes to act in the OT environment.

Note: VPN "split tunneling" protections are intended to defeat this kind of remote control of a remote access laptop, but sophisticated attacks can bypass these protections.

Evaluating Defenses

In this section we evaluate these attacks against five defensive postures

Software Defense in Depth

A software remote access defense in depth system is one that uses all of:

- A DMZ consisting of two next-gen firewalls from two different vendors, one firewall at the IT/DMZ interface, and the other at the OT/DMZ interface,
- A protocol break in the DMZ for all protocols,
- An RDP or other jump host within the DMZ, with no TCP connections permitted to pass through both layers of firewalls,
- A VPN that terminates in the IT/DMZ firewall, and
- The jump host configured with two-factor authentication.

How well does this design defeat each of the benchmark attacks in the previous section?

- High-Connectivity VPN – the VPN terminates in the IT/OT firewall and thus with many VPNs, any common malware on the remote laptop can attack any host in the DMZ, not only the jump host – not reliably defeated.
- Server Pivot – if an attack gains a foothold in the jump host, for example by exploiting a zero-day or other vulnerability in the 2FA system of the jump host, the adversary can use the compromised jump host to pivot and attack OT hosts through the OT/DMZ firewall – not reliably defeated.
- Session Hijacking – in this scenario, we postulate an RDP-style remote access system in the jump host, protected across the Internet by a VPN, not a web-based system. Such sessions are reasonably difficult to hijack, and so we credit this style of attack as reliably defeated.
- Client pivot – the very sophisticated "hijack the window" client pivoting attack is very difficult to accomplish, and equally difficult to defeat reliably – not reliably defeated.

Software remote access solutions provide extremely convenient remote access solutions – remote users can log in to any site for which they have authorization, any time they are authorized – but so can attackers.

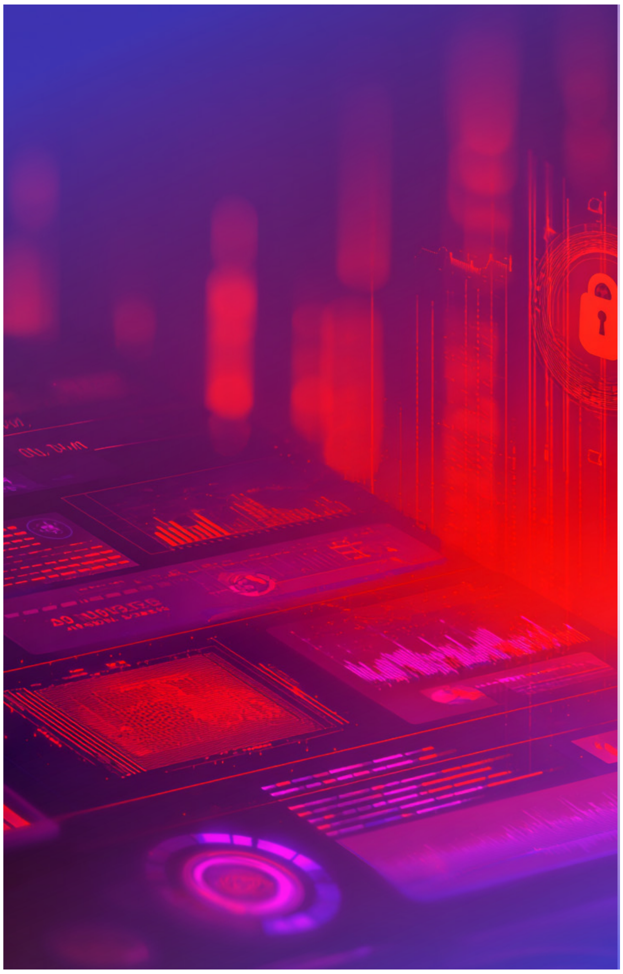

OT Software Remote Access

- This defensive posture is assumed to be any of the many rendezvous-style OT remote access solutions on the market. The OT site is permanently connected to a cloud/web service. A web browser connects to that same service, provides 2FA and other credentials, and is connected to an OT asset rendezvous-style via the cloud-based software remote access system.

- How well does this design defeat each of the benchmark attacks in the previous section?

- High-Connectivity VPN – there is no VPN in these modern zero-trust-based systems. Legitimate remote users can connect to the cloud service only via a web browser and can rendezvous only with systems for which they have authorization – high-connectivity VPN attacks are not possible in this scenario.
- Server Pivot – if an attack gains a foothold in the cloud-based software system, for example by exploiting a zero-day or other vulnerability in cloud-based web server, in the worst case the adversary can use the compromised cloud system to rendezvous with any OT asset connected to the cloud – not reliably defeated.
- Session Hijacking – these cloud-based services are nearly universally browser-based, making session hijacking attacks straightforward – not reliably defeated.
- Client pivot – the very sophisticated "hijack the window" client pivoting attack is very difficult to accomplish, and equally difficult to defeat reliably – not reliably defeated.

Software remote access solutions provide extremely convenient remote access solutions – remote users can rendezvous with any site for which they have authorization, any time

Timed Switch

In this defensive posture we assume a software DiD-style or software rendezvous-style defensive posture, with the addition of a timed hardware switch. When the switch is engaged, remote users have between 30 and 90 minute to interact with the software remote access system before the switch automatically disengages again, isolating the software system.

This kind of hardware-augmented software remote access system is strictly stronger than pure software remote access, because the software attack target is not physically accessible to the Internet most of the time, provided that the switch is in the "disconnected" state most of the time. We illustrate this in the summary diagram as yellow status – stronger than the red "failure" indicators, but not as strong as the green "reliably defeats" indicators.

This kind of remote access is also marginally less convenient than software remote access, because there needs to be a person at the OT site who is able to and permitted to "turn the key" or "press the switch" on the timed hardware device.

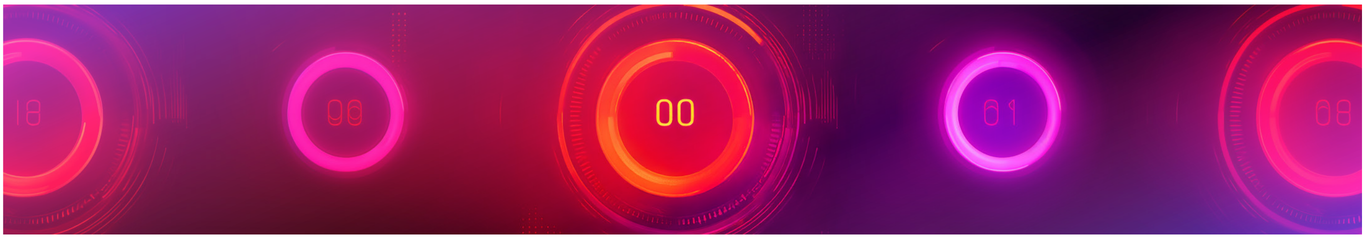

Hardware-Enforced Remote Access

In this defensive posture we assume a hardware-enforced unidirectional remote access system. How well does this design defeat each of the benchmark attacks in the previous section?

- High-Connectivity VPN – there is no VPN in a true hardware-enforced remote access solution, and so this attack scenario is not possible – reliably defeated.
- Server Pivot – if an attack gains a foothold on the Internet-exposed CPUs of a hardware-enforced system, that attack has gained very little in terms of being able to pivot through the compromised server CPUs into the OT network – reliably defeated.
- Session Hijacking – when these hardware-enforced services are not browser-based, session hijacking is very difficult – reliably defeated.
- Client pivot – the very sophisticated "hijack the window" client pivoting attack is very difficult to accomplish, and equally difficult to defeat reliably – not reliably defeated.

Hardware-enforced remote access solutions provide extremely convenient remote access – remote users can rendezvous with any site for which they have authorization, any time they are authorized – and make doing the same much more difficult for attackers than do all-software remote access systems.

Unidirectional Remote Screen View

In this defensive posture we assume that a NIST-style hardware-enforced unidirectional gateway, "pointing" from the OT network out to the IT network, is the sole online connection between any OT and any IT asset. Further, we postulate that remote screen view can be enabled through the unidirectional gateway to a web server built into the IT / Internet side of the gateway. And we assume that a remote support person operating the remote laptop is on the phone talking to a person inside the OT network, giving them advice as to how to move the mouse or use the keyboard to diagnose, correct or otherwise improve the operation of the OT system.

- How well does this design defeat each of the benchmark attacks in the previous section?
- High-Connectivity VPN – there is no VPN – reliably defeated.
- Server Pivot – if an attack gains a foothold on the Internet-exposed CPUs of a hardware-enforced unidirectional gateway pointing from OT out into IT, that compromised CPU is not physically able to send any attack information back into the OT network – reliably defeated.
- Session Hijacking – while the web-based session that observes the screen image feed can be hijacked, the hijacked session is still not physically able to send any attack information into the OT network to affect the OT network in any way – reliably defeated.
- Client pivot – When the very sophisticated "hijack the window" client pivoting attack, or any other client pivoting attack is launched, the compromised remote laptop still has no physical way to send any attack information into the OT network – reliably defeated.

While unidirectional remote screen view is less convenient than either software remote access or hardware-enforced remote access, because of the need for an attendant at the OT site assisting with the remote access session, unidirectional remote access is the most secure kind of remote access. No compromised Internet asset or IT asset is physically able to send attack information into the protected OT network through the outbound unidirectional hardware.

Waterfall Security Solutions

Waterfall Security Solutions has produced OT-focused Unidirectional Security Gateways and related products for nearly 20 years. We are trusted by heavy industry, critical industrial infrastructures and government installations all over the world.

Truly Unidirectional Gateways

Waterfall's Unidirectional Security Gateways are truly unidirectional – physically able to send information in only one direction, usually from OT networks out to IT networks or the Internet, allowing no information whatsoever to flow in the reverse direction. Waterfall's hardware-enforced unidirectional gateways comply with the NIST 800-82 definition in every respect: optical isolation, server replication, and device emulation.

In addition, Waterfall's latest WF-600 series gateways are true network appliances with a wide arrange of server replication software built into the devices and with no need to install hardware or software in other parts of the OT or IT networks. The only user interface is web-based, used for configuration, management, monitoring and even troubleshooting. And the connector library that replicates servers and emulates devices is not only wide but deep, with standard features and options including no-single-point-of-failure high availability, metadata synchronization, and deep content inspection and filtering.

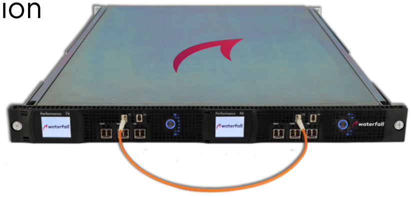

HERA: Hardware-Enforced Remote Access

Waterfall's HERA unattended remote access product is a full-featured unattended remote access product that is stronger than software remote access. The HERA gateway consists of two cooperating Waterfall Unidirectional Gateways, one oriented to send screen images out to HERA clients running on remote laptops and workstations, and the other oriented to send encrypted keystroke and mouse information into the OT network.

HERA uses Trusted Platform Module (TPM) hardware, in both the remote laptop client and the HERA gateway hardware itself, to protect HERA login credentials from theft. This makes the remote laptop hardware itself another kind of 2FA, since logging into HERA gateways protecting OT sites is only possible using the configured laptop hardware to authenticate the connection.

Furthermore, the HERA communications protocol is designed to be filtered by hardware, and the HERA gateway hardware and software all assure that only encrypted HERA information can pass through the gateway, blocking today's sophisticated and ubiquitous IP and TCP/IP-based pivoting attacks.

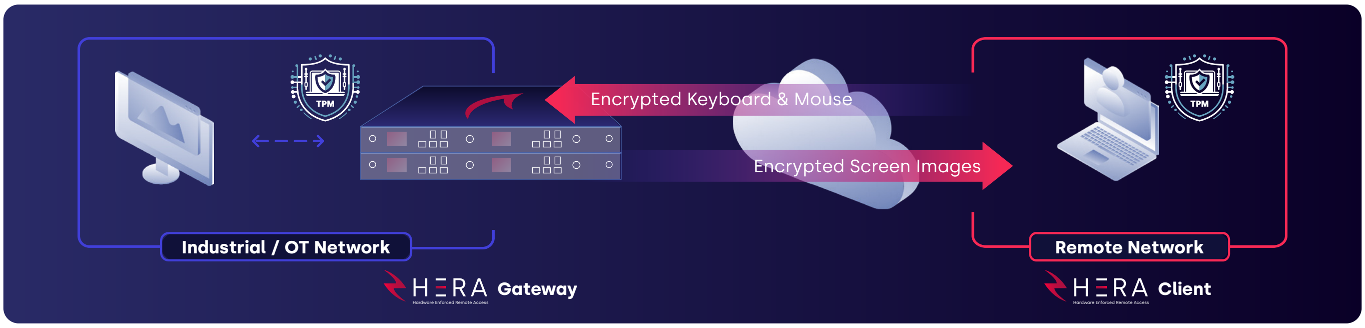

Truly Unidirectional Remote Screen View

Waterfall's Remote Screen View product is hardware-enforced unidirectional attended remote access, allowing OT insiders to activate screen sharing via the one-way unidirectional gateway. Remote users and vendors can see the screens of OT systems, for example engineering workstations, without any ability to send keystrokes, mouse movements or any other information back into the OT network through the outbound Unidirectional Gateway product. Instead of controlling and potentially attacking OT assets directly, remote support personnel typically communicate with an OT insider over the phone, providing real-time instructions as to where to move the mouse, which buttons to click on and what data to enter via the keyboard.

Waterfall Secure Bypass

The Waterfall Secure Bypass (SBP) unit is a timed A/B switch that is activated by a physical key. The SBP requires minimal effort to activate. Most often a remote support person calls someone in the OT network – such as the SCADA system operator – and requests access. That person hangs up and calls the support person back at a known phone number to confirm that the request was legitimate, walks over to a key closet, selects the SBP key and activates the unit.

The operator can then return to their normal duties while the remote expert activates their 2FA / software remote access technology to log into the OT network. After a pre-programmed interval of typically 30-90 minutes, the SBP returns to its normal "disconnected" state, preventing attackers from interacting with the software remote access system when the system is not in use.

Risk-Based Decisions

The landscape of remote access is changing. The latest advice from authorities and experts is that in today's threat environment, VPNs and jump hosts are recommended for neither today's IT networks nor today's high-consequence OT networks. With large numbers of remote workers manipulating sensitive corporate information in their homes, on their Internet-connected laptops, IT networks should use cloud-based solutions that are stronger than VPN/jump host access to IT networks. But with today's sophisticated attacks pivoting through VPN-based software remote access solutions, important OT networks should use hardware-enforced remote access solutions.

With Waterfall's HERA now on the market, the decision to use hardware-enforced security has never been simpler. With HERA providing all the features of unattended remote access that software SRA provides, at essentially the same cost, why use the weaker software-only solutions when a hardware-enforced alternative is easily available?

For more information about Waterfall's HERA,
please contact us for a free consultation with a remote access expert.