Supply Chains in the Crosshairs: Five Ways to Help Defend Against the Next Big Cyberattack

If You Think You've Seen the Last of the SolarWinds-Type Attack, Think Again

In less than a year, organizations worldwide have found themselves victims of multiple high-profile attacks, including Microsoft Exchange Server exploits, Kaseya VSA ransomware attacks, and countless other ransomware attacks. With cryptocurrencies at record prices, threat actors like REvil (implicated in the Kaseya VSA attack) are emboldened to launch ransomware campaigns to attack and infect a broader range of secondary targets. A continuous string of headline-generating cyberattacks suggest that organizations need to do everything they can to safeguard their networks sooner rather than later.

This paper offers recommendations and provides guidance to help reduce vulnerabilities and threat exposure across enterprise environments while protecting adjacent supply chains.

Introduction

As a result of digital transformations, we've seen companies embrace countless technology and software vendors, building out their supply chain to increase productivity and leverage emerging technologies to support their growth and gain efficiency at scale. Threat actors are increasingly targeting the software supply chain to exploit the trust companies have in their vendors and software providers because they have access to the enterprise network. With this access, attacks can even include compromising a vendor's digital certificate signing process to bypass the victim's defenses.

The threat of a software supply chain attack is compounded by the current challenges faced by security teams, including an endless stream of disjointed, unconnected security alerts that lack the necessary context needed to determine what action to take. It is imperative that companies adopt detection and response tools with integrated machine-learning capabilities to be able to carefully watch for malicious behavior, regardless of whether the source is from a signed or trusted supplier.

Furthermore, implementing a Zero Trust framework to inform an organization's network architecture can help successfully migrate from traditional perimeter-based security models to one based on a continual verification of trust.

Superb Tradecraft from the Adversary Requires a Reckoning from Our Industry

In March 2020, SolarWinds sent out a regular software update that unknowingly contained hacked code to its 33,000 Orion® customers, with approximately 18,000 installing it, creating further compromise that affected a wide swath of private and public entities, including parts of the Pentagon, the State Department, the Department of Energy, the Department of Homeland Security, the Treasury, the National Nuclear Security Administration, as well as several cybersecurity solution providers. Needless to say, the repercussions are still unfolding as companies and federal sectors conduct post-mortem assessments.

These recent attacks leveraging SolarWinds demonstrated the potential magnitude of a sophisticated supply chain attack. Given the extent to which organizations use SolarWinds software and servers to manage their networks, systems, and IT infrastructure, the threat actors had unparalleled access. By all accounts, it's estimated the attack went undetected for approximately 10 months.

SolarWinds announced that an adversary had successfully integrated malicious code into a properly signed SolarWinds software update. Once the update was applied, the malicious code was extraordinarily stealthy. In this attack, the malware, known as SUNBURST, would sit dormant for about two weeks and then ping a subdomain of avsvmcloud[.]com—the command-and-control (C2) server for the backdoor—from where it would obtain further instructions and additional payloads to execute. Some of the largest and most advanced companies and government agencies in the United States identified SUNBURST on their production servers, jeopardizing and exposing confidential data.

One aspect that made the hack interesting is how the attacker stood up the infrastructure over such a long period of time, compromising adjacent organizations with most of the attacker's traffic appearing as normal. That said, 72 out of 72 malware detection vendors indicated avsvmcloud[.]com domain as benign.

Given how central SolarWinds is to IT automation, they were able to leverage the strategic positioning and connectedness of the solution to gain access to potentially anything else in the customers' environment by compromising the SAML infrastructure and leveraging compromised SAML certificates to move laterally and access sensitive data in email. One method of detection known to be successful was based on behavioral abnormalities or techniques common with known malware and exploit kits, which is exactly what Cortex XDR did.

"Recently, we experienced an attempt to download Cobalt Strike on one of our IT SolarWinds servers. Cortex XDR instantly blocked the attempt with our Behavioral Threat Protection capability and our SOC isolated the server, investigated the incident, and secured our infrastructure. We also deployed a set of IOCs to our customer-facing Palo Alto Networks products as a result of this."

– Nikesh Arora, CEO, Palo Alto Networks

On December 17, 2020, Palo Alto Networks CEO Nikesh Arora disclosed that their internal deployment of Cortex XDR successfully blocked a DNS request from their SolarWinds Orion server thanks to the Behavioral Threat Protection capability in Cortex XDR, allowing the Palo Alto Networks SOC team to isolate the server and initiate an investigation. They concluded that because of Cortex XDR, the attack was unsuccessful, no data was compromised, and the security of their infrastructure was secure.

This critical integration of threat prevention, detection and response, and Cortex XDR's use of machine learning and AI to automatically integrate endpoint, network, and cloud data allowed them to stop this unprecedented attack.

As supply chain attacks continue to grow in number and scale, cybersecurity has clearly entered a new era of extremely well-funded, focused, and disciplined nation-state threat actors whose objective is to gain access to their targets, maintain persistence over time, and accomplish a variety of goals, including data theft. Thwarting these types of sophisticated campaigns will require new approaches and technologies to stay ahead of threat actors, who, every day, are growing bolder in their determination to launch even more advanced attacks—all the while leveraging the scale of the cloud and the power of automation to wage attacks that legacy technology and risk management practices don't stand a chance against.

Preparing for the Next SolarWinds-Type Attack: Companies Need to Take Action Now

It is important to note that even before this attack, it was already clear that a large percentage of security operations centers depend too much on human intervention and a range of siloed security products. Analysts were already under stress, struggling to stay on top of the alerts produced by the multitude of products they have deployed. Most are still relying on manual review of alerts that are indicators of suspected attacks, with tens of thousands of alerts each week pouring in from dozens of security products, saying, "look at me." Manual investigation and response practices often require an inordinate amount of time gathering context (what do I know about the endpoint, user, and time of suspected attack activities) and conducting incident analysis (what are related events to the hostname and IP address, the traffic, domain, application being used, etc.), all of which result in incomplete assessments and a lack of security efficacy.

> A 2019 survey of CISOs reported that "over 41% see more than 10,000 and that some claim to see more than 500,000 alerts daily."
>
> The same report noted that respondents revealed only 24% of investigated alerts were considered legitimate, down from 34% in 2018. The report also observed a substantial drop in the number of legitimate alerts that were in fact remediated—from 51% in 2018 to 43% in 2019.[1]

And with many organizations still operating legacy security solutions, including legacy anti-virus (AV), endpoint detection and response, and other security technologies, the risks are even greater. With an overwhelming amount of mostly low-quality data coming in, many analysts respond by de-tuning sensors or simply ignoring some of the alerts, which of course, raises the risk level. Many alerts, without context, are dismissed as false positives because they are not sufficient to warrant an investigation. Yet, when placed into context with further observation from additional data sources, the data may be the key to understanding truly malicious intent from otherwise benign activity.

The future of efficient security operations lies in replacing legacy, siloed security tools with those that are well integrated and provide robust analytics, machine learning, and automated detection to accelerate response times while increasing accuracy. When the right tools are integrated with data that is relevant and consolidated, organizations can reduce response times and get a holistic view of incidents with rich details to better inform investigations.

So, what is the path forward? What should companies and organizations do to prepare for the next supply chain attack along with the ever-evolving threats that are sure to come in the near future?

Consider the next five actions and how each can help influence activities, from conducting a cyber risk assessment to creating a broader security operations strategy.

1. *Anticipating the Unknowns*, Cisco, March 2019, https://ebooks.cisco.com/story/anticipating-unknowns/page/6/6.

1. Know Your Attack Surface

With employees, partners, and vendors working outside the enterprise network perimeter, organizations face a greater risk to internal systems and data being exposed and attacked. Options include penetration testing, vulnerability scanning, and an emerging technology called attack surface management.

An Overview of Attack Surface Management (ASM)

Defined by SANS Technology Institute, attack surface management is:

> "An emerging category of solutions that aims to help organizations address this challenge by providing an external perspective of an organization's attack surface … An organization's attack surface is made up of all internet-accessible hardware, software, SaaS and cloud assets that are discoverable by an attacker. In short, your attack surface is any external asset that an adversary could discover, attack and use to gain a foothold into your environment."[2]

SANS lists some common use cases for adoption of an ASM solution, including:

- Identification of external gaps in visibility
- Discovery of unknown assets and shadow IT
- Attack surface risk management
- Risk-based vulnerability prioritization
- Assessment of M&A and subsidiary risk

Whether one chooses to deploy ASM solutions or perform pen testing or vulnerability scanning, what is clear is the need to identify both product and operational requirements to determine best fit, including functionality, feature/s, capability, and evaluation criteria.

2. Prevent Everything You Can

With a marketplace full of numerous security solutions, cyberattacks are still occurring and getting more sophisticated (and seemingly well-funded) in their level of complexity.

That said, numerous notable attacks have been found to utilize fairly common attack vectors such as embedded malware, phishing emails, and privilege escalation. Therefore, utilize technology and best practices to prevent everything and focus your attention on what matters.

Prevention Basics

- Invest in securing your endpoints. Endpoint Protection Platforms (EPPs) use multiple techniques for prevention, including static analysis to evaluate potential malware based on file inspection, heuristics rules to block exploits, and behavioral analysis to evaluate file maliciousness based on the functions they perform.
- Whether deployed in development, QA, or a "live" production environment, make sure you have an integrated security solution that allows you to send vulnerability data to tools that you are already using, such as bug trackers for fast remediation.
- Use complex passwords. Simply having a password that is at least 10 characters long will help harden your security. Even if not prompted, change your password two to three times a year.
- Keep security software up to date. Install patches regularly.
- Restrict network access to trusted hosts and networks. Only allow internet access to required network services. And if absolutely necessary, don't deploy systems that can be directly accessed from the internet. If remote access is required, consider using VPN, SSH, or other secure access methods.
- Prevent phishing and other email-borne infections by training employees on current best practices and policies, such as deleting suspicious attachments.
- Use multi-factor authentication (MFA) if possible.
- Configure your spam filters for maximum coverage.

2. Pierre Lidome, "The SANS Guide to Evaluating Attack Surface Management," SANS Institute, October 26, 2020, https://www.sans.org/reading-room/whitepapers/analyst/guide-evaluating-attack-surface-management-39905.

Many breaches happen due to a combination of human error, unpatched systems, and the sheer persistence of modern threat actors employing their own versions of digital transformation to their advantage.

If you suspect you've been compromised:

- If only a few systems are infected, immediately disconnect them (physically) from your internal network to prevent and contain the infection. If this cannot be accomplished in a timely manner or more than a few systems are infected, and you have not implemented strong firewall egress filtering and proxy servers, immediately block ALL outbound traffic to external networks.
- Implement filters on internal routers, firewalls, and other networking equipment as appropriate to isolate infected segments and to monitor network traffic to ensure internal containment or identify how this infection is spreading and which hosts are infected.
- Monitor all network traffic to address possible multifaceted attacks.
- Review appropriate log files to attempt to identify the first system infected and what the attack vector was if possible.
- It is vital to determine if any of the infected systems successfully connected to any site on the internet and what information, if any, was exposed.

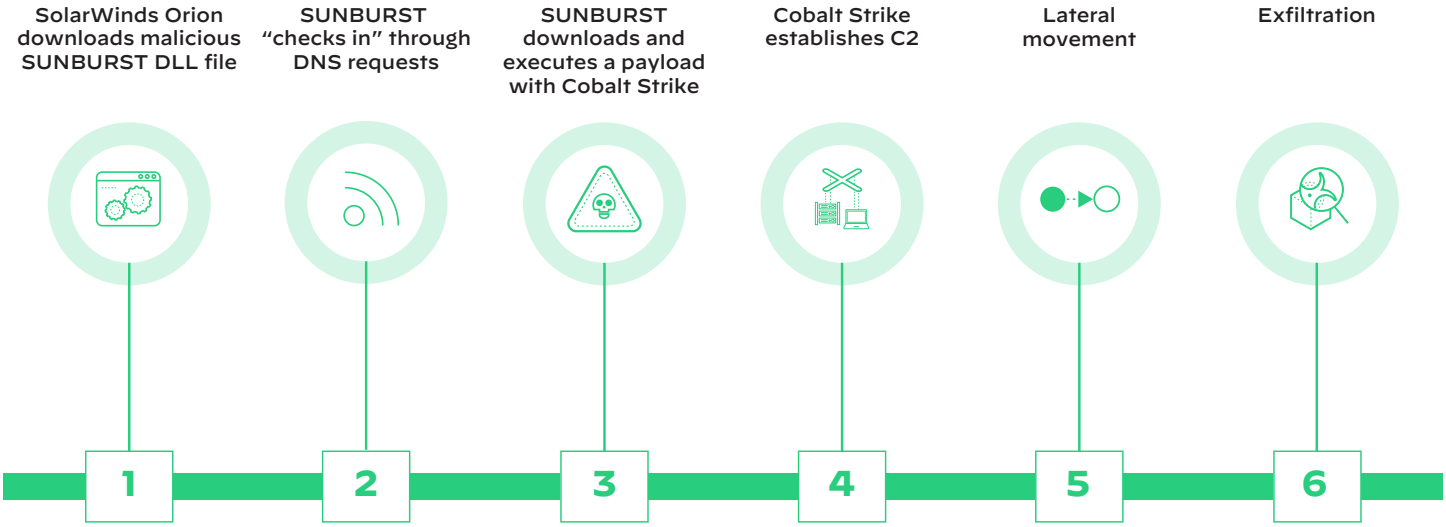

Figure 1: SolarStorm attack was blocked at Stage 3 by Cortex XDR

With Cortex XDR, you save the time and cost of building out your own global endpoint security infrastructure. This simplified deployment, which requires no server licenses, databases, or other infrastructure to get started, enables organizations to quickly protect their endpoints.

3. Gain Maximum Visibility

Enterprise visibility, with a unified approach to data throughout a supply chain, can provide a holistic view of applications and infrastructure. Visibility can also include the telemetry from endpoints, networks, and cloud environments. Moreover, it must correlate those data sources to understand how various events are linked and when a certain behavior is, or isn't, suspicious based on context.

Telemetry data and forensics combined with machine learning can provide a clear view into an attack chain for rich analysis during triage and verification of an alert—simplifying and speeding up investigation and response.

In the case of the SolarWinds attack, even though the unauthorized malicious code was already built into the supplier's software, Palo Alto Networks own Cortex XDR was able to block an attempt to download Cobalt Strike on one of their IT SolarWinds servers because of the Behavioral Threat Protections (see Figure 1).

For advanced attacks like the one on SolarWinds, having full visibility can enable organizations to detect and stop all stages of the attack lifecycle (even if the host has already been compromised). If an attack is so advanced that it bypasses your preventative measures, you need to be able to detect the post-intrusion activity the attacker needs to achieve their objectives.

4. Act Quickly

While hackers had gained entry into SolarWinds' system around January 2019, access to their infrastructure had reportedly been available for sale on the dark web in the Exploit Cybercrime forum on October 13, 2017, underscoring the financial incentive to launch APT campaigns.

By March 2021, both Microsoft and FireEye reported new indicators of compromise (IoCs), including backdoor and other malware implants, to establish sustained access to affected networks. According to Microsoft, they:

> "... discovered these new attacker tools and capabilities in some compromised customer networks and observed them to be in use from August to September 2020. Further analysis has revealed these may have been on compromised systems as early as June 2020. These tools are new pieces of malware that are unique to this actor. They are tailor-made for specific networks and are assessed to be introduced after the actor has gained access through compromised credentials or the SolarWinds binary and after moving laterally with TEARDROP and other hands-on-keyboard actions."[3]

Once an attacker gains entry, evading initial detection and maintaining persistence becomes essential to a campaign's "success." As the named threat actor behind the SolarWinds breach, SolarStorm had taken advantage of using stolen credentials to gain access to cloud services, as well as exploiting compromised identities to gain and maintain access to networks via VPNs and remote access tools.

To that end, the urgent necessity to reduce dwell time (or the breach detection gap) and subsequent lateral movement is imperative to contain, remove, and recover from an attack. In addition to potential reputational damage, fines for non-compliance, and loss of critical business data, the longer a breach goes undetected and uncontained, the greater the financial impact.

In their report "Quantifying the Value of Time in Cyber-Threat Detection and Response," Aberdeen Group noted that when dwell time was confined to seven days, the impact is reduced by 77%, and if shortened to one day, business impact can be reduced by 96% (see Figure 2).[4]

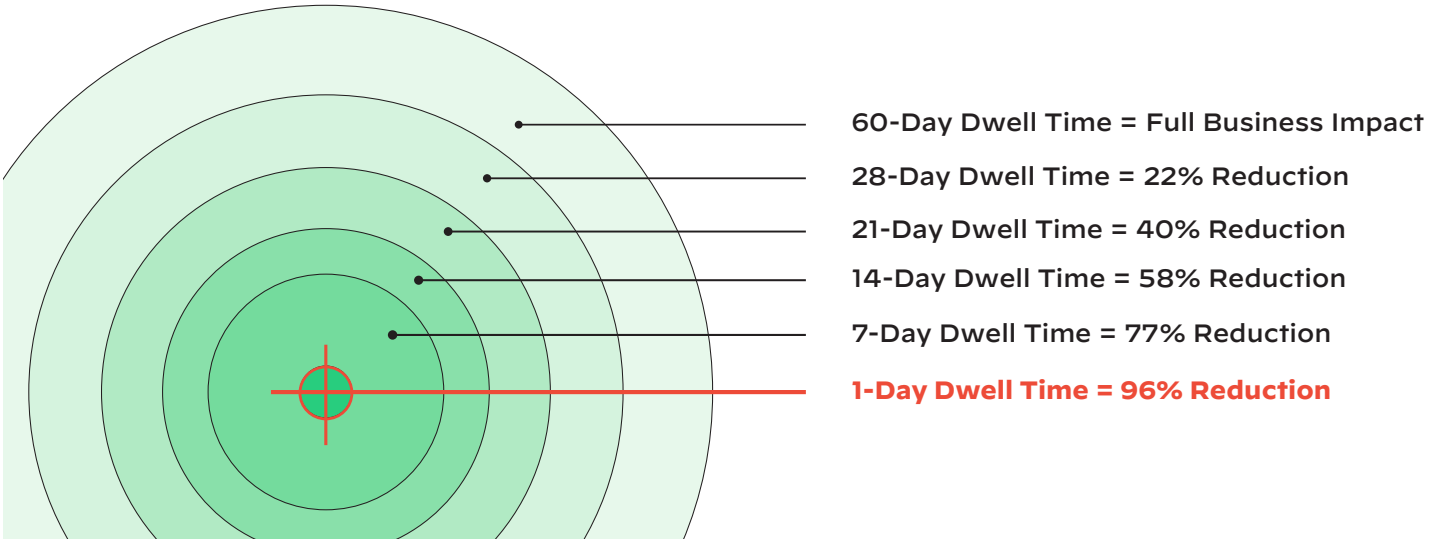

Figure 2: "Quantifying the Value of Time in Cyber-Threat Detection and Response," Aberdeen Group, February 2016

3. Ramin Nafisi et al., "GoldMax, GoldFinger, and Sibot: Analyzing NOBELIUM's Layered Persistence," March 4, 2021, https://www.microsoft.com/security/blog/2021/03/04/goldmax-goldfinder-sibot-analyzing-nobelium-malware/4.
4. Aberdeen Group, "Quantifying the Value of Time in Cyber-Threat Detection and Response," February 2016.

Beyond the importance of responding immediately to attacks is to consider how fast and how often threat actors are scanning and locating potential threat vectors. Advancements in scanning technologies allow attackers to locate attack vectors quickly and easily, revealing abandoned, rogue, or misconfigured assets that can become backdoors for compromise.

In their 2021 Cortex Xpanse Attack Surface Threat Report: Lessons in Attack Surface Management from Leading Global Enterprises, Palo Alto Networks outlined key findings from their research of the public-facing internet attack surfaces of some of the world's largest businesses. From January to March, their team monitored scans of 50 million IP addresses associated with 50 global enterprises to understand how quickly adversaries can identify vulnerable systems for fast exploitation.

Some interesting takeaways from that research include:

- **Adversaries work nonstop**: In a never-ending game of "cat and mouse," threat actors were found to conduct a new scan once every hour, whereas global enterprises can take weeks.
- **Adversaries jump on new vulnerabilities**: Attackers began scanning within 15 minutes following announcements of new common vulnerabilities and exposures (CVEs) released between January and March and launched scans within five minutes of the Microsoft Exchange Server zero-days security update.
- **Cloud comprised the most critical security concerns**: Cloud footprints were responsible for 79% of the most critical security issues found in global enterprises, reiterating the inherent risk of cloud-hosted/based services, compared to 21% for on-premises.

MTTD (mean time to detect) is the amount of time it takes a company to identify a potential security incident.

MTTF (mean time to failure) is how long a defective system can run until it shuts down.

MTTR (mean time to respond) is how long it takes a team to get a grip on, remediate, or eliminate a threat after it has been identified.

MTBF (mean time between failures) reflects the reliability and availability of a system. It is used to evaluate the system's performance under predetermined conditions for a set amount of time.

5. Bake in Zero Trust

The concept of Zero Trust has been around for a while, yet it garnered acceptance in 2009 when Forrester Research formalized the "Zero Trust Model" of cybersecurity.[5] In short, it is both an architectural model for networks and a framework for setting security policies.

Zero Trust relies on strict verification and validation for every person, device, or entity attempting to access network resources. The primary goal is to prevent successful breaches or corruption of data, applications, and business-critical systems from attacks and exploits.

The principles in Zero Trust are designed to reduce exposure and unauthorized access across the threat landscape. They've been thoughtfully developed to address the security of critical applications and sensitive data across an enterprise organization. These principles can easily become a part of any security strategy. Some of them include:

- **Policy of Least Privilege (PoLP)**: A policy in which end users are given the minimum amount of access they need to carry out their jobs. This helps reduce pathways and exposures to malware, attackers, and the chances of data exfiltration.
- **Microsegmentation**: A network is divided into separate segments or "secure zones" in data centers or in cloud deployments that require different access credentials to help isolate workloads. This also helps limit lateral (or east-west) movement in internal networks if breached.
- **Multi-factor authentication (MFA)**: A security protocol that requires individuals to be authenticated with more than one required security procedure. Typically, this is a combination of things one knows (e.g., passwords or a PIN), things one has, such as a fob, badge, etc., and physical markers such as biometrics, voice recognition, or fingerprints.

5. John Kindervag et al., "No More Chewy Centers: The Zero Trust Model of Information Security," March 23, 2016, https://www.forrester.com/report/No+More+Chewy+Centers+The+Zero+Trust+Model+Of+Information+Security/-/E-RES56682?objectid=RES56682.

Recommended steps to help mature Zero Trust capabilities include:

- Monitor all activity and collect all data—not just suspicious events.
- Detect anomalous behavior with analytics and machine learning.
- Detect and block malicious behavior on the endpoint.
- Segment access with host firewall.
- Monitor and restrict access to unauthorized USB devices with device control. Users cannot connect any storage device to the machine, except for authorized cases for authorized devices for a limited time only.
- Block malicious remote hosts.
- Orchestrate security controls. Automation and orchestration can help identify gaps in Zero Trust architectures and automatically resolve them or trigger workflows to help analysts remediate them. According to the Forrester report The Zero Trust eXtended (ZTX) Ecosystem, "Avoid solutions that function in isolation and opt for those that integrate to form an ecosystem to aid better visibility and control across the ecosystem and robust orchestration of security defenses."

Some additional guidance to help inform a Zero Trust approach can be found here:

- National Institute of Standards and Technology: Zero Trust Architecture: NIST Publishes SP 800-207
- The National Cyber Security Centre: Zero trust principles - beta release
- The National Security Agency: Embracing a Zero Trust Security Model
- Forrester Research: Five Steps To A Zero Trust Network* where they advise the following to implement a successful Zero Trust roadmap:
    1. Identify your sensitive data.
    2. Map the flows for sensitive data.
    3. Architect a Zero Trust microperimeter.
    4. Monitor the Zero Trust environment, in detail, with security analytics.
    5. Embrace security automation and orchestration.

** The Forrester report is available to Forrester subscribers and for purchase.*

A Not-so-Secret Weapon: Cortex XDR

Fortunately, there is an efficient way to apply the above principles to any security strategy. Modern extended detection and response (XDR) platforms collect and automatically correlate data across multiple sources—endpoint, server, cloud workloads, and network—so threats can be detected faster and security analysts can improve investigation and response times.

Palo Alto Networks Cortex® XDR™ is the industry's first extended detection and response platform that natively integrates endpoint, network, and cloud data to stop sophisticated threats, enabling the user to instantly eliminate network, endpoint, and cloud threats from one console.

Only Cortex XDR can help security teams:

- **Automatically detect stealthy attacks** across threat vectors using behavioral analytics across network, endpoint, and cloud data. Cortex XDR's core technology is focused on learning the behavior of each device and user and how it compares to other devices and users in the network of the individual organization. These behavioral profiles are used to detect deviations from past behavior, peer behavior, or the expected behavior of the entity.
- **Shorten investigations** by stitching security telemetry and alerts from multiple endpoints and additional data sources into a single incident that reveals the root cause. This visibility across data sources eliminates security blind spots, improves detection accuracy by identifying sophisticated attacks across data layers, and provides additional context to simplify investigations. By consolidating data, it also reduces the number of distinct detection and response products that customers need to manage.
- **Continually adapt defenses** by applying knowledge gained from investigations to prevent future threats. Your analysts can quickly stop the spread of malware, isolate endpoints, sweep across all endpoints to delete malware in real time, and even directly access endpoints and investigate threats without disrupting end users.

A Final Word

Attacks like the one on SolarWinds are a clarion call to cybersecurity professionals in all sectors and all roles to stay vigilant in protecting their networks from persistent and sophisticated threats. While the industry takes somewhat of a collective "breather," conducting post-mortem assessments and resolving any latent damage from the SolarWinds attack, organizations must realize the urgent requirement to do better—because the next big attack could hit tomorrow.

Attackers are already devising techniques that defy what we believe is possible today, searching for creative ways to launch stealthy campaigns undetected and unfettered from doing the most damage possible.

Now is the time to take advantage of the wealth of modern security solutions and best practices to harden your security stance. Incremental steps can be taken to ensure security teams move in the right direction, including gaining visibility into the attack surface, preventing as much as possible, and applying Zero Trust principles.

Perhaps most notably is considering next-generation solutions like XDR that are a logical evolution of proven tech like EDR. XDR provides the necessary visibility and control to adjacent business components via integrations that combine EDR data with other types of telemetry.

The ability of XDR to leverage security data streams toward continually improving machine-learning models—as well as facilitating coordinated detection and response—holds enormous promise for keeping threat actors at bay for some time to come.

Interested in learning more about new threat hunting skills or how Cortex XDR can defend your environment? Schedule a hands-on workshop with our experts and/or request a Cortex XDR demo.

Additional Resources for Understanding XDR

Download our e-book, The Essential Guide to XDR.

Download our datasheet on Cortex XDR.

Visit our Cyberpedia page, What is XDR?

For information on Round 3 of the MITRE ATT&CK® Evaluation download our e-book.

We are recognized by Forrester Research as a leader in The Forrester Wave™: Zero Trust eXtended Ecosystem Platform Providers, Q3 2020. Click on the link to download the report.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com