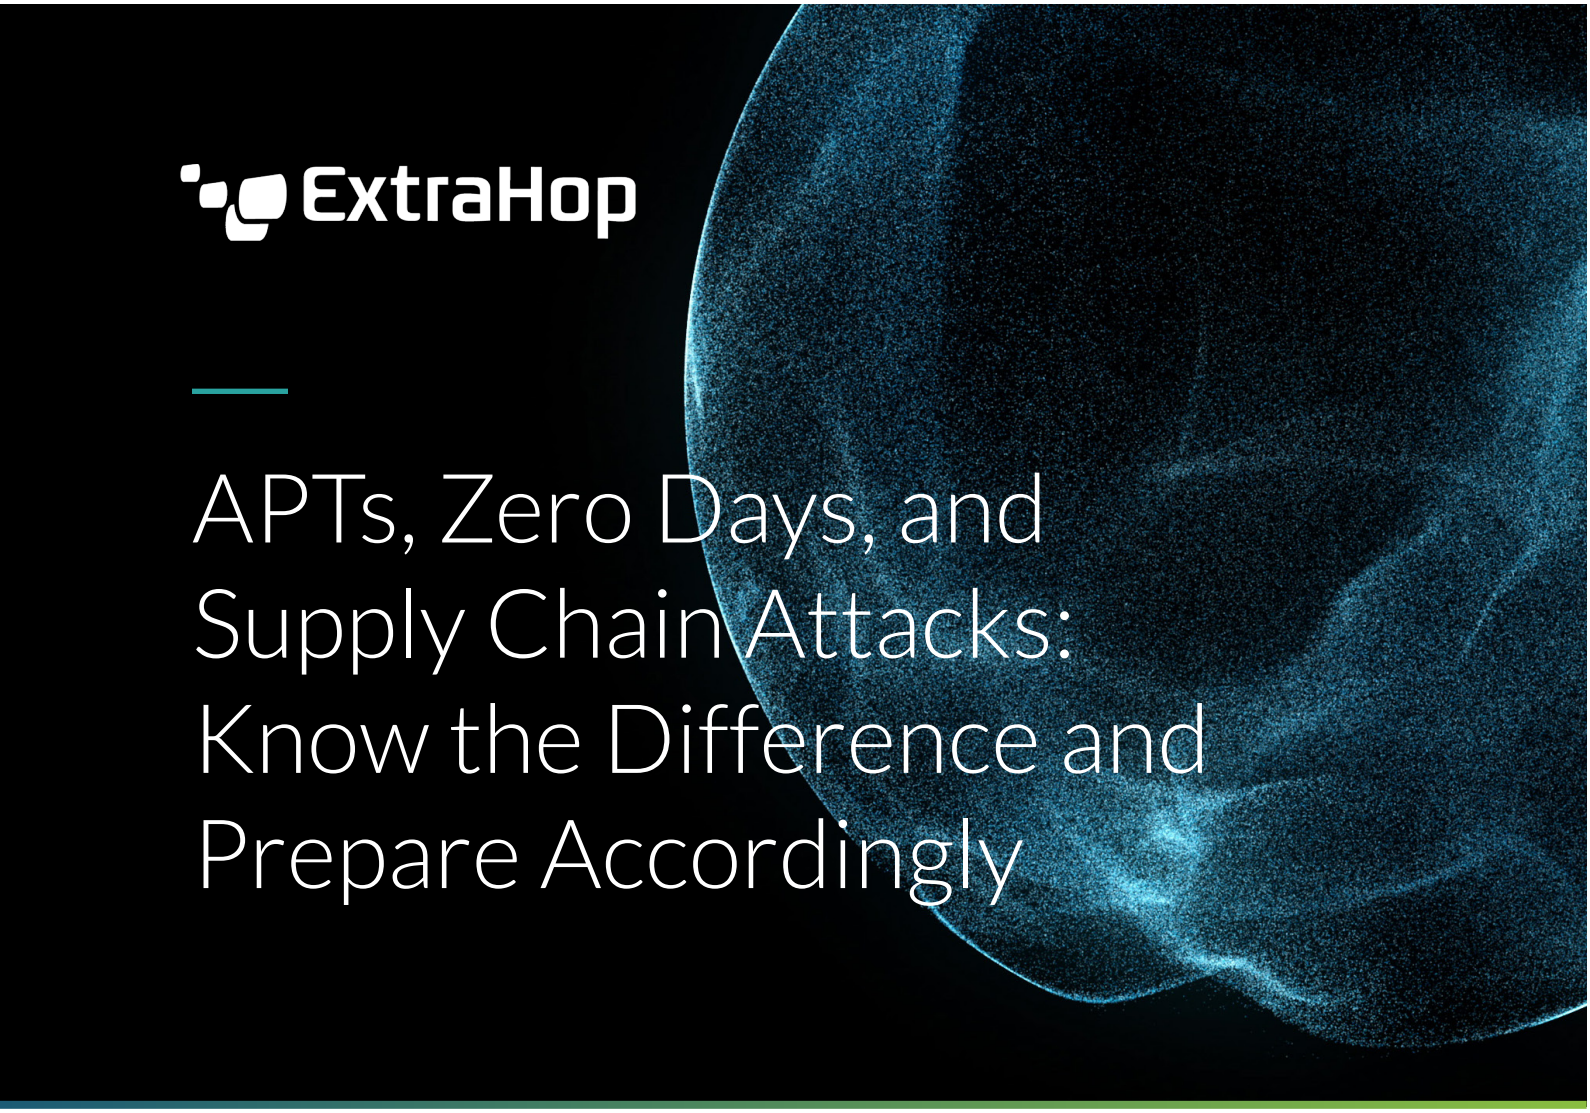

ExtraHop

APTs, Zero Days, and Supply Chain Attacks: Know the Difference and Prepare Accordingly

EXECUTIVE SUMMARY

Zero days. Software supply chain attacks. Advanced persistent threats. Cyber attacks perpetrated by nation-states and cybercriminals alike are on the rise, impacting organizations across geographies and industries. Increasingly, the catch-all term for these mechanisms of cyber warfare, espionage and crime is "advanced threats." But what, exactly, makes a threat advanced?

In this paper, we explore three manifestations of advanced threat activity: zero days, advanced persistent threats, and software supply chain compromise. These three manifestations are not mutually exclusive, often overlapping and enabling each other as nation-state adversaries and cybercriminals work to execute on their nefarious missions undetected. This paper examines how to classify these attacks, the interplay between them, and provides guidance on how to apply controls.

TABLE OF CONTENTS

INTRODUCTION

Cybersecurity issues receive plenty of media attention. Major data breaches involving the large-scale exposure of personal, financial, or health information grab headlines. Nation-state meddling that compromises elections and insider threats that expose state secrets have brought security to the fore of the public conscience.

But recently, the tone of cybersecurity coverage in mainstream news outlets has started to change. Major daily and business press publications and cable news channels are covering issues once reserved for security and technical trades. Outlets like The New Yorker and The Wall Street Journal are publishing articles on advanced persistent threats, zero day vulnerabilities (such as in Microsoft Exchange servers), and a combination of both attack types used in advanced software supply chain compromises like the SolarWinds SUNBURST attack.

Increasingly, the catch-all term for these mechanisms of cyber warfare, espionage and crime is "advanced threats." But what, exactly, makes a threat advanced?

In this paper, we explore three manifestations of advanced threat activity: zero days, advanced persistent threats, and software supply chain compromise. These three manifestations are not mutually exclusive, often overlapping and enabling each other as nation-state adversaries and cybercriminals work to execute on their nefarious missions undetected.

An Overview of Advanced Threat Types

Zero days have been around as long as software has. Due to the access they provide, they are prized among cybercriminals and nation-state hackers alike. Zero days are considered an advanced form of cyber attack because finding and exploiting these unknown vulnerabilities typically requires a high degree of skill and research. As such, zero days are a prized commodity, and often sell for five and six figure sums on the black market. Zero days, by nature of them being unknown to responders, are usually considered advanced. And attackers increasingly use other advanced attack methods to deploy zero days, such as with the attacks against Microsoft Exchange servers reported in March 2021. The attackers started out quietly searching for and exploiting four zero day vulnerabilities they had discovered to use against those systems because they wanted to take advantage of the vulnerabilities as long as possible. Then, once vulnerabilities were discovered and reported, the attackers unleashed a loud, broad-scale search and attack sequence to take advantage of the window of opportunity between the time the zero day was reported, and the time impacted systems were patched against them. In just a few days, attackers had scanned and infected an estimated 250,000 Exchange servers, according to the WSJ article cited above.

Advanced Persistent Threats (APTs) combine vulnerabilities, social engineering, and known malware families.

Advanced Persistent Threats (APTs) are another form of advanced attacks which combine vulnerabilities, social engineering, and known malware families to infiltrate systems, escalate privileges, and exfiltrate information over long (persistent) periods of time. The goal is to dwell within an environment for as long as possible and gather as much information as they can on the target (and to exfiltrate as much valuable data as they can) before being discovered.

An early example of an APT is Stuxnet, a malware variant created by the NSA in 2011 supposedly to target Iran's nuclear program. A more recent example is the SolarWinds SUNBURST attack, which used multiple advanced methods to hide during the course of deploying the initial trojan and setting up and maintaining command and control channels. As a result, nation-state adversaries lingered in SolarWinds network for months, taking slow, deliberate actions to compromise its development environment and trojanize a signed and authenticated SolarWinds Orion update. Then it used the same advanced methods to install and maintain persistence in the victim organizations.

This brings us to software supply chain attacks. Just as APTs and zero days overlap in today's attack methodologies, software supply chain attacks can use a variety of techniques—including zero day exploits and advanced persistent techniques—to breach the perimeter and carry out malicious activity down the software supply chain. But, as in the case of the SUNBURST attack, the nation-state actors didn't simply wait for a vulnerability to become apparent. They gained access to the build server for SolarWinds Orion and created their own vulnerability within a software update channel, using advanced techniques to embed their malware into the Orion patch without notice. Once customers installed the patch, the attackers had access to victim servers and went undetected until an outsider, FireEye, alerted SolarWinds to the breaches.

This is where supply chain attacks take advanced persistence and zero day sophistication to a new level. Once the malicious code had been inserted into the software update, the attack spread to other victims in the supply chain and hid by remaining inert until after the update was tested and installed by SolarWinds customers—more than 18,000 of them in total.

At the time of this writing, the SUNBURST trojan continues to persist and spread, while researchers are still uncovering new advanced methodologies for hiding and perpetuating this supply chain attack, and new families of malware used in the attack continue to be uncovered, according to a March advisory by Microsoft. Unfortunately, functional testing against the downloaded Orion update didn't catch the trojan because it wasn't doing anything to interact with the test. Instead, the SUNBURST back door trojan waited to activate until after testing and installation, and even then, it moved low and slow staying under the radar while also disabling security and monitoring tools on victim networks and endpoints.

Advanced attacks are prevalent and successful against enterprises and organizations of all sizes. In the following sections, we describe how to classify these attacks and where to apply controls.

Advanced Adversaries

SUNBURST represents the work of thousands of developers associated with the Russian adversary group Cozy Bear, while the Microsoft Exchange Server zero day attacks tie back to well-funded Chinese hacking groups. That means attackers and researchers developing these advanced attack methods likely have the means to far outspend the defenders in IT networks, who are usually strapped for the budget and skills they need to detect and respond adequately to APTs, zero days, and advanced supply chain attacks.

Well-funded attacks like these succeed because of this disparity. Everyone's a potential target, from hospitals and other critical infrastructures to tech companies to government agencies. As adversaries up their games, so, too, must the defenders, lest they become the victim. IT security teams need to look beyond the low-hanging fruit they know how to protect and consider new ways advanced attackers are bypassing their controls.

What's in a Name? NIST Definitions

A lot of terms used interchangeably to describe advanced attacks actually mean different things. Below, NIST defines the difference between a threat, vulnerability, and attack, all of which pertain to advanced attack methods:

Threat: NIST defines a cyberthreat as "any circumstance or event with the potential to adversely impact organizational operations, organizational assets, or individuals through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service. Also, the potential for a threat-source to successfully exploit a particular information system vulnerability."

Vulnerability: A vulnerability is a weakness in a system or systems (misconfigurations, software bugs, or actions of a system) that exposes it to unauthorized access.

Cyber Attack: An attack is the "targeting an enterprise's use of cyberspace for the purpose of disrupting, disabling, destroying, or maliciously controlling a computing environment/infrastructure; or destroying the integrity of the data or stealing controlled information."

Adversary Sophistication

Advanced Persistent Threats (APTs) indicate an adversary with sophisticated levels of expertise and significant resources. These attacks and their perpetrators are persistent, adaptive, and determined. NIST writes that in advanced attacks, adversaries "use multiple different attack vectors (e.g., cyber, physical, and deception) to generate opportunities to achieve objectives, which are typically to establish and extend footholds within the information technology infrastructure of organizations for purposes of continually exfiltrating information and/or to undermine or impede critical aspects of a mission, program, or organization, or place itself in a position to do so in the future."

NIST defines a supply chain attack as those in which an adversary implants malware or other vulnerabilities prior to installation of applications, updates and components "in order to infiltrate data, or manipulate information technology hardware, software, operating systems, peripherals or services at any point during the life cycle."

According to a 2020 supply chain report, supply chain attacks rose 430% in just one year.

Supply chain attacks are also usually well-resourced with the backing of nation states or organized crime rings. These adversaries see supply chain attacks as a means to burrow deep into the technology and industrial infrastructures of organizations and to exfiltrate state and tech secrets of value—or launch a debilitating cyber attack. According to a 2020 supply chain report, supply chain attacks rose 430% in just one year between 2019 and 2020. These attacks are now targeting environments where software is being developed and shared, which is the very core of the software supply chain. They're doing so by spoofing and infecting code

components in public repositories widely used by third party developers of software. So, in that scenario, tier-1 software providers in the supply chain have been tricked into downloading malicious code to their development teams, for example.

Finally, NIST defines a zero day attack as "an attack that exploits a previously unknown hardware, firmware, or software vulnerability." Zero Days are frequently used in APTs and supply chain attacks, including SolarWinds Orion SUNBURST and the Exchange Server attacks.

The Rarity of Zero Day Exploits

Searching for new vulnerabilities and developing the code to actively exploit them requires a high degree of knowledge and sophistication on the part of the adversary. However, zero day vulnerabilities are not always identified by well-funded nation-states or organized criminals. In the case of a lone wolf who discovers a new operational flaw or coding bug from his basement, for example, it may not become a malicious attack unless that lone wolf runs scripts against the vulnerabilities for criminal gain, or shares or sells the vulnerabilities to criminal gangs. New vulnerabilities are discovered by white hat and gray hat hackers every day and reported legitimately, often through bug bounty programs that pay generously for the bug finds.

But for nation states with deep resources, these zero days are highly valued and developed through teams of experts skilled at finding and exploiting vulnerabilities. They are also expertly executed using advanced techniques for deployment and hiding, while these nation-states keep the vulnerabilities and their exploits secret for as long as possible.

Overlapping and Interacting

While attacks generally leverage weaponized exploits, an exploit exists independent of attacks—an exploit may or may not ever get used in an attack. So, for consistency, the rest of this paper will refer to zero day attacks, where zero day exploits have been actively used for criminal purposes.

The types of advanced attacks this paper discusses can work independently or together as part of a larger attack chain. Not all advanced attacks include zero days, but nearly all zero-day attacks are considered advanced as they are not common in the wild. Not all supply chain attacks deploy zero days, but supply chain attacks are among the most advanced attacks businesses and agencies must defend against. Table 1 compares differentiators and commonalities as they apply to well-funded, nation-state adversaries across each of these three types of attacks.

Table 1. Advanced Attack Types and How They Align

ADVANCED ATTACK TYPE	EXPLANATION & EXAMPLES	HOW THEY'RE ADVANCED
Zero Days exploit new unknown vulnerabilities to take advantage of a window of opportunity when there is no awareness of or patches for said vulnerabilities.	Zero days are usually kept secret for maximum impact. The goal is to infect target systems and maintain control by keeping the zero days secret until the attackers have exploited their target systems. A recent example, the Microsoft Exchange Server zero days, shows how fast attackers move to exploit new vulnerabilities before target victims know about the vulnerabilities and/or install patches.	Zero days require research and testing to discover and create web shells used to exploit the vulnerabilities. Attacks against zero days can include advanced techniques to exploit the vulnerabilities and persist without notice. Zero days may also be included in a larger advanced persistent attack as the initial attack vector or to spread internally in new unknown ways. For example, threat actors exploiting the Exchange zero days are also using homemade web shells to establish persistence and launch secondary attacks, according to reports.
Advanced Persistent Attacks are good at installing, hiding, and performing other malicious actions inside a victimized organization for long periods of time without notice.	Advanced Persistent Attacks are methodical, multi-staged, and well-funded. They are designed to linger without detection for long periods of time and are often only caught by a third party. These methods are well-thought out and well-funded, such as with CozyDuke reported in 2015 by the Russian sponsored group, CozyBear, and with the SolarWinds SUNBURST supply chain attack.	Advanced Persistent Attacks use mostly known vulnerabilities and don't always involve zero days. The sophistication is in how criminals string together multiple steps, from initial phish email to infiltration, to lateral spread, establishing command and control (C2) services and copying and sending data out. They are also advanced in how they hide these activities by turning off or circumventing detection tools.
Supply chain attacks take advantage of trusted relationships between software providers and their customers. They infiltrate a supplier and then spread to its customers through what should be secure and trusted applications, programs, API's, and updates.	The SUNBURST attack that launched from infected SolarWinds build servers to its update servers and then installed on downstream client servers is a good example of a software supply chain attack that utilizes advanced persistent attack methodologies across the supply chain to spread and persist. Open source software repositories are also being targeted to hide malware in components or outright counterfeit components and modules.	Supply chain attacks are perhaps one of the most sophisticated forms of advanced attacks because they abuse the trusted relationships between software developers and their repositories, and between the software developer and their customers down the supply chain. Supply chain attacks can use any number of zero days and APTs to install, set up C2 channels and maintain their presence by turning off or circumventing security tools and intelligence feeds.

Supply chain attacks take advantage of trusted relationships between software providers and their customers.

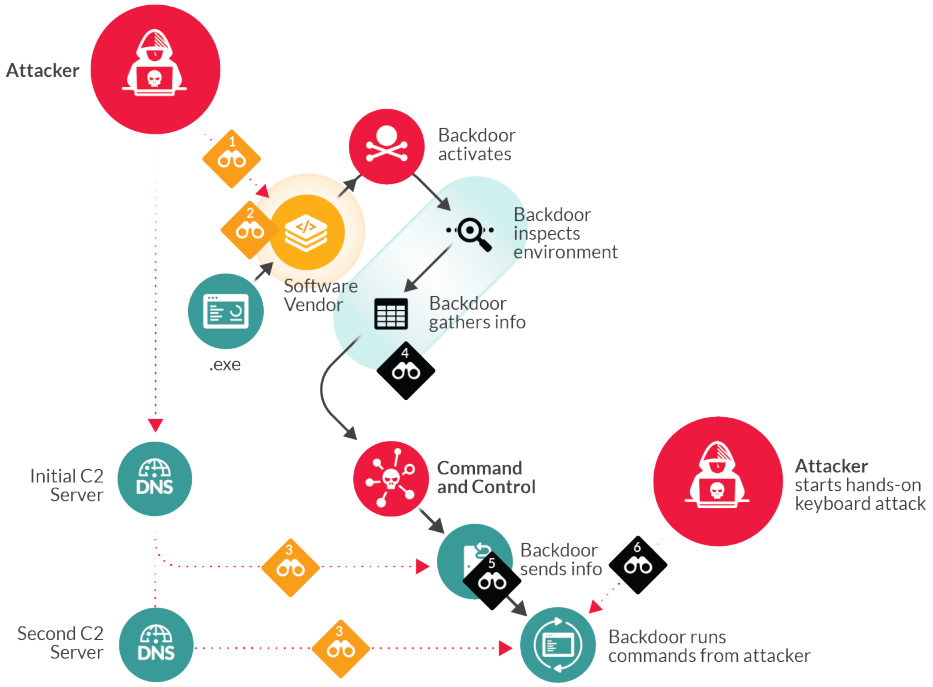

Shared Responsiblity: Optimal Visibility into a Zero Day Attack

Software Vendor Responsibility

1. The software vendor should be monitoring their perimeter to catch attackers trying to get in. They almost certainly were monitorin th is, but failed to catch the attacker's initial entry.
2. Within the perimiter, the software vendor should be monitoring E/W traffic for lateral movement and watching these areas especially closely. Because they are a software provider that ships updates to hundres of thousands of customers, their DevOps pipelines and build servers are very important. If unusual lateral movement or other activity touches the build servers, they should know.
3. The software vendor should be watching their perimeter for C2 traffic. This is detectable via ML behavioral modeling against DNS traffic.

Enterprise Responsibility

4. When the "Backdoor Inspects Environments and Gathers Info" that's probably via network scanning and various queries that generate network traffic, visible through NDR.
5. When the "Backdoor sends info" it generates network traffic, which is visible through NDR.
6. Hands-on-keyboard attacks often generate anomalous interactive traffic. The attacker may be directly manipulating endpoints through remote acces tools such as VDI, RDP, WinRM, remote powershell, etc. This traffic is also visible through NDR.

Figure 1. From Zero Day to Supply Chain: Extended Attack Map

As we can see from the table above, the lines between these attacks are blurred as well-funded adversaries continue to one up the enterprises and SMBs they're targeting. In the most recent abuse of the Exchange server zero days, for example, at least one of the attackers linked together all four zero days to execute remote code and take control of Exchange email accounts belonging to infectious disease researchers, law firms, higher education, defense contractors and policy think tanks, according to reports. See Figure 1 to visualize all these advanced attack methods working together in a 'super' attack.

Abusing Trust

Even with layers of security on the endpoints (for example, antivirus and EDR), the network (IDS, firewalls, network traffic inspection) and around applications (such as email security for Exchange), advanced attacks still succeed primarily because they mask their behavior under the guise of trusted system behaviors. In most advanced attacks, trust is abused at all stages, enabling attackers to break in, beacon out to their command-and-control server, exfiltrate data, and obfuscate all of this activity for as long as possible. This "kill chain" typically involves multiple exploits that enable activities including:

1. Gaining Initial Access on an Endpoint—usually through spearphishing, a web download, or misconfigured Web, DNS and other public-facing servers.
2. Establishing a Foothold by conducting internal system and network reconnaissance with passive scanning for vulnerable assets.
3. Opening backdoors and encrypted tunnels to establish C2 connections.
4. Maintaining persistence by hiding in memory and other hard-to-monitor locations, rewriting itself, shutting down when it knows it's being watched, hiding from security tools (or even manipulating the tools).
5. Gaining administrative access for higher level privilege and gaining control of more endpoints and systems on the network.
6. Spreading to other systems with administrator access (commonly because all too often, admins use passwords that are easy to guess).
7. Establishing internal control servers to hide stolen data inside the network before transmitting it out to secondary servers maintained by attackers.
8. Establishing C2 channels for external transmission of stolen data and credentials.
9. Transferring data out of the network, often encrypted, in low slow bursts to fly under the radar.

Kill Chain and IOCs

The cyber kill chain emphasizes stopping an attack as early in the chain as possible.

These stages of advanced attacks are often referred to as the cyber kill chain, which emphasizes stopping the attack as early in this chain as possible, preferably before the phish succeeds, and if not, then before C2 communications are established and sensitive data is transferred to internal host servers and then exfiltrated out of the organization altogether. Early detection is more difficult with advanced attack methods and zero days. Absent advanced intelligence on these advanced attacks, many tools are blind to these indicators. Figure 2 shows an advanced attack kill chain.

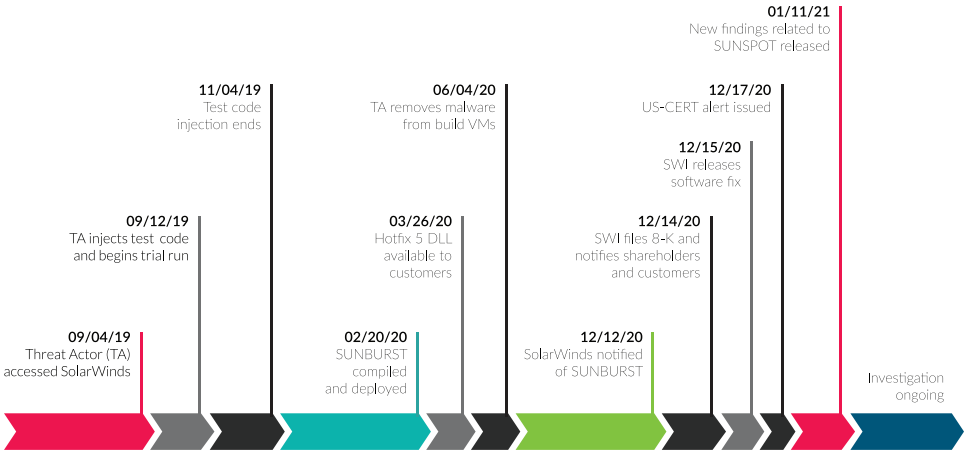

All events, dates, and times approximate and subject to change; pending copmleted investigation.

Figure 2. Shutting Down Advanced Persistent Attacks Along the Kill Chain

Detecting these actions as indicators of compromise (IOCs) is a common methodology used to try and catch advanced forms of attacks occurring in systems and networks. But because these indicators are so well-hidden and obfuscated in advanced attack methods, security teams need visibility with context to observe and classify unusual behaviors and traversals from endpoints and systems on the network and in the cloud. This requires the ability to passively monitor traffic patterns inside the network, combined with visibility across devices and workloads on-premises and in the cloud. Since the traffic in many advanced attacks, including the C2 traffic, is encrypted, network operators also need the ability to securely decrypt traffic and identify malicious activity in the traffic packet's metadata.

Many victim organizations have not fully detected, let alone cleared the SUNBURST attack modules from their systems, so SUNBURST still persists and spreads. Detection came sooner in the case of the Exchange zero days, but not before attackers scanned and breached some 250,000 servers, according to the WSJ. Once discovered, Microsoft quickly released patches for the vulnerabilities as they were being actively exploited. They also released a script for organizations to test the security status of their Exchange servers. Even after these resources were made available to Exchange users, attacks against these vulnerabilities continued to spread.

Beat Them at Their Own Game

Stealthy attacks require stealthy defenses, particularly since advanced persistent attack methods are 'security aware' and turn off security or hide themselves at the first hint of being observed. These attack methods can turn off logging and erase logs, disable agents and antivirus, and co-opt security tools to use for malicious activity. Advanced attackers also hide their activities in network, DNS, and HTTPS traffic while tweaking network firewalls and egress filtering.

Advanced attackers can hide their activities in network, DNS, and HTTPS traffic while tweaking network firewalls and egress filtering.

Because advanced attacks of all types are on the rise, IT teams should prepare for new vectors and hiding techniques that bypass their traditional security methods. This protection starts with strong vulnerability management and access and identity programs, but it shouldn't stop there.

Because most tools are blind to advanced attacks, organizations should look to network activity (both on-premises and in the cloud) to flag activity that differs from the normal baseline. Abnormal traffic inside and outside of the enterprise, including its cloud services, is a good indicator of malicious activity.

When suspicious behavior is identified, IT teams need the ability to drill down into the traffic itself. Packet-level data allows administrators to identify IP addresses, ports, and domains as a starting place. Decryption and full stream reassembly bolstered by protocol analysis allows for the identification of services, applications, malformed traffic, and much more. This level of data coupled with behavioral baselines allows for the detection and response to off hours or unusual activity that is indicative of malicious activity, while providing analysts the context to differentiate between the unusual (but benign) and the truly malicious.

Table 2. Network Inspection of APTs, Supply Chain and Zero Days

EXAMPLE	UNUSUAL BEHAVIOR	DEEPER EXAMINATION
Advanced Persistent Attack	Unusual connections between internal systems that don't usually communicate, such as the devops network and the finance servers Lateral movement using different credentials Increase in encrypted (SSH) traffic with remote execution commands Requests for multiple HTTPS connections inside and outside of the network Unusual and vulnerable ports and services outside of approved security policy (ports and traffic associated with SSH, TelNet, SMTP, DNS, NetBIOS, etc.) Abnormal file reads	Examine packet data to reveal metadata about the traffic, including login information, source and destination IP address, ports and services, protocols such as SMTP (for mail transfer), and more. This metadata should be viewable even if packets are encrypted. Also monitor any traffic to and interactions with security software and confirm status of endpoint agents, network firewalls and DNS services.
SUNBURST Supply Chain Advanced Persistent Attack	All of the above, plus: • A sensitive server, the Orion network management server, connects to an external host • Attacker hostnames match victim environment • Known malware like TEARDROP and BEACON included in components of the attack • DNS Connections from internal systems to unusual external domains, including country extensions outside of the normal geographic region • Temporary file replacements and task modifications	Ensure antimalware is up to date and able to detect known malware like TEARDROP and BEACON since known malware is often used as components of advanced persistent attacks. Monitor the network for unusual DNS requests and behaviors, compare host and domain names, follow traversals and ports, and open packet headers for metadata comparison: compare host names, destination and host IPs, file transfer requests and more.

Security teams need a high level of visibility into unusual behaviors across their physical and cloud networks.

EXAMPLE	UNUSUAL BEHAVIOR	DEEPER EXAMINATION
Exchange Zero Day Attacks	Unusual command executions on Exchange New and unusual forms of ASPX and other files in temporary and user file systems Directories with nonexistent resources Spoofed HTTP agents Unusual Exchange PowerShell Snapin Requests	Examine log folders on endpoints and network devices. Followup with network traffic analysis to detect unusual behaviors, actions, destinations, services, and files such as ASPX unusually located in temporary and user file systems. Also follow previous advice to detect advanced persistent attacks because zero days are often exploited and extended using advanced persistent attack methodologies.

Passive Monitoring

Security teams need a high level of visibility into unusual behaviors across their physical and cloud networks, but monitoring efforts should not be detectable to their adversary. A covert position defeats tactics like those used in the SUNBURST attack, where endpoints running an agent for security tools like EDR were identified and simply avoided, allowing attackers to hide in unmonitored spaces and avoid detection.

Conclusion: Invisible Visibility

Software supply chains are rich hunting grounds for well-funded attackers because they attack once and spread to many along the supply chain, particularly to other software and technology companies, infrastructure agencies and even researchers working on valuable intellectual property. Expect to see more advanced persistent attacks that will merge with zero days and set their aim on software supply chains as well-funded attackers up the ante against lesser-funded enterprises and SMBs with limited resources to protect and respond.

By understanding how advanced attacks operate, who's behind them, and what their targets are, IT teams can update their processes and controls to detect and block these attacks early in the kill chain before lasting damage is done. Layered security at the endpoint, application, and network are still critical protection and detection techniques, but IT teams need invisible visibility into network and endpoint activity if they're to beat the attackers at their own games.

Deb Radcliff

ABOUT THE AUTHOR

Deb Radcliff has more than 25 years of experience in the cybersecurity industry, starting out as the first investigative journalist to make cybercrime a beat where she followed the FBI, DoD, Secret Service, CIA, local and state law enforcement as they were building their own cyber units. Her articles are cited in numerous research papers and college textbooks. She spoke at West Point, won two Neal Awards, and was runner up for a third. In 2005, Radcliff stood up a new Analyst Program for the SANS Institute and oversaw focused whitepaper and webcast content developed by SANS experts for 15 years. She still writes for CSO, her own blog OnlineCrimeBytes. The first book in her cyberthriller fiction series, "Breaking Backbones: A Hacker Trilogy," is available on Amazon.