Examining Exchange Exploitation and its Lessons for Defenders

Background

On 02 March 2021, Microsoft released out-of-band updates for Microsoft Exchange to cover four actively-exploited vulnerabilities:

- CVE-2021-26855: a pre-authentication Server-Side Request Forgery (SSRF) vulnerability enabling access to a vulnerable Exchange server. This specific vulnerability, identified by researchers at DEVCORE, is also referred to as ProxyLogon.
- CVE-2021-26857: a privilege escalation vulnerability allowing an attacker with code execution to run commands as SYSTEM.
- CVE-2021-26858: a post-authentication arbitrary file write vulnerability in Exchange allowing an attacker to write contents to any accessible part of the victim system.
- CVE-2021-27065: another post-authentication arbitrary file write vulnerability.

Used together, these vulnerabilities allow for remote access to an exposed Microsoft Exchange instance, follow-on code execution at privileged levels, and the ability to establish persistence on the victim system. DEVCORE research started in October 2020, with

acknowledgement from Microsoft that the SSRF vulnerability ProxyLogon existed on 06 January 2021.

In coordination with Microsoft's release, information security company Volexity released its own report covering intrusion activity utilizing this exploitation chain, referred to as Operation Exchange Marauder, since 03 January 2021—interestingly, a few days prior to Microsoft's acknowledgment of the vulnerability . Independently of Microsoft and Volexity, FireEye also released reporting on overlapping exploitation activity taking place since January 2021, although no precise date was provided in their post.

Following disclosure on 02 March 2021, multiple parties reported odd activity prior to release and substantial increases in Exchange targeting shortly thereafter. Most notably, several entities reported widespread scanning of Microsoft Exchange servers just prior to Microsoft's vulnerability disclosure, from 27 to 28 February 2021.

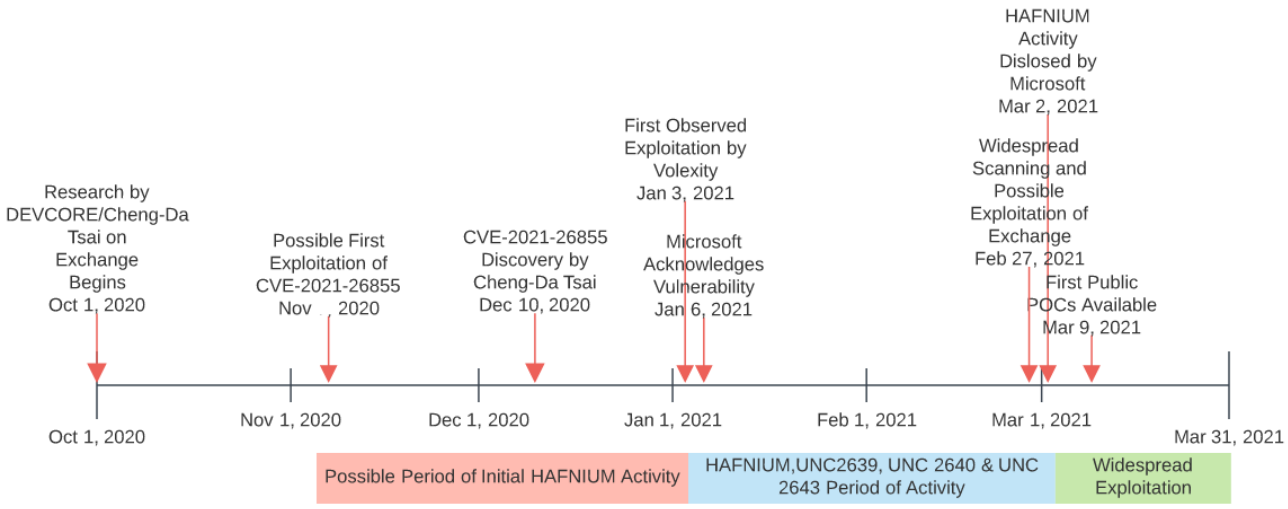

Furthermore, multiple sources revealed that, while public reporting indicated initial exploitation in tracked campaigns started in January 2021, such activity may extend as far back as November 2020. Given that public timelines from DEVCORE indicated research and analysis only began in October 2020 with vulnerability discovery in December 2020, the possibility of public exploitation in late 2020 raises a number of questions—sadly none of which can be answered with current evidence.

Irrespective of when public exploitation of CVE-2021-26855 started, based on the spike in scanning activity identified just prior to Microsoft's announcement and subsequent activity, operations appear to have increased rapidly after disclosure. Within days of Microsoft's announcement and corresponding blogs from DEVCORE, Volexity, and others, various public Proof of Concepts (POCs) appeared from independent researchers and security firms starting 09 March 2021. If not already the case previously, exploitation of CVE-2021-26855 and related vulnerabilities in the exploit chain took off such that multiple entities—from opportunistic state-sponsored organizations through likely criminal elements—are actively looking for and taking advantage of these security issues.

Webshell Payloads

Regardless of how or when malicious actors learned of these vulnerabilities in Microsoft Exchange and began using them, through at least 09 March 2021 adversary actions remained relatively static: either leverage process execution to gather system information and dump credentials and other items for memory; or utilize the exploitation chain to install a webshell on victim Exchange instances. The former utilized a number of common, publicly-available (and even legitimate) tools such as ProcDump, Covenant, and Nishang. The latter, while opening up possibilities for a number of actions for webshell installation and function, stands out as the majority of observed instances across multiple vendors reflect a long-lived, well-known, essentially publicly-available framework: China Chopper.

The China Chopper webshell framework first appeared no later than 2010. Since China Chopper's discovery, researchers linked the tool to operations from a variety of entities, ranging from state-sponsored espionage campaigns through cyber criminal elements. Although historical China Chopper use is associated with threats physically located in China, subsequent disclosures and widespread availability mean, as noted by researchers at Cisco Talos, that:

> *"This web shell is widely available, so almost any threat actor can use [it]. This also means it's nearly impossible to attribute attacks to a particular group using only [the] presence of China Chopper as an indicator."*

While initial access vectors to victims included the exploitation of four zero day vulnerabilities until disclosure on 02 March 2021, this activity concluded with deployment of a commodity, widely known webshell capability. We therefore observe a significant disconnect between intrusion methodologies (technically complex and non-public) and follow-on actions on objective (use of widely known, commodity tools). Although adversaries are ultimately judged on how successful their operations are, and this particular campaign appears to be very successful, as opposed to their technical complexity, this divergence between access and entrenchment capabilities is nonetheless curious.

More significantly still, as pointed out by various researchers examining intrusion data, China Chopper deployments linked to Exchange exploitation are not uniform. Such observations strongly indicate that more than one adversary—likely operating independently of each other—is associated with Exchange exploitation operations.

A Note on Attribution

Initial reporting from Microsoft noted that HAFNIUM is "state-sponsored and operating out of China, based on observed victimology, tactics and procedures." While the statement notes operations *out* of China and that the entity is assessed to be "state-sponsored," the sentence as constructed does not explicitly make the claim that HAFNIUM is a Chinese state-directed operation. Yet despite the very careful wording in Microsoft's blog, multiple media reports quickly made the direct link to China. While such a link is certainly possible and has not been ruled out, as of this writing no conclusive evidence has emerged linking HAFNIUM operations to the People's Republic of China (PRC).

Yet HAFNIUM is far from the only entity assessed to be targeting this vulnerability. Independent reporting from FireEye indicates at least three clusters—referred to as UNC2639, UNC2640, and UNC2643—actively targeting at least CVE-2021-26855 if not the complete exploitation chain since January 2021, without specifying links to any known threat actors or state interest. However, subsequent public comments from Kevin Mandia, CEO of FireEye, to the Associated Press indicated "two groups of Chinese state-backed hackers...installed backdoors known as 'web shells' on an as-yet undetermined number of systems." As of this writing, DomainTools is not aware if this is a revision to FireEye's earlier technical reporting.

In addition to FireEye, multiple security firms identified multiple actors exploiting these vulnerabilities. Security company Red Canary noted two distinct clusters separate from HAFNIUM behaviors, including one labeled Sapphire Pigeon active since 05 March 2021, along with other activity that could not be clustered based on limited evidence. Antivirus vendor ESET noted an astounding 10 separate groups targeting the Exchange vulnerabilities in their telemetry, including nine cases overlapping with existing threat groups and one cryptocurrency mining campaign.

While none of the reports beyond Microsoft and public comments from FireEye leadership link identified activity to China, it is worth noting that several of the groups identified in ESET's analysis have previously been linked to PRC-sponsored activity. This includes:

- Tick, also referred to as BRONZE BUTLER.
- LuckyMouse, also referred to as Emissary Panda or APT27.
- Calypso, which is assessed to have PRC-origins in some analysis.
- The "Winnti Group," although this is a wide-ranging classifier that may encompass many distinct entities.
- Tonto Team, also referred to as Karma Panda and CactusPete.

- Mikroceen APT Group, also referred to as Vicious Panda, is assessed to have PRC-origins in some analysis.

While this reporting indicates that PRC-related entities are tied to Exchange exploitation activity, ESET's analysis and telemetry shows that such activity started on 28 February 2021 at the earliest, with most entities commencing exploitation following Microsoft's public release. Given these observations, while PRC-linked entities appear to be targeting the set of vulnerabilities since disclosure, it remains unclear with any degree of certainty what entities were doing so prior to late February 2021.

Ultimately, evidence at this time only supports the following conclusions:

- CVE-2021-26855 was under active exploitation since January 2021 by multiple groups, with the possibility of some exploitation activity prior to this time.
- Since 27 February 2021 and especially following public disclosure by Microsoft on 02 March 2021, multiple additional entities have opportunistically leveraged these vulnerabilities as part of multiple, independent campaigns.
- While a number of entities linked to the Exchange exploitation activity have previously been linked to PRC-directed or -sponsored operations, multiple additional entities are also involved.
- Precise identification and origin of the initial groups targeting these vulnerabilities, including HAFNIUM and the FireEye UNC clusters, remains unavailable as of this writing.

Network Detection Possibilities

The best advice to mitigate the vulnerabilities disclosed by Microsoft is to apply the relevant patches. However, given the speed in which adversaries weaponized these vulnerabilities and the extensive period of time pre-disclosure when these were actively exploited, many organizations will likely need to shift into response and remediation activities to counter existing intrusions. Red Canary and Microsoft provided excellent guidance for host-based detection, analysis, and recovery. The remainder of this article focuses on network-specific avenues available to defenders.

As of this writing, nearly all instances of identified adversary post-exploitation activity relate to webshell deployment. One potentially easy mitigation strategy prior to patching would be to eliminate direct access to Exchange from the internet over HTTPS, a necessary condition for remote exploitation. While this would limit accessibility to services such as Outlook Web Access (OWA), such services can be provided via a Virtual Private Network (VPN) or similar portal to reduce attack surface.

While strongly recommended, defenders must also appreciate that this is a threat reduction step and not an elimination. Given the desirability of Exchange as both a source of intelligence collection itself and as an effective way to pivot throughout a victim network,

defenders should anticipate savvy attackers attempting to exploit on-premise, vulnerable Exchange deployments post-intrusion where possible as well. Therefore, patching is ultimately necessary to eliminate this intrusion operation, while webshell monitoring and defense is recommended to both counter this event as well as future security concerns.

Webshells are a difficult security problem to resolve as they take advantage of the inherent nature of the servers on which they are installed to listen for and accept remote traffic via HTTP or HTTPS. For services that must remain accessible, simply blocking these services and related connectivity is not an option. Yet defenders retain several possible avenues to detect this activity through Network Security Monitoring (NSM) and similar practices.

For externally accessible servers with known specific functionality (such as Exchange OWA), NSM looking for odd, unusual, or simply new Uniform Resource Identifiers (URIs) can alert defenders to a potential webshell. In this case, a server that should only be accepting traffic to a few narrowly-defined resources (such as the URI to reach an accessible OWA resource), can be monitored for new, unusual URIs in network traffic. Identifying successful communication to a different URI at minimum reveals a misconfiguration or potentially insecure service, and at worst can identify functionality put into place by an attacker.

Another NSM possibility focuses on follow-on lateral movement or expansion from initial access on an Exchange (or other) server. Typical server functionality would indicate receiving and responding to significant traffic, but not normally initiating connections to clients within the network. Provided an intruder desires to move beyond their initial point of access, they will need to transition to other hosts. Identifying anomalous traffic flows from servers can indicate a potentially compromised host and an intruder attempting to move deeper into the victim network.

Finally, rapid enrichment and analysis of source traffic to servers may be able to identify suspicious or anomalous connections. While not especially useful for services designed for general public access (such as a web server), this approach may work reasonably well with more circumscribed items such as mail services. Depending on scope and geographic reach, organizations can identify typical source Autonomous System Numbers (ASNs) or Internet Service Providers (ISPs) for legitimate connectivity. Using this list as a baseline, defenders can then monitor for connections from new or unique ASNs, ISPs, or hosting providers. For example, in the case of the Exchange exploitation activity, multiple vendors reported use of Virtual Private Servers (VPSs) from providers such as DigitalOcean (see Appendix). Identifying traffic to an Exchange server or similar service from a VPS node would likely be anomalous compared to traffic from typical, legitimate user activity.

Overall, a combination of visibility and information enrichment can be applied to gain greater insight into network traffic and external connectivity, while potentially revealing malicious behaviors such as webshell installation or communication. Defenders are

cautioned that none of the above approaches are universal in scope or applicability, and would require a combination of testing, baselining, and similar evaluation to avoid implementing detection or alerting logic which may lead to significant false positives.

Conclusion

The rapid expansion in Microsoft Exchange exploitation is extremely concerning for a variety of organizations using this software. Starting with narrowly tailored targeting in January 2021 (and possibly earlier), activity exploded from late February onward as an increasing number of threats learned about and either developed or gained access to exploit code. Based on this rapid expansion in activity, threat attribution and similar evaluation will be difficult if not impossible, especially as public POCs become available for widespread use.

While concerning, defenders are not completely without recourse in this situation. A combination of timely patching, attack surface reduction, and active threat hunting within environments can be applied to reduce the likelihood of intrusion and identify potential breaches that have already taken place. Although certainly not easy, given the scale and rapid expansion of Exchange exploitation, organizations running such software are strongly encouraged to enter into response and recovery mode now as an increasingly diverse set of threats are quickly subverting any accessible system.

Appendix: Infrastructure Linked to Exploitation Activity

IP	ISP	Location	Function	Source	Actor
103.77.192[.]219	Multibyte Info Technology Limited	HK	Exploit Source	Volexity	HAFNIUM
104.140.114[.]110	Eonix	US	Exploit Source	Volexity	HAFNIUM
104.248.49[.]97	DigitalOcean	US	Exploit Source	Various	N/A
104.250.191[.]110	PERFORMIVE	US	Exploit Source	Volexity	HAFNIUM
108.61.246[.]56	Choopa	JP	Exploit Source	Volexity	HAFNIUM
112.66.255[.]71	Chinanet	CN	Exploit Source	Various	N/A

IP	ISP	Location	Function	Source	Actor
139.59.56[.]239	DigitalOcean	IN	Exploit Source	Various	N/A
149.28.14[.]163	Choopa	US	Exploit Source	Volexity	HAFNIUM
157.230.221[.]198	DigitalOcean	US	Exploit Source	Volexity	HAFNIUM
161.35.1[.]207	DigitalOcean	US	Exploit Source	Various	N/A
161.35.1[.]225	DigitalOcean	US	Exploit Source	Various	N/A
161.35.45[.]41	DigitalOcean	GB	Exploit Source, Scanning	Swiss CERT, Rapid7	N/A
161.35.51[.]41	DigitalOcean	US	Exploit Source	Various	N/A
161.35.76[.]1	DigitalOcean	DE	Exploit Source	Various	N/A
165.232.154[.]116	DigitalOcean	US	Exploit Scanning	FireEye, Rapid7	UNC2639
167.99.168[.]251	DigitalOcean	US	Exploit Source	Volexity	HAFNIUM
167.99.239[.]29	DigitalOcean	US	Exploit Source	Various	N/A
182.18.152[.]105	CtrlS Datacenters Ltd	IN	Unknown	FireEye	UNC2639
185.250.151[.]72	Innovation IT	US	Exploit Source	Volexity	HAFNIUM
188.166.162[.]201	DigitalOcean	DE	Exploit Source	Various	N/A
192.81.208[.]169	DigitalOcean	US	Exploit Source	Volexity	HAFNIUM
194.87.69[.]35	LLC Baxet	RU	Webshell C2	Rapid7	N/A

IP	ISP	Location	Function	Source	Actor
203.160.69[.]66	China Unicom	HK	Exploit Source	Volexity	HAFNIUM
211.56.98[.]146	Korea Telecom	KR	Exploit Source	Volexity	HAFNIUM
45.77.252[.]175	Choopa	SG	Exploit Source	Various	N/A
5.2.69[.]14	The Infrastructure Group	NL	Exploit Source	Various	N/A
5.254.43[.]18	Voxility	US	Exploit Source	Volexity	HAFNIUM
77.61.36[.]169	KPN	NL	Exploit Source	Various	N/A
80.92.205[.]81	Innovation IT	US	Exploit Source	Volexity	HAFNIUM
86.105.18[.]116	WorldStream	NL	Unknown	FireEye	UNC2643
89.34.111[.]11	23Media	DE	Unknown	FireEye	UNC2643
91.192.103[.]43	Datasource	CH	Exploit Source	Various	N/A