Google Cloud Security

Mandiant

M-Trends

2024 Special Report

Chinese Espionage Operations Targeting The Visibility Gap

Endpoint detection and response (EDR) platforms have become commonplace among companies seeking to expand the visibility into endpoint activity necessary to provide a baseline of security monitoring. This increase in visibility has forced attackers to evolve in order to maintain operational efficacy. While some attackers have invested in EDR bypass techniques, others have, instead, chosen to focus on areas of the corporate environment where in-depth visibility remains uncommon. While EDR agents have become a standard part of security deployments, many specialized appliances that either segment or host assets critical to the organization often lack similar levels of visibility. These systems have become a new preferred safe haven for attackers as it enables them to maintain long-term persistence with lower risk of detection due to the gap in visibility.

Common examples of devices that rarely support EDR deployment are firewalls, email filtering products, virtualization platforms, and virtual private network (VPN) solutions. To further complicate matters, the platforms on which these appliances are built may be proprietary or otherwise locked down, such that forensic analysis efforts are hindered. Exploits for these devices are exceedingly valuable to attackers, primarily because they typically require no user interaction to succeed, which helps to minimize the chance of detection. If an attacker possesses an exploit for a zero-day vulnerability on these devices, they are often able to gain access to a target environment and remain undetected for an extended period of time. Furthermore, the attacker can use the exploit to gain access to additional targets or reestablish access to the same target if it is disrupted.

Zero-day: Vulnerabilities disclosed before patches are made available.

N-day: Vulnerabilities first exploited after patches were made available.

Mandiant observed a range of attackers targeting devices that matched this profile in 2023. Sandworm[1] continued to leverage access via compromised network edge infrastructure to enable their wartime operations in Ukraine. The financially motivated group FIN11[2] exploited a zero-day in MOVEit Transfer software to steal data as part of their data theft extortion operations. In the past year alone, Mandiant has investigated several high-profile cases of suspected Chinese espionage operations leveraging zero-day and n-day vulnerabilities to target systems where visibility has been difficult to instrument.

Custom Malware for Edge Devices

Security and networking devices that sit at the logical perimeter of a network and host services on the internet are often referred to as “edge devices.” Mandiant has observed a trend in which China-nexus attackers have gained access to edge devices via exploitation of vulnerabilities, particularly zero-days, and subsequently deployed custom malware ecosystems. These malware ecosystems have typically consisted of several distinct code families that attackers operate in unison, and are usually custom developed or tailored for the target edge device and underlying operating system. Developing malware for these managed appliances is a non-trivial task. Vendors typically do not enable direct access to the operating system or filesystem for appliance device owners or users. In order to operationalize attacks within this class of platform, attackers must maintain a resource intensive malware development lifecycle that, by necessity, maintains flexibility and a high degree of technical acumen. While this process requires a substantial investment, it also produces clear results when leveraged successfully by attackers.

Remaining Undetected

In general, custom malware may go undetected for long periods of time since there is unlikely to be specific detections in place for the malware. This is particularly true for edge devices, where network defenders may have little to no means for monitoring and detection of malware activity. Malware authors may also take special care to ensure that the malware hinders forensic investigation by circumventing or clearing logging systems in place on the device. Even after the malware has been discovered and exposed by the security community, it is often a non-trivial task for a device owner to identify if they have been impacted, since off-the-shelf security products typically will not support edge devices. Furthermore, it is sometimes the case that exploitation attempts of zero-day vulnerabilities leave little or no reliable evidence behind. This is often exacerbated by the fact that the attack may have occurred months or even years prior to detection. Additionally, there are often challenges with traditional forensic techniques as these devices are typically kept under tight control by manufacturers, adding to the already complex nature of investigating these types of compromises.

Example: BOLDMOVE

BOLDMOVE is a backdoor used by suspected Chinese espionage groups that has both Windows and Linux variants containing a core set of features. Mandiant identified a custom variant of the Linux version of BOLDMOVE, which contained an extended set of features to remain undetected on Fortinet devices. This variant of BOLDMOVE disabled the `miglogd` and `syslogd` logging daemons on the appliance, and contained a command to patch memory address space for these logging functions. These customizations of the BOLDMOVE backdoor are suspected to have enabled the attacker to remain undetected for a longer period of time than they would have otherwise been able to through traditional means.

Reduced Complexity, Increased Reliability

Edge devices such as email security gateway appliances and VPNs are typically high-availability devices that run for months or years at a time without being rebooted. As such, many of these devices are put through rigorous testing regimes by the manufacturer during development to ensure their stability. China-nexus malware developers take advantage of the built-in functionality included in these systems, which benefits them in several ways. In general, leveraging native capabilities will enable attackers to reduce the overall complexity of the malware by instead weaponizing existing features within that have been rigorously tested by the organization. For example, for devices that use proprietary software components such as custom file formats or configuration files, attackers may be able to leverage built-in functions to parse or process these files rather than developing their own implementation. This concept is analogous to living-off-the-land, and is particularly effective on edge devices since these native device operations are likely not being monitored by network defenders and, as such, may go unnoticed.

> **Living off the land:** Attacker use of legitimate, pre-installed tools and software within a target environment, notably to evade detection.

Example: THINCRUST

During an UNC3886 compromise, Mandiant discovered a backdoor deployed to FortiAnalyzer and FortiManager devices named THINCRUST, which disguised its command and control (C2) communications as legitimate API calls to the devices. UNC3886, a suspected Chinese espionage group, appended the Python-based backdoor code into legitimate web framework files that were responsible for providing the API interface for the appliance. This gave UNC3886 the ability to harness the native API implementation to access and send commands to THINCRUST by simply interacting with a new endpoint URL, which they had added. By leveraging existing capabilities built into the appliance, UNC3886 was able to simplify their malware while maintaining the reliability necessary for continued operations.

Tailored Capabilities and Smaller Footprint

Custom malware for edge devices may only include the capabilities required to achieve the attacker's mission objectives. Malware used to exploit vulnerabilities may never be used again once the vulnerabilities are discovered and patched, as the cost to maintain and repurpose the code often outweighs the benefits. By developing relatively simple malware that serves only to provide the attackers with the desired functionality on the target device, attackers are able to achieve their goals while minimizing their overall footprint. In the same vein, the requirement for complex obfuscation is likely lessened since the primary objective of the attacker is to remain undetected entirely, rather than hinder the analysis of the malware once it has been discovered. By the time the malware is discovered, the vulnerability and campaign would have already been exposed, and the attacker's operations typically come to a close.

Activity after internet access restrictions implemented to FortiManager

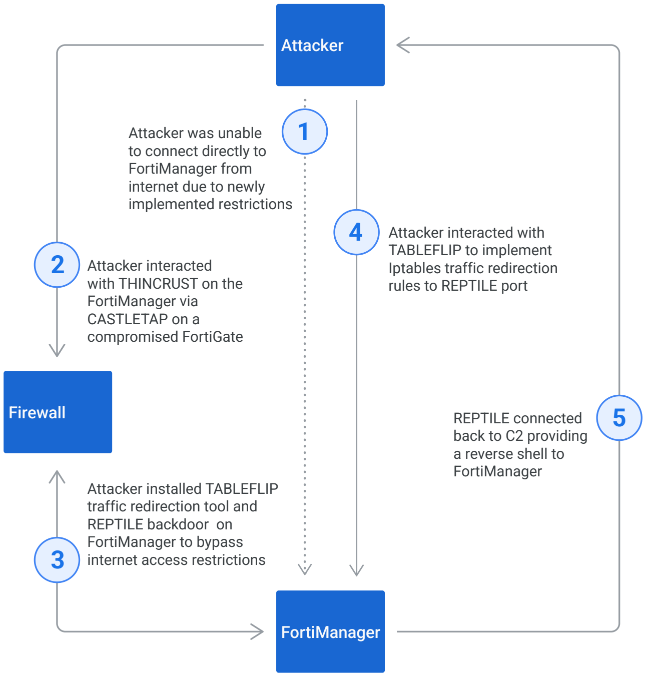

1 Attacker was unable to connect directly to FortiManager from internet due to newly implemented restrictions

2 Attacker interacted with THINCRUST on the FortiManager via CASTLETAP on a compromised FortiGate

3 Attacker installed TABLEFLIP traffic redirection tool and REPTILE backdoor on FortiManager to bypass internet access restrictions

4 Attacker interacted with TABLEFLIP to implement Iptables traffic redirection rules to REPTILE port

5 REPTILE connected back to C2 providing a reverse shell to FortiManager

Example: TABLEFLIP

After losing access to a FortiManager device during one incident due to access control lists change, UNC3886 adapted to the situation by deploying a network traffic redirection utility named TABLEFLIP. TABLEFLIP passively listens on all active interfaces for specialized command packets that contain an XOR encoded IP address and port to redirect traffic to using iptables commands. UNC3886 deployed TABLEFLIP alongside the publicly available REPTILE rootkit to act as a reverse shell, and successfully gained access back to the FortiManager device. The ability to produce purpose-built malware in response to changes in an ongoing operation places defenders at a disadvantage when facing capable and agile attackers with nation-state backing.

Attribution Challenges

Custom malware developed for edge devices may stifle attribution for cyber threat intelligence analysts. These malware families, and potentially the entire ecosystem, could be almost entirely unique when compared to existing malware because of the target operating system and tailored capabilities. As such, they may not contain code or other overlaps analysts traditionally find between related malware families that contribute to the technical attribution analysis.

Example: SEASPRAY and WHIRLPOOL

SEASPRAY is a launcher written in Lua that UNC4841 injected into legitimate Barracuda Email Security Gateway (ESG) modules. SEASPRAY registers an event handler for incoming emails, and launches an external binary, which Mandiant tracks as WHIRLPOOL, when certain markers are present. WHIRLPOOL is a simple TLS reverse shell utility that receives a C2 IP address and port to connect to from SEASPRAY at runtime. Because SEASPRAY was a relatively simple implementation that consisted of a few lines of code that were specific to the Barracuda ESG appliances, it did not offer much value in terms of attribution. Similarly, WHIRLPOOL was a simple and generic TLS reverse shell that did not contain any embedded C2 server information that could be analyzed. Usage of such malware on edge devices presented significant challenges for analysts performing attribution analysis.

In-Depth Knowledge of Edge Devices

Mandiant has observed several instances where China-nexus attackers demonstrated a high level of in-depth knowledge when targeting edge devices. The degree of knowledge spanned not only the malware used during the attack, but also the zero-day vulnerabilities used to gain access to these devices.

Example: DEPTHCHARGE

DEPTHCHARGE is a passive backdoor that Mandiant observed UNC4841 begin to deploy about one week after Barracuda's initial public notification of the ESG zero-day campaign. This was followed by more rapid deployment to what Mandiant assessed were high-value targets, once Barracuda announced plans to replace affected devices. The timing of the accelerated deployment suggests that UNC4841 may have anticipated this, and was prepared for remediation efforts with tooling and tactics, techniques and procedures (TTPs) designed to enable them to continue operations in the face of attempts to disrupt their access to target networks.

Several aspects of DEPTHCHARGE and its execution chain demonstrated an intimate knowledge of the Barracuda ESG device and its software components. Most notable was that the attacker had identified a method for malware persistence inside the configuration database for the appliance, which would result in it being present in exported backup configurations. This meant that device owners looking to set up a clean device would unknowingly export backup configurations containing DEPTHCHARGE persistence, and in-turn infect their clean appliances when attempting to restore their configuration.

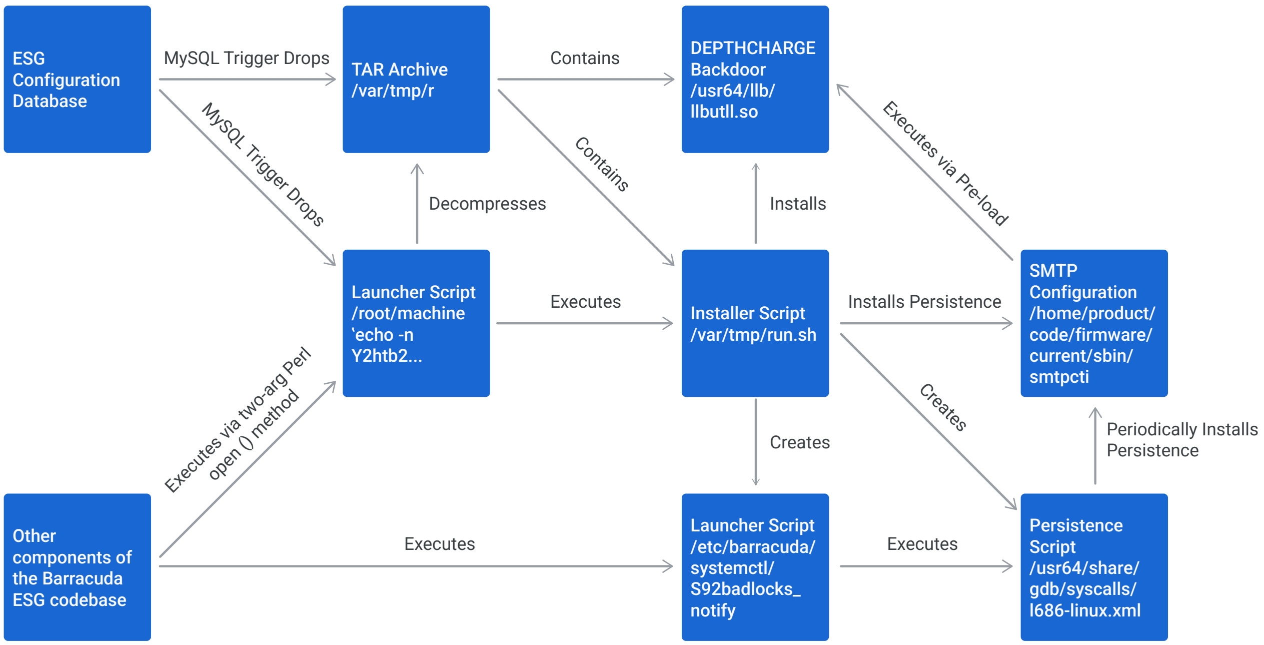
ESG Configuration Database
MySQL Trigger Drops
TAR Archive /var/tmp/r
Contains
DEPTHCHARGE Backdoor /usr64/llb/llbutll.so
MySQL Trigger Drops
Decompresses
Contains
Installs
Executes via Pre-load
Launcher Script /root/machine 'echo -n Y2htb2...
Executes
Installer Script /var/tmp/run.sh
Installs Persistence
SMTP Configuration /home/product/code/firmware/current/sbin/smtpcti
Executes via two-arg Perl open () method
Creates
Creates
Periodically Installs Persistence
Other components of the Barracuda ESG codebase
Executes
Launcher Script /etc/barracuda/systemctl/S92badlocks_notify
Executes
Persistence Script /usr64/share/gdb/syscalls/I686-linux.xml

Perhaps an even more intricate display of UNC4841's knowledge of the inner workings of the appliance was the method through which DEPTHCHARGE achieved command execution from a trigger inside the MySQL configuration database after it had been imported. UNC4841 understood that completely separate components of the ESG's codebase accessed files using the two-argument form of Perl's open() function, and that they could craft a special filename that would result in commands being executed on the appliance. By having the MySQL trigger drop a file that induces this command execution, UNC4841 was able to chain these techniques together to have DEPTHCHARGE dropped and executed upon import of a backup configuration on a new appliance. This effectively enabled UNC4841 to survive through a complete device replacement in the small number of cases where this occurred.

Custom Malware for Hypervisors

As cloud computing has grown in popularity over the years, hypervisors have subsequently become commonplace within modern infrastructures. However, while instrumenting endpoint visibility on the guest virtual machines is relatively easy, instrumenting visibility on the hypervisor itself can present substantial challenges. As with network edge devices, Type 1 hypervisors commonly run on operating systems versions that are rarely supported by EDR vendors. Not surprisingly, given the low visibility yet excessively high target value of a hypervisor, Mandiant has observed China-nexus attackers targeting hypervisors as well.

Hypervisor technologies, such as VMware's ESXi, use Virtual Machine Communication Interface (VMCI) sockets to facilitate communication between the bare-metal host and the guest operating systems. Mandiant has observed attackers leverage VMCI sockets for lateral movement and continued persistence within targeted environments. UNC3886 utilized backdoors such as VIRTUALPITA, and took advantage of VMCI-based channels to communicate from the ESXi host to the guest virtual machines (VM). Since the traffic over the virtualized layer is localized to the bare metal machine, there are no security mechanisms restricting any guest VM or ESXi host from initiating a connection with the other, essentially bypassing any network segmentation. Additionally, traffic cannot be monitored outside of the guest VMs and ESXi hosts present in the virtualized environment. While the client/server communication socket has connection-oriented and

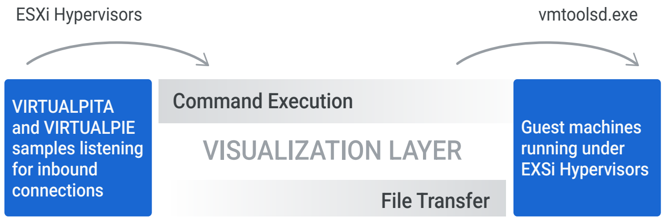
ESXi Hypervisors
vmtoolsd.exe
VIRTUALPITA and VIRTUALPIE samples listening for inbound connections
Command Execution
VISUALIZATION LAYER
File Transfer
Guest machines running under EXSi Hypervisors

connectionless variants very similar to TCP and UDP, it is invisible to commonly used networking tools such as tcpdump, netstat, nmap, and Wireshark without custom configurations, as it belongs to a different socket address family.

UNC3886 used a novel persistence technique to deploy the VIRTUALPITA backdoor using vSphere Installation Bundles ("VIBs"). Mandiant had not previously observed this technique for deploying malware or persistence, which suggests UNC3886 spent significant resources to understand the inner workings of VMware technologies, especially the VMCI sockets for circumventing security restrictions. Furthermore, VIRTUALPITA was used to pass arbitrary commands to guest VMs without being logged on the host. These commands then get executed on the guest VM with the vmtoolsd.exe process where they are then observed in the Windows event logs. VIRTUALPITA also sets the HISTFILE to 0, which would remove any terminal history on the host system, leaving behind little forensics evidence.

Recommendations

The most critical strategy for protecting against such attacks is maintaining proper patch management to mitigate the risk of exploitation of known vulnerabilities. Applying the most recent patch is the best way to limit any unexpected tampering or modification of the appliance. For zero-day vulnerabilities, where exploitation is unlikely to be detected, a defense-in-depth approach provides the best chances of surfacing evidence of the malicious activity further in the attack lifecycle.

If an organization has identified that they were operating vulnerable devices and may have been compromised, Mandiant recommends performing an investigation and hunting activities within their networks. An investigation may include, but is not limited to, the following:

- Scanning potentially impacted devices with publicly available tools such as IOC scanners to identify evidence of compromise.
- Sweeping the entire environment for known IOCs.
- Reviewing network logs for signs of data theft and lateral movement.
- Reviewing network logs for abnormal logins or internal traffic from edge devices.
- Capturing a forensic image of the impacted appliance and conducting a forensic analysis.
- Applying any malware signatures (e.g. YARA rules) to appliance images to assist forensic investigators.

Organizations should also consider implementation of security controls detailed in architecture hardening guidance provided by security vendors. Mandiant has previously provided such documentation for the Barracuda ESG event.[3]

Outlook

Despite the amount of resourcing required, Chinese espionage groups are almost certainly going to continue investing in the acquisition of zero-day exploits and platform-specific tooling. Mandiant expects that we will continue to see targeting of edge devices and platforms that traditionally lack EDR and other security solutions due to the challenges associated with discovery and investigation of compromise. Exploitation of these devices will continue to be an attractive initial access vector for Chinese espionage groups to remain undetected and maintain persistence into target environments.

It is also likely that we will continue to see the deployment of custom malware ecosystems by Chinese espionage groups that are tailored for the device and operation at hand. This approach provides several advantages such as the increased ability to remain undetected, reduced complexity and increased reliability, and a reduced malware footprint. Additionally, it presents challenges to technical attribution being performed by threat intelligence analysts.

Organizations must remain vigilant and ensure that they're not only monitoring their networks at the operating system layer, but also continue to patch, maintain and, where possible, monitor the appliances that are running the underlying infrastructure of their networks.

[1] https://cloud.google.com/blog/topics/threat-intelligence/gru-disruptive-playbook
[2] https://cloud.google.com/blog/topics/threat-intelligence/zero-day-moveit-data-theft
[3] https://services.google.com/fh/files/misc/barracuda-esg-rpt-en.pdf

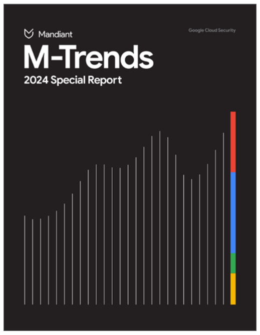

Read the full report: M-Trends 2024 Special Report

For more information, visit cloud.google.com/security.