TXOne Networks | Keep the Operation Running

The OT Zero Trust Handbook

txOne networks

Table of Contents

Introduction

How much are you willing to risk? That's the first and most important question you must answer when making budgeting decisions for your organization's cyber defenses. Cybercriminals don't attack everyone, so why not buy insurance and bet that your systems will not be targeted? If you get hit, will you pay the ransom or rebuild your systems?

Here's the bad news: if your organization is making money it's already being targeted by attackers. Worse, in addition to setting up defenses against attacks from outside you must also consider how to address attacks from sources that would normally be considered trustworthy. Even trusted and well-meaning employees can carry malware on-site in an unsecured USB drive, and attackers exploit trust in equipment manufacturers by hiding malware in updates or newly-purchased assets. How can you maintain productive operations in such a high risk environment?

OT Zero Trust 101

OT zero trust is an approach to protecting devices, data, and smart machines under the premise that all operational assets and asset activities should be treated as untrustworthy. Traditional cybersecurity, originating in IT, is designed to secure networks based on the needs of user activity. The standard IT "CIA" model prioritizes confidentiality first, integrity second, and availability third. In contrast, OT network security has to be based on asset activity and an "AIC" model where availability is #1, integrity is #2, and confidentiality is #3. By never making assumptions about credibility and continually evaluating trust on the network, architecture based on OT zero trust restricts communication on a "need-to-know" basis.

Traditionally, perimeter defense has been the focus of cybersecurity. Think of your perimeter network like the hull of a submarine, such that one hole in your defenses can sink it. Submarines can resist being sunk because their hulls have compartments that seal off a hull breach, stopping water from flooding the entire submarine. A flat network is like a submarine with no compartments – one security breach and you're sunk very well. When the OT network is segmented into productivity-based zones, cyber threats are more easily contained and suspicious activity is more easily identified.

In our rapidly changing world, it is no longer feasible to base trust on incidental verification – instead, verification must be a regular process carried out in real time by dedicated appliances or software. Policies based on OT zero trust empower cybersecurity teams with real-time threat intelligence, enabling them to comply with regulations while keeping up with productivity goals and mission-critical tasks. You must be vigilant in asking questions about assets, like "Should this robot arm in maintenance mode be transmitting data to my HMIs right now?" The ability to observe asset status and restrict privileges situationally, such as when an asset is in maintenance mode, is one benefit of OT zero trust that helps technicians identify and take decisive action against suspicious behavior.

The 4 Cornerstones of OT Zero Trust

OT Zero Trust relies on four cornerstones: security inspection, trust lists, network segmentation, and asset shielding.

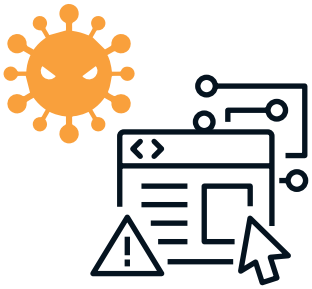

Inspect

Scan all inbound devices brought on site by personnel to stop insider threat, and scan assets before onboarding to prevent supply chain attacks.

Lock Down

Trust lists secure endpoints and networks alike by specifying what is allowed and blocking everything else.

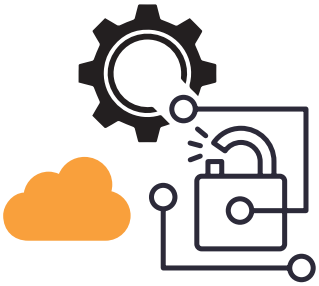

Segment

Network segmentation groups vulnerable assets into operations-friendly safe zones, preventing attackers from moving and malware from spreading.

Reinforce

Shield assets at a network level to secure vulnerabilities in legacy and other unpatched assets without interrupting their work.

Inspect

Scan all inbound devices brought on site by personnel to stop insider threat, and scan assets before onboarding to prevent supply chain attacks.

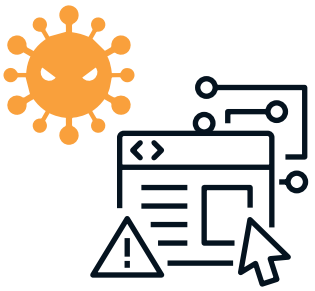

OT zero trust begins with the asset life cycle the moment a device enters the work site. Attackers often get threats into OT environments by making use of a trusted third party's ability to introduce threats. Employees can create insider threat such as this by negligence or by malice. Additionally, supply chain attacks can be set off by attackers hiding malware in an asset before the vendor ships it. These supply chain attacks can impact thousands of companies overnight. Prevent insider threat and supply chain attacks introduced through onboarding assets by implementing a procedure for scanning devices as they enter the worksite.

For a company high on the supply chain, this could cause a catastrophe as every link of the chain gets hit. It's likely that in the near future vendors will need to take a more standardized and rigorous approach to guarantee their products are secure, but until that day every OT organization must take this matter into their own hands. Each worksite needs a checkpoint that is equipped to quickly conduct and log the results of malware scanning and removal. This scan procedure can also be applied to outbound devices to protect your organization from spreading a reputation-searing supply chain attack.

Integrate asset scans and inventories to ensure scan completion, ease audits, and make it easier to confirm vulnerability status of each asset. This creates a record that every device that connects to the network or assets is malware-free. By conducting these scans with a portable scanning device that can be taken from machine to machine, scans of stand-alone and air-gapped systems can be tracked from one centralized console. While some stand-alone systems can be extremely difficult to scan due to strict industry regulations that disallow software installation, the use of a handheld portable scanning device enables malware scanning and cleanup to be conducted from memory.

Lock Down

Trust lists secure endpoints and networks alike by specifying what is allowed and blocking everything else.

Trust lists secure endpoints and networks alike by specifying what actions, assets, applications, or users are trusted to do and blocking everything else. OT zero trust is based on using trust lists to manage application privileges on endpoints and asset privileges on networks. At the endpoint level, they can stop unauthorized applications from executing and ensure that only trust listed users can make changes to configurations or data.

For fixed-use and legacy assets, a trust list is used to guarantee that only applications related to the endpoint's operational role can run. For modernized machines that have more resources and must stay flexible when conducting a wider variety of tasks, trust listing can be informed by machine learning and a library of approved ICS applications and certificates. This allows applications that are already approved and verified to be excluded from scans so they can work undisturbed. Meanwhile, machine learning has the ability to spot suspicious behaviors related to cyber threats without disturbing trust listed processes.

By deploying network trust lists, asset communication privileges can be strictly limited to those necessary to carry out their work. They can also be based on common OT protocols, by which only approved commands can be sent at a network level. This prevents both malicious tampering and misoperation.

It's hard to stop malware that's hidden in a trustworthy-seeming update from a vendor, as nearly 1500 companies learned in the Kaseya VSA supply chain attack of 2021. With an application trust list, updates can't run on your systems until an administrator has scanned, approved, and scheduled them. A deployed trust list puts a stop to any malware that might be introduced to an endpoint by excluding all unlisted applications and all attempted changes to configuration or data by unlisted users.

Determine your OT zero trust policies and enforce them with lockdown software that can be centrally managed from a single pane of glass. Security teams are empowered to monitor the OT environment in real time and take the best course of action.

Attacks on Critical Infrastructure from the Dark Side

In the cyber incident involving a well-known pipeline company, attackers used DarkSide's Ransomware-as-a-Service (RaaS) to stop the pipeline responsible for 45% of the eastern United States' fuel resources for the first time in its 57 years of op-eration,[1] successfully extorting a USD $4.4 million payout.[2] DarkSide has seemingly gone quiet since this attack, likely due to increased attention from government and law enforcement organizations. We can expect, however, continued development on RaaS – RaaS has become an industry, and like any industry will gradually pool its in-novations. The BlackMatter ransomware, for example, includes tools and techniques from the Darkside, REvil, and LockBit 2.0 ransomware families.[3] Our researchers have not yet confirmed but suspect that BlackMatter is actually the DarkSide group resuming operations under a changed name (a common practice among cybercrim-inal organizations).

---

[1] William Turton and Kartikay Mehrotra, "Hackers Breached a Major Pipeline Using Compromised Password", Bloomberg June 5 2021

[2] Associated Press, "A major pipeline confirms it paid $4.4m ransom to hacker gang after attack", The Guardian, May 20 2021 [3] Pedro Tavares, "A full analysis of the BlackMatter ransomware", The Infosec Institute, Nov 10 2021

Segment

Network segmentation groups vulnerable assets into operations-friendly safe zones, preventing attackers from moving and malware from spreading.

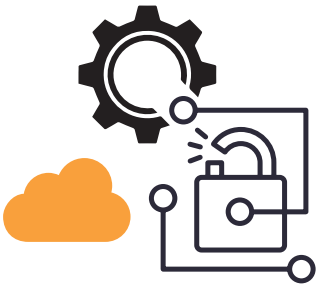

With network segmentation you can split your network up into more easily-defended zones based on which assets do or do not need to communicate. Using OT-native solutions, create protocol-based policies to determine which kinds of commands can go in or out, and which assets can communicate with each other. Well-wrought command-level policy can prevent misoperation in addition to malicious traffic, but this requires support for the OT protocols that the work site's assets use.

The fundamental tools of network segmentation are appliances – IPSes and firewalls. A next-generation IPS micro-segment critical assets or groups of assets that require 1-to-1 protection, while next-generation firewalls transparently create segmentation and broader definition of network security policies. Work site-friendly "OT-native" IPSes and firewalls can be deployed transparently without changes to existing architecture. With this capability, micro-segmentation can be conducted using trust lists set at the network level and at the protocol level using OT-native IPSes or firewalls.

Adopting network segmentation has an immediate benefit – it allows for isolating or aggregating vulnerable assets into a safe zone that is more easily kept away from zero day attacks and other potentially devastating cyber threats. In some cases, such assets play an important role in the production line, so taking them off the grid is not allowed even when there is risk exposure. A network segmented with OT zero trust-based policies prevents attackers from traveling within your network and puts a stop to the spread of malware. Network segmentation sets up network architecture with availability, risk mitigation, and malware containment in mind from the ground floor.

Reinforce

Shield assets at a network level to secure vulnerabilities in legacy and other unpatched assets without interrupting their work.

Updating assets depends on a lot of factors. Is the patch available? If it's available, is it compatible? Does the OT environment allow for the asset to be patched? Asset status and patch status are major obstacles to the maintenance process. For assets that will eventually receive an official patch, they can use asset shielding to reduce risk until it's the right time for an update and a vendor-supplied patch has been released and tested.

For technicians, this makes mean time to patch (MTTP) much less of a concern and makes it possible to prioritize both productivity and security at the same time. Through asset shielding, one can secure assets regardless of whether or not their creator has released a security update. Unpatchable and legacy assets can be secured without changes to their configurations. The OT-native IPSes and firewalls that make this kind of asset-centric cyber defense possible have rule sets specifically designed to repel attacks without forcing endpoints to conduct an update, meaning no system reboots and no production downtime. Engineers can keep assets operational and secure while they prepare the patch for deployment during a scheduled maintenance window.

Deploying OT Zero Trust

When cybersecurity originated, it was based on the fact that in IT environments bad actors conduct attacks by stealing trust from legitimate users. Conversely, in OT environments bad actors conduct attacks by stealing trust from legitimate assets. This allows them to send the commands that break machines, install hacker tools and malware on assets, or steal sensitive information. OT zero trust defeats these strategies by zeroing out all unnecessary trust from the equation.

In 2020 and 2021, many major cyberattacks were the result of causes such as a trusted technician bringing an infected device into a worksite, an unlocked endpoint that ran a malicious application, or an attacker moving through and compromising the network. By minimizing trust on the network with defenses that are OT-native, cyber incidents like these are preventable and the cyber defenses that prevent them are manageable.

Phase 1

Inspect inbound devices

Don't trust inbound devices. Set up security checkpoints to inspect them. In the 2020 SolarWinds incident, hackers dropped a malware bomb inside a vendor software update and then sat back and let the vendor distribute the malware throughout their supply chain to all their customers. For a company high on the supply chain, this causes a catastrophe as every link of the chain gets hit.

OT zero trust relies on a fast handheld malware scanner that can verify that every inbound device is clean, protecting work sites against supply chain attacks. This is the only way to catch malware that attackers have hidden inside devices before they left the factory. The scanning process can be integrated with the asset inventory process to ensure scan completion, ease audits, and make it easier to confirm vulnerability status of each asset.

When attackers take advantage of a trusted employee's ability to introduce threats to the operational network, they create what's known as 'insider threat.' Eliminate the possibility of insider threat by putting a procedure into place for scanning devices as they arrive on-site.

Insider threat and supply chain attacks have been increasingly common factors in cyber attacks over the last few years – prevent both by verifying that every device that connects to assets or the network is malware-free. In the future, it's likely that vendors will need to follow a more standardized and even regulated approach to guaranteeing products as malware-free, but until then every OT organization must take this matter into their own hands: every work site needs a check point equipped to quickly conduct and log the results of malware scanning and removal.

Phase 2

Lock down endpoints

Securing endpoints means having a solution attuned to each endpoint's needs – this can be challenging due to operational conditions, such as assets from many different vendors and of many different lifespans working together. While legacy endpoints can be extremely sensitive and are often found handling limited, fixed tasks with limited resources, modernized endpoints have more resources to perform more tasks flexibly. Each endpoint needs defenses that can accommodate its needs and purpose without disrupting availability.

For fixed-use and legacy systems, the nature of the work is more predictable, so deploy a straightforward trust list that only allows work-necessary applications to run. Malicious files can't execute or spread themselves, and even updates from vendors can't run until given approval. TXOne's specialists think this is an ideal solution because its straightforward nature is extremely resource-friendly.

Modernized endpoints that perform a wider variety of tasks have historically been difficult to secure without interrupting their availability. The key to solving this problem is application filtering based on a library of verified OT applications and licenses. This allows next-generation antivirus to exclude trusted applications from scans and give their work priority while resource-friendly malware scans detect and stop known threats. Unknown threats can potentially be an issue on modernized systems, so next-generation antivirus includes machine learning that can identify and then report or stop suspicious behaviors.

Supply Chain Attacks

One way attackers trigger supply chain attacks is by disguising a malware payload as an update. To initiate the Kaseya supply chain attack, attackers took advantage of a "zero-day authentication bypass" vulnerability (CVE-2021-30116) to inject REvil ransomware into customers' systems disguised as an update and ultimately affect nearly 1500 organizations worldwide.[4] The REvil group claimed to have infected over a million devices when they demanded a ransom of USD $70 million.[5]

[4] Shaun Nichols, "Kaseya ransomware attacks: What we know so far", Tech Target, 6 Jul 2021
[5] Lance Whitney, "Kaseya supply chain attack impacts more than 1,000 companies" Tech Republic, Jul 6 2021

Phase 3

Secure the network

Your network is like a busy freeway. Most traffic obeys the rules, but malicious messages also go with the flow looking for an on-ramp through a security hole in your system. When you seal unnecessary roads and on-ramps in your network and make specific rules for traffic, attackers will find it much more difficult to infect or access assets. By making productivity-based rules about which assets can communicate with each other, security experts segment the network into smaller zones. After these zones are created, give zones customized security policies based on the specific needs of each group of assets.

ICS-CERT advisories bring new vulnerabilities into the public eye every day, along with mitigations that technicians must sometimes carry out quickly due to high risk. Managing the security of these vulnerabilities without disrupting the operation can be extremely challenging, and sometimes they are even found in mission-critical legacy devices that no longer receive updates. Create a network-based asset shield around the device's vulnerabilities so that attackers can no longer exploit them – asset shielding acts in place of a security update, even for otherwise unpatchable legacy systems.

The Shamoon Shutdown

In 2012, 2016, and 2018, the malware known as 'Shamoon' was used to attack oil and gas companies in the UAE.[6] Attackers calling themselves "Cutting Sword of Justice" attacked Saudi Aramco's network, wiping out more than 30,000 computers by taking advantage of the organizations' "flat" networks. Flat, unsegmented networks are much easier for attackers to exploit and take over, such that one security hole allowed the virus to spread to all the computers where it erased data and destroyed their ability to reboot. While industrial systems were not directly attacked, oil production stopped because the internet was down and there was no way to monitor the equipment.

Malware such as this is unable to infect endpoints or spread itself within OT network secured with the OT zero trust approach. Application trust lists block any unapproved applications from running while network segmentation only allows communication with or between assets on the simplified, limited basis necessary to do their work.

[6] Darknet Diaries, Ep. 30: Shamoon, 22 Jan 2019

OT Zero Trust Risk Assessment

Reduce trust to reduce risk, trust is risk – the problem is that reducing trust increases effort for availability.

Which devices are most vulnerable? What are the attack surfaces? How can we mitigate them? The goal of the risk assessment is to predict security holes and fill them. The NIST Guide to Industrial Control Systems (ICS) Security recommends building a risk framework then continually assessing risks, responding to threats, and monitoring vulnerabilities.

Traditional risk assessments have two vectors: likelihood that the risk event will happen and the impact if it does. IT cybersecurity assessments revolve around the CIA triad: confidentiality, integrity, and availability. These are not sufficient for OT, which instead is based on an "AIC" triad - availability first, followed by integrity, with confidentiality last.

Machines work all day, every day, 24 hours per day. Productivity is key. Safety is critical. For example, the NERC-CIP defines a "bright-line" criteria for categorizing bulk electrical systems based on the impact if they were rendered unavailable for more than 15 minutes. Detectability is important. If you can't see the threat then you cannot respond.

The environment plays an important role. Inclement weather is a risk to outdoor operations. Finally, risk thresholds are unique to every company. For each risk you identify you decide what response meets your comfort level. Responses generally include avoiding, transferring, sharing, mitigating, or accepting the risk.

Risks specific to ICS

Consider what damage could occur if your sensors or actuators were hijacked. Think about what would happen if an attack propagated to connected systems. If your digital controllers stop functioning, what happens to your non-digital assets? These risks are specific to ICS:

- Safety of your team and your community
- Physical impact to property and the environment
- Consequences for non-digital control components

Cyber Attacks Unique to OT and ICS Environments

Researchers at the National Institute of Standards and Technology (NIST) investigated the most common types of cyber threat events.[7] Here we've added information to show how OT zero trust addresses each threat event.

Threat Event	What is it?	How can we secure systems against it?
Denial of Control Action	Control systems are disrupted by delaying or blocking the flow of data and creating bottlenecks.	Network trust lists block commands sent from the wrong location and restrict the privileges of unwelcome guests.
Control Devices Reprogrammed	Unauthorized changes to programmed instructions in assets such as PLCs, RTUs, DCS, or SCADA controllers that change alarm thresholds, change equipment behavior, or prematurely shut down systems. Any of these could cause spills, fires, equipment damage, or other harm to workers and the environment.	Trust lists disallow unprivileged users from changing configurations on the network, and the status of each asset can be confirmed from a centralized console on a single pane of glass. With trusts lists in place on endpoints, configurations and data can only be altered by trust listed users or applications.
Control Logic Manipulation	Malicious changes to software or configuration settings could produce disruptive results.	
Safety Systems Modified	Safety systems could be turned off or programmed to take incorrect actions that damage or destroy systems and threaten workers or the environment.	

Threat Event	What is it?	How can we secure systems against it?
Spoofed System Status Information	False data sent to control system operators could disguise attacks.	Network segmentation allows for better inspection of data on the network, making the out-of-place characteristics of false data easier to detect so that the malicious command is rejected and the source of the breach can be tracked ("Why is my robot arm in maintenance mode uploading a file to my HMI?").
Malware injected into control systems	Ransomware, worms, and other malware can disrupt operations in various ways including denying control of assets or destroying equipment.	Application trust lists come in two forms to cover the needs of legacy and modernized endpoints. For fixed-use or legacy endpoints, a straightforward trust list that disallows all unlisted applications is enough to protect operations without causing crashes or delays by taking up too much of the computer's resources. For modernized endpoints that use more resources to more flexibly perform a wider variety of tasks, protection from malware is based on intelligently identifying which operations-related applications are present on the device, then excluding those applications from scans, empowering them to make changes, and giving those applications priority for resources. Meanwhile, machine learning is able to detect suspicious behavior, providing the system with protection from unknown threats.

[7] Keith Stouffer, Victoria Pillitteri, Suzanne Lightman, Marshall Abrams, Adam Hahn, "NIST Special Publication 800-82 Revision 2, Guide to Industrial Control Systems (ICS) Security", National Institute of Standards and Technology, May 2015

The Business Case for OT Zero Trust

Sleep well at night knowing that your hard work is protected. Look at the benefits of improved system reliability and availability. Prioritize the potential costs and impacts of a cyber incident. Consider the cost to deploy, run, monitor, review, maintain, and improve your cybersecurity. Be aware of the benefits and the potential consequences.

Benefits

- Improve control system safety, reliability and availability
- Improve employee morale, loyalty, and retention
- Reduce community concerns
- Increase investor confidence
- Reduce legal liabilities
- Meet regulatory requirements
- Enhance corporate image and reputation
- Ease insurance coverage and cost
- Improve investor and bank relations

Potential Consequences

- Reduction or loss of production at one site or multiple sites simultaneously
- Endangerment of human lives
- Damage to equipment
- Environmental damage including release, diversion, or theft of hazardous materials
- Violation of regulatory requirements
- Product contamination
- Criminal or civil legal liabilities
- Loss of proprietary or confidential information
- Loss of brand image or customer confidence

txOne networks

txone.com