



Table of Contents


Executive Summary	4	Targeted Attacks	19
Expanding Knowledge by Sharing Intrusion Realities	5	Threat Groups	20
By the Numbers	8	Malware	21
Data From FireEye Mandiant Investigations	9	Threat Techniques	28
Detection by Source	10		
Dwell Time	11		
Industry Targeting	17		



FIREEYE MANDIANT SERVICES | SPECIAL REPORT

Ransomware	36	UNC2452	62
Ransomware Evolves Into Multifaceted Extortion	37	Mapping UNC2452 Activity to the Targeted Attack Lifecycle Framework	63
Steps toward Proactive Hardening Against Ransomware in Multiple Environments	41	Case Studies	72
Recovery and Reconstitution Challenges in Post-Ransomware Scenarios	45	Insider Threat Risks to Flat Environments	73
Newly Named Threat Groups	49	Red Team Makes the Most of Social Engineering and System Misconfigurations	76
FIN11	50	Conclusion	80
Pandemic-Related Threats	56	More Security Awareness to Build Best Practices	81
Threats Against Organizations Working with COVID-19 Information and Research	57		

EXECUTIVE SUMMARY



Expanding Knowledge by Sharing Intrusion Realities



Security practitioners faced a series of challenges in this past year which forced organizations into uncharted waters. As ransomware operators were attacking state and municipal networks alongside hospitals and schools, a global pandemic response to COVID-19 necessitated a move to remote work for a significant portion of the economy. Organizations had to adopt new technologies and quickly scale outside of their normal growth plans.

As organizations settled into a new understanding of “normal,” UNC2452, a suspected nation-state threat actor, conducted one of the most advanced cyber espionage campaigns in recent history. Many security teams were forced to suspend wide-ranging analyses around the adoption of remote work policies and instead focus on a supply chain attack from a trusted platform.

Nation states taking a cyber espionage approach to COVID research, threat groups working together to achieve their objectives, exploitation of quickly adopted work-from-home strategies and a wake-up call for global supply chain compromise – experiences in 2020 will shape security policies for years to come.

Themes covered in M-Trends 2021 include :

- 59% of the security incidents investigated by Mandiant last year were initially detected by the organizations themselves, an improvement of 12% from the prior year.
- Ransomware has evolved into multifaceted extortion where actors not only deploy ransomware encryptors across victim environments, but also employ a variety of other extortion tactics to coerce victims into complying with demands.
- FIN11, a recently named financially motivated threat group, was responsible for widespread phishing campaigns, that conducted several multifaceted extortion operations.
- Pervasive ransomware campaigns drove down the median dwell time as threat actors sought to capitalize on shifting trends in the workspace and a global crisis.

- UNC2452, a suspected state-sponsored group, undertook a broad-scale espionage campaign after injecting a trojanized DLL into the SolarWinds Orion build process. Mandiant identified the campaign and worked with law enforcement agencies and industry partners to protect organizations and respond to the adversary.
- Mandiant experts observed the use of 63% of MITRE ATT&CK techniques, and just over a third of techniques observed were seen in more than 5% of intrusions.
- Threat actors took advantage of infrastructure supporting work-at-home with an increased focus on vulnerability exploitation.

One of the most striking trends for the period of October 1st, 2019 to September 30th, 2020 was the significant reduction in the global median dwell time. At 24 days, this is the first time Mandiant has observed the global median dwell time dip below one month. While this reduction in dwell time may correlate to better visibility and response, it is also likely the preponderance of ransomware helped drive down the time between initial infection and identification.

With the inclusion of all the observations listed above, the addition of new metrics reported in By The Numbers, the introduction of the named threat group FIN11, new case studies, and many other topics, M-Trends 2021 builds on our transparency to continue providing critical knowledge to those tasked with defending organizations. The information in this report has been sanitized to protect identities of victims and their data.

BY THE NUMBERS



Data from FireEye Mandiant Investigations



The metrics reported in M-Trends 2021 are based on FireEye Mandiant investigations of targeted attack activity conducted between October 1, 2019 and September 30, 2020.



Internal detection is when an organization independently discovers it has been compromised.

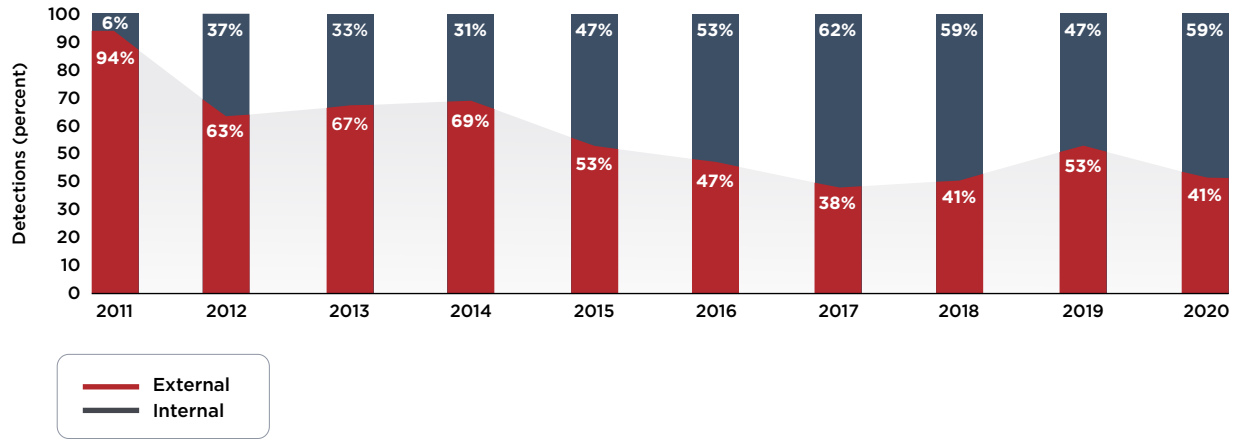


External notification is when an outside entity informs an organization it has been compromised.


Detection by Source

Organizations continue to improve their ability to discover compromises within their environments. While M-Trends 2020 noted a drop in internal notifications for 2019 compared to 2018, Mandiant experts observed a return to organizations detecting the majority of incidents internally in 2020. Organizations increased internal incident detection to 59% in 2020—a 12-point increase compared to 2019. This return to organizations detecting the majority of intrusions within their environments is in line with the overall trend towards increased internal detection observed over the last decade. It shows a continued dedication to the expansion and enhancement of organic detection and response capabilities. The increase in ransomware activity affects this category as well.

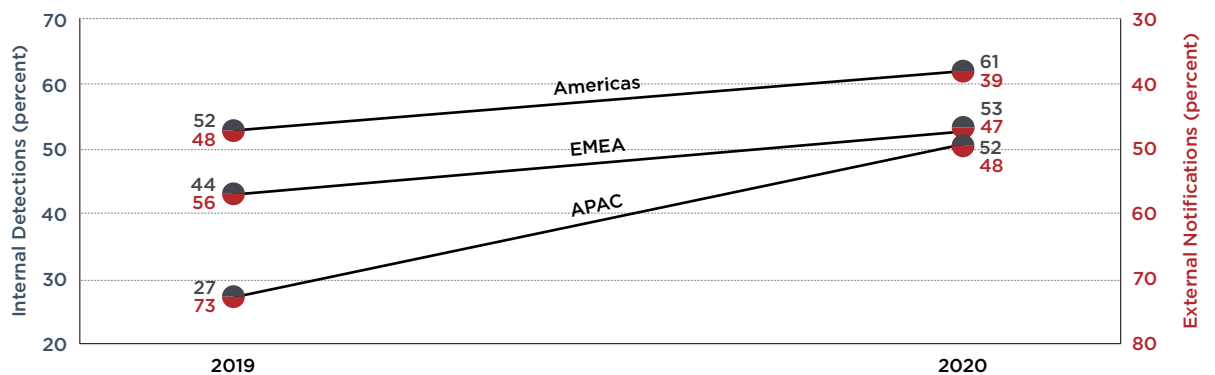
DETECTION BY SOURCE, 2011-2020



DETECTION BY SOURCE BY REGION, 2020



DETECTION BY SOURCE BY REGION, 2019-2020 COMPARISON





Dwell time is calculated as the number of days an attacker is present in a victim environment before they are detected. The median represents a value at the midpoint of a data set sorted by magnitude.

Dwell Time

Organizations continue to find and contain adversaries faster than in previous years. Over the past decade, there has been a marked reduction in median dwell time, from just over one year (2011) to just under one month (2020).



Median Dwell Time



Global Dwell Time

In 2020, the global median dwell time dropped below one month for the first time. Organizations are now detecting incidents in only 24 days—more than twice as fast as 2019. These improvements in detection hold true regardless of the notification source. Global median dwell time for incidents which were detected internally dropped to just 12 days and incidents with external notification sources came in at 73 days.

GLOBAL MEDIAN DWELL TIME, 2011-2020


Compromise Notifications	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
All	416	243	229	205	146	99	101	78	56	24
External Notification	—	—	—	—	320	107	186	184	141	73
Internal Detection	—	—	—	—	56	80	57.5	50.5	30	12


Global Median Dwell Time Distribution

Globally, organizations are detecting more incidents within the first 30 days of an intrusion and fewer incidents with a dwell time longer than 700 days. The distribution of global dwell time continues to show an increased proportion of incidents with a dwell time of 30 days or fewer. In 2020, 52% of the compromises investigated by Mandiant experts had dwell times of 30 days or fewer, compared to 41% in 2019 and 31% in 2018. There were also improvements at the other end of the spectrum; Mandiant observed a 3% decrease in investigations with dwell times greater than 700 days.

The overall trends across multiple years could be explained by continued development and improvement of organizational detection capabilities and an evolution of the threat landscape.

GLOBAL MEDIAN DWELL TIME DISTRIBUTION, 2018-2020







Americas Median Dwell Time


The Americas saw median dwell time continue to decrease in 2020. The dwell time for incidents which were discovered internally improved the most—from 32 days down to nine days. This is the first time Mandiant has observed the median dwell time in any region dip into single digits.

Median dwell time in the Americas was 3.5 times shorter in 2020 than in 2019. Companies were detecting incidents internally 3.6 times faster and receiving external notification of compromises 2.1 times faster.

In 2020, 27.5% of incidents investigated in the Americas involved ransomware. The large number of investigations which involved ransomware undoubtedly drove down the median dwell time. Ransomware incidents in the Americas had a median dwell time of just three days and accounted for 41% of incidents with a dwell time of 14 days or fewer.

AMERICAS MEDIAN DWELL TIME, 2016-2020



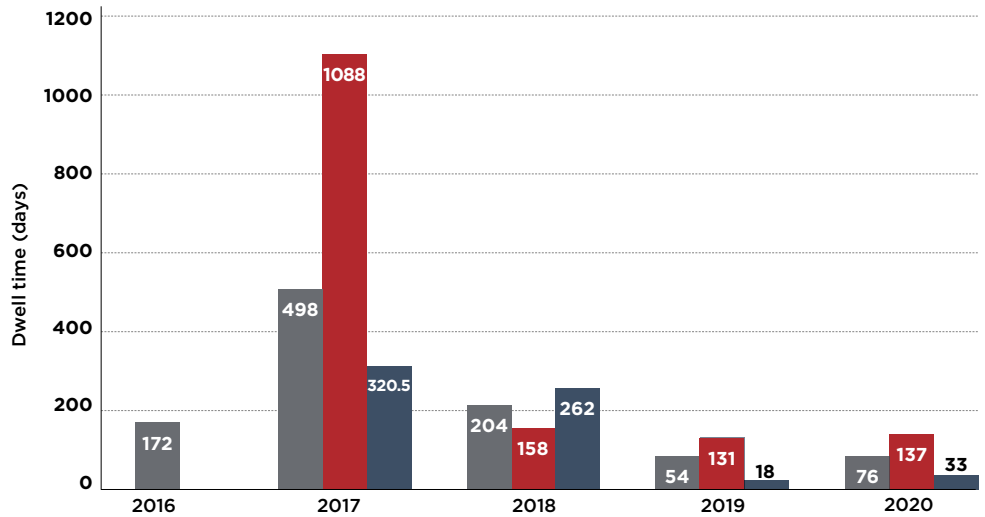




APAC Median Dwell Time

The median dwell time for APAC increased from 54 days in 2019 to 76 days in 2020. APAC saw a decrease in the number of ransomware-related breaches which accounted for 12.5% of incidents investigated in 2020 as compared to 18% in 2019. The reduction in ransomware-related incidents was a likely contributor to the overall increase in median dwell time for APAC.

Adversaries continue to maintain access in compromised organizations in APAC for extensive periods of time. Consistent with observations in 2019, 10% of breaches investigated in APAC during 2020 showed dwell times of more than three years and 4% were greater than nine years.

APAC MEDIAN DWELL TIME, 2016-2020



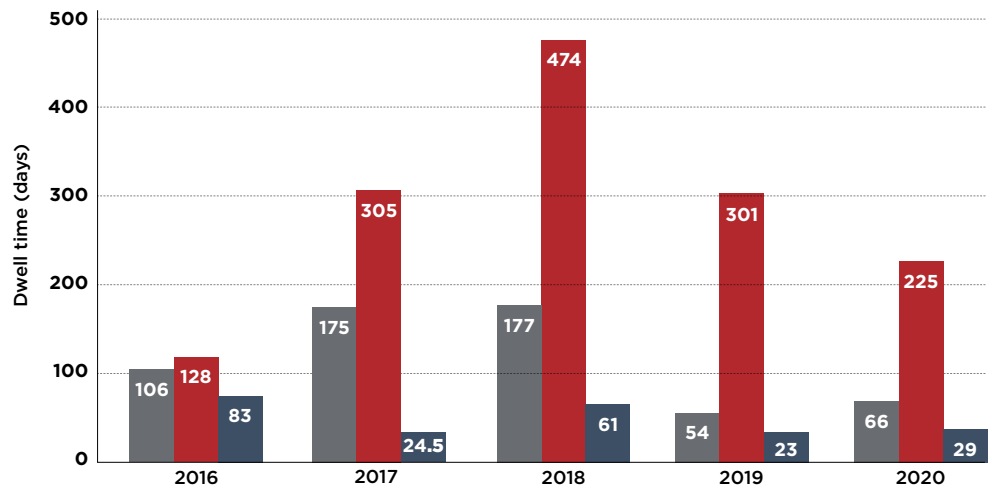


EMEA Median Dwell Time

The median dwell time for EMEA increased from 54 days in 2019 to 66 days in 2020. Mandiant experts observed that 28% of incidents in EMEA had a dwell time of one week or less, and 8% of incidents had dwell times longer than three years. Organizations in EMEA continue to respond to long-standing intrusions while also contending with faster paced compromises such as ransomware.

When separated by notification source, median dwell time for EMEA increased for incidents discovered internally but decreased when companies were notified of a compromise by an external entity. For incidents that were detected internally, EMEA saw median dwell time increase by 20%, from 23 days in 2019 to 29 days in 2020. Conversely, compromises in EMEA with an external notification source had a 25% decrease in median dwell time, from 301 days in 2019 to 225 days in 2020.

EMEA MEDIAN DWELL TIME, 2016-2020



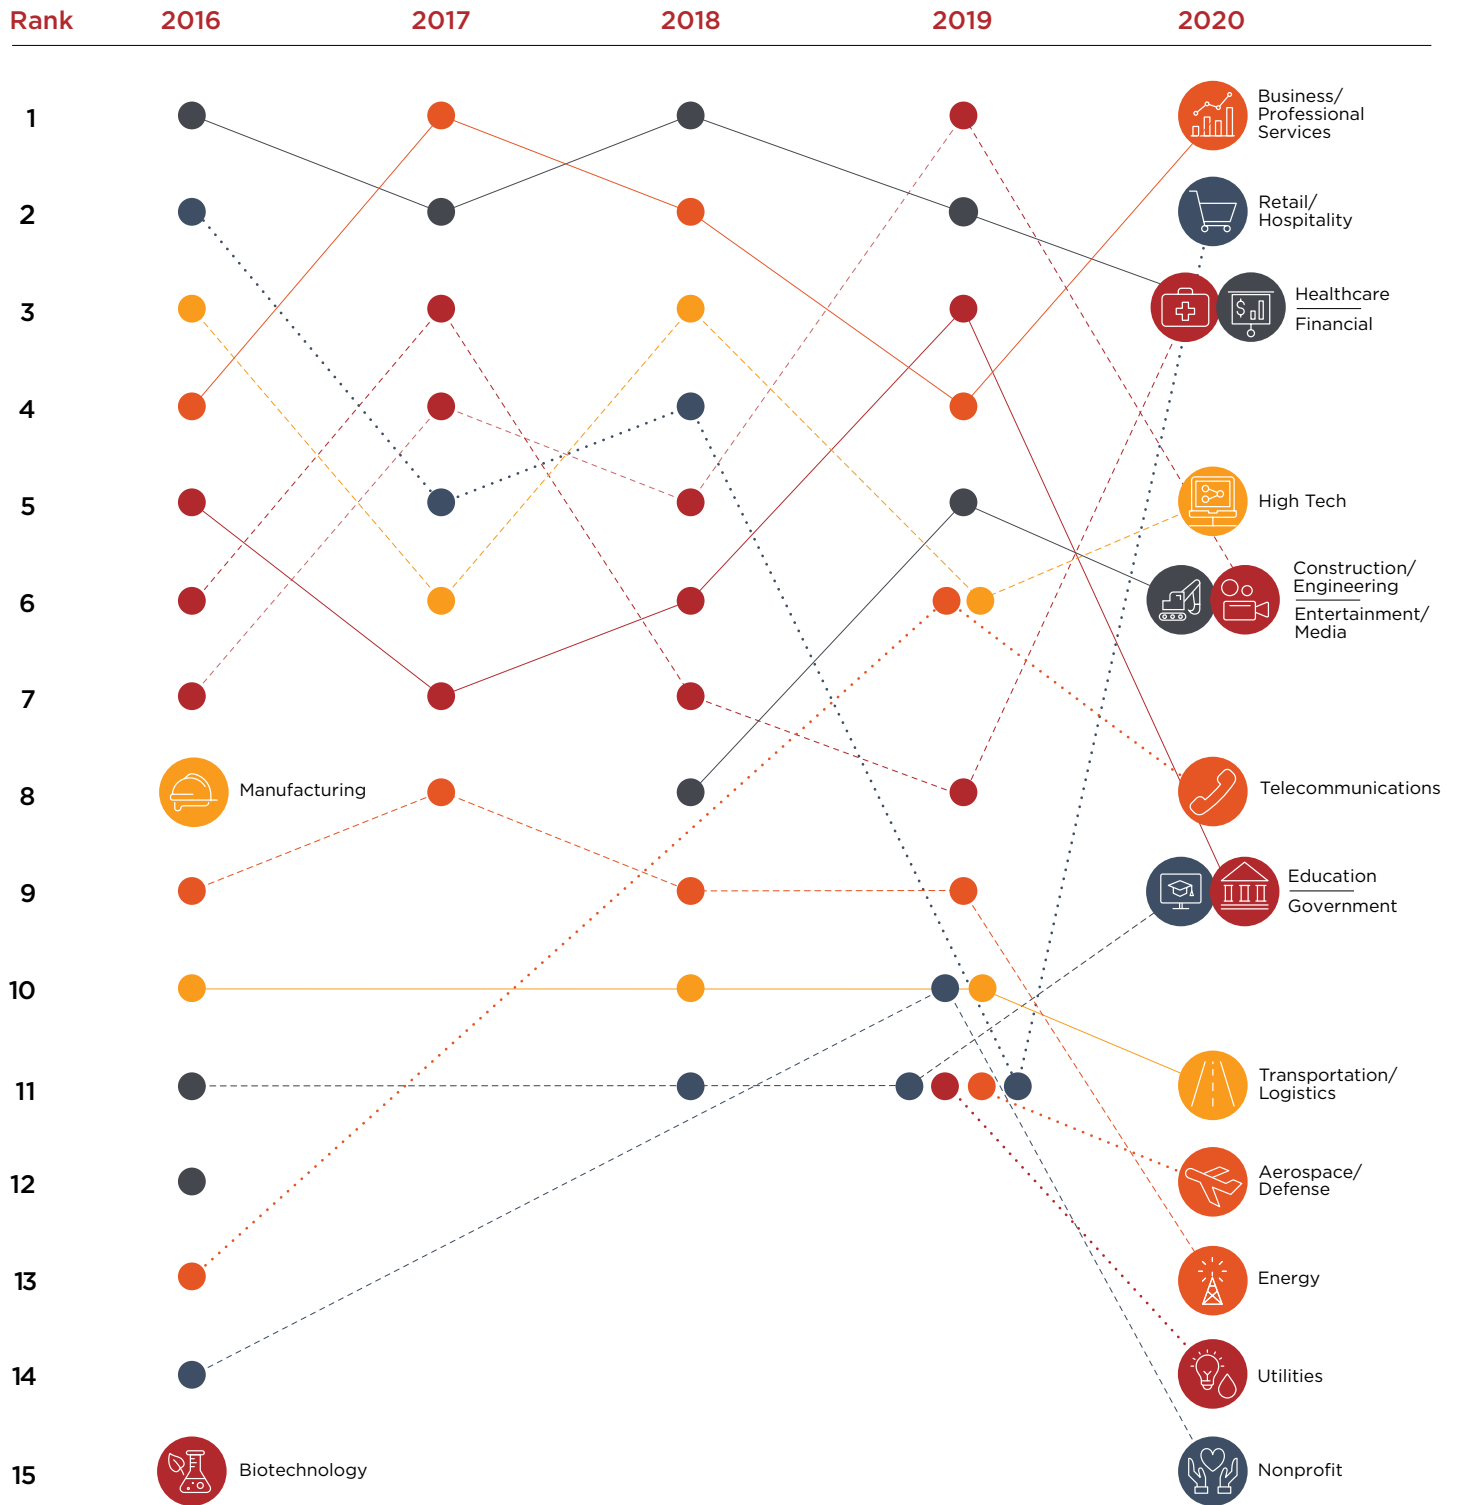
Industry Targeting

Mandiant has observed that the most targeted industries continue to remain consistent year over year. The top five most targeted industries in 2020 were business and professional services, retail and hospitality, financial, healthcare, and high technology. Over the past decade, business and professional services and financial have consistently placed in the top five most targeted industries. Overall, the top targeted industries change little while position in the rankings is somewhat fluid.

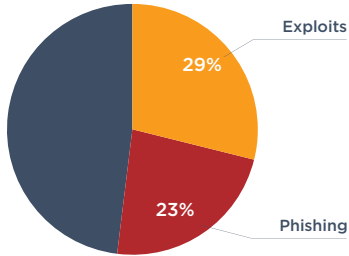
Big Movers

Mandiant experts observed that retail and hospitality organizations were targeted more heavily in 2020, coming in as the second most targeted industry, compared to 11th in 2019. Healthcare rose to 3rd most targeted industry in 2020, compared to 8th in 2019. In the other direction, Mandiant experts observed a decrease in targeting of entertainment and media which dropped from the most targeted industry in 2019 to 6th in 2020.


TARGETED INDUSTRIES, 2015-2020




Initial Infection Vector
(when identified)



Objective: Financial Gain



Objective: Data Theft



Multiple Threat Groups Identified
(per environment)



Targeted Attacks

Mandiant experts responded to a wide variety of intrusions in 2020, making observations about initial infection vectors, adversary operations and victim environments.

Initial Infection Vector

While phishing remains an effective vector for initial compromise, in 2020, Mandiant observed adversaries leveraging exploits more often than other vectors. In cases where the initial vector of compromise was identified, evidence of exploits was found in 29% of intrusions whereas phishing accounted for 23% of intrusions. Mandiant experts also observed adversaries used stolen credentials or brute forcing as the initial attack vector in 19% of the investigations. Prior compromise accounted for 12% of the intrusions in which the initial compromise was identified.

Adversary Operations

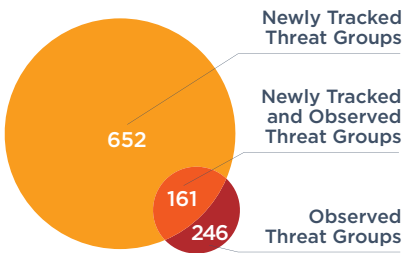
Adversaries continue to use intrusions for monetary gain through methods that include extortion, ransom, payment card theft and illicit transfers. Direct financial gain was the likely motive for 36% of intrusions and an additional 2% of intrusions were likely perpetrated to resell access.

In 2020, data theft remained an important mission objective for threat actors. In 32% of intrusions adversaries stole data and in 29% of those cases (9% of all cases) the data theft likely supported intellectual property or espionage end goals.

Approximately 3% of intrusions likely only served to compromise architecture for further attacks, and insider threats remain rare, represented by fewer than 1% of intrusions.

Environment

In 29% of cases, Mandiant experts identified more than one distinct threat group in the victim environment—nearly twice the percentage noted in 2019.

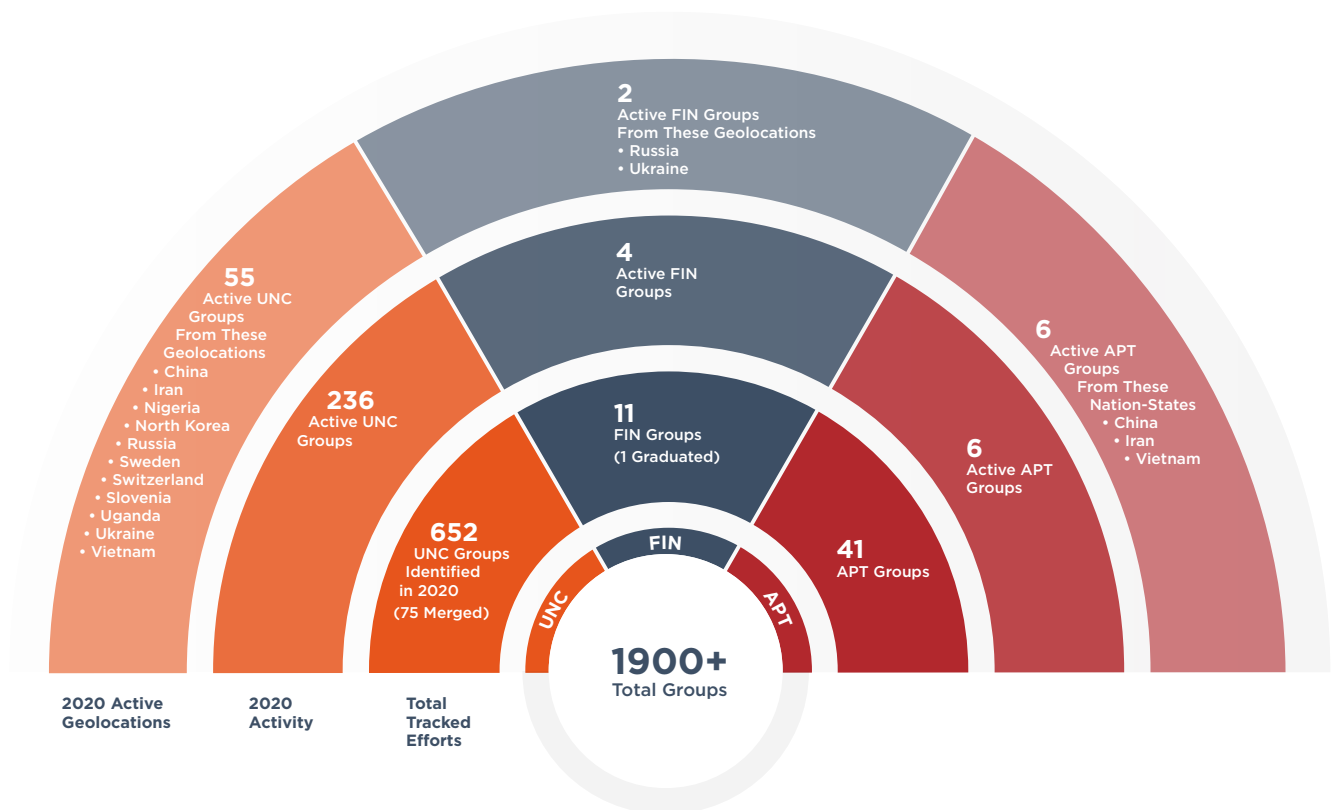


Threat Groups

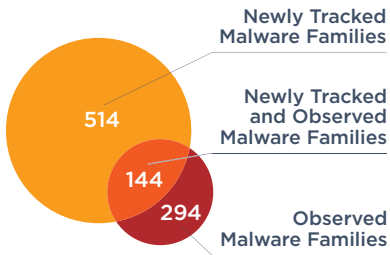
Over the course of Mandiant’s history, Mandiant experts have tracked more than 2,400 threat groups, which includes 650+ newly tracked threat groups in 2020. Mandiant experts have combined or eliminated approximately 500 of these groups over the years, leaving more than 1,900 distinct threat groups tracked at this time. By expanding and refining a vast threat actor knowledgebase, Mandiant can support a broad spectrum of investigations while maintaining fidelity within that dataset. In 2020, Mandiant experts graduated one group to a named threat group and merged 75 threat groups based on extensive research into activity overlaps. For details on how Mandiant defines and references UNC groups and merges, please see, “How Mandiant Tracks Uncategorized Threat Actors.”¹

In 2020, Mandiant experts investigated intrusions that involved 246 distinct threat groups. Organizations faced intrusions by four named financial threat (FIN) groups; six named advanced persistent threat (APT) groups, including groups from the nation-states of China, Iran and Vietnam; and 236 uncategorized threat (UNC) groups. Of the 246 threat groups observed at intrusion clients, 161 of these threat groups were newly tracked threat groups in 2020.

THREAT GROUPS, 2020



1. FireEye (December 17, 2020). [DebUNCing Attribution: How Mandiant Tracks Uncategorized Threat Actors.](#)



A malware family is a program or set of associated programs with sufficient “code overlap” among the members that Mandiant considers them to be the same thing, a “family”. The term family broadens the scope of a single piece of malware as it can be altered over time, which in turn creates new, but fundamentally overlapping pieces of malware.

Malware

Mandiant continually expands its knowledgebase of malware families based on insights gained from frontline Mandiant investigations, public reporting, information sharing and other research. In 2020, Mandiant began tracking more than 500 new malware families. This is on par with the number of newly tracked malware families compared to the previous year.

Mandiant responds to hundreds of diverse intrusions each year where adversaries provide organizations with unique challenges. In 2020, Mandiant experts observed 294 distinct malware families in use during investigations into compromised environments. Of the nearly 300 malware families observed by Mandiant experts during intrusions, 144 were malware families which Mandiant began tracking in 2020. Adversaries not only use established malware but also continue to innovate and adapt to be effective in victim environments.




A malware category describes a malware family's primary purpose. Each malware family is assigned only one category that best describes its primary purpose, regardless if it has functionality for more than one category.

Malware Families by Category

The malware category distribution remains relatively consistent year over year. Of the 514 newly tracked malware families in 2020, the top five categories were backdoors (36%), downloaders (16%), droppers (8%), launchers (7%) and ransomware (5%).

Malware category	Primary purpose
Backdoor	A program whose primary purpose is to allow a threat actor to interactively issue commands to the system on which it is installed.
Credential Stealer	A utility whose primary purpose is to access, copy or steal authentication credentials.
Downloader	A program whose sole purpose is to download (and perhaps launch) a file from a specified address, and which does not provide any additional functionality or support any other interactive commands.
Dropper	A program whose primary purpose is to extract, install and potentially launch or execute one or more files.
Launcher	A program whose primary purpose is to launch one or more files. Differs from a dropper or an installer in that it does not contain or configure the file, but merely executes or loads it.
Ransomware	A program whose primary purpose is to perform some malicious action (such as encrypting data), with the goal of extracting payment from the victim in order to avoid or undo the malicious action.
Other	Includes all other malware categories such as utilities, keyloggers, point of sale (POS), tunnelers and data miners.

NEWLY TRACKED MALWARE FAMILIES BY CATEGORY, 2020






An observed malware family is a malware family identified during an investigation by Mandiant experts.

Observed Malware Families by Category

Backdoors are a mainstay for adversaries and consistently comprise the largest malware family category observed during investigations. Mandiant experts observed that attackers deployed at least one backdoor in more than half of the intrusions investigated. Of the 294 malware families observed in 2020, the top five categories were backdoors (41%), downloaders (9%), droppers (9%), ransomware (8%) and launchers (6%).

OBSERVED MALWARE FAMILIES BY CATEGORY, 2020



Newly Tracked Malware Families by Availability

Mandiant experts observed that 81% of newly tracked malware families were non-public whereas 19% were publicly available. While adversaries do use publicly available tools and code, the majority of malware families tracked were likely privately developed or their availability is restricted.


NEWLY TRACKED MALWARE FAMILIES BY AVAILABILITY, 2020



A publicly available tool or code family is readily obtainable without restriction. This includes tools that are freely available on the Internet, as well as tools that are sold or purchased, as long as they can be purchased by any buyer.



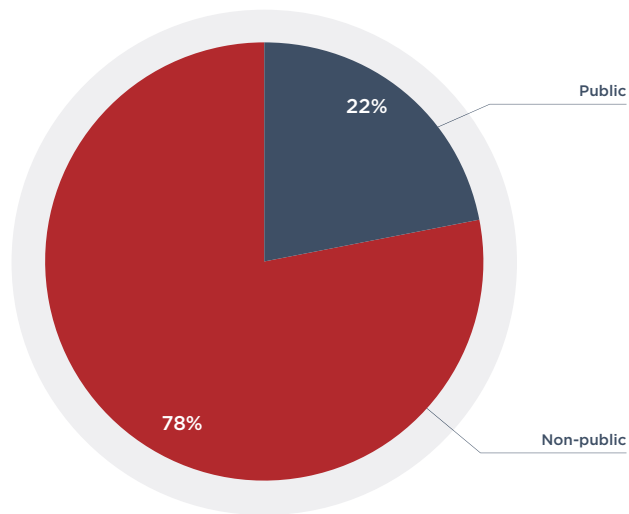
A non-public tool or code family is, to the best of our knowledge, not publicly available (either for free or for sale). They may include tools that are privately developed, held or used, as well as tools that are shared among or sold to a restricted set of customers.



Observed Malware Families by Availability

Similar to the availability for newly tracked malware families, in 2020, Mandiant experts observed that 78% of malware families used by adversaries during an intrusion were non-public and 22% were publicly available.

OBSERVED MALWARE FAMILIES BY AVAILABILITY, 2020

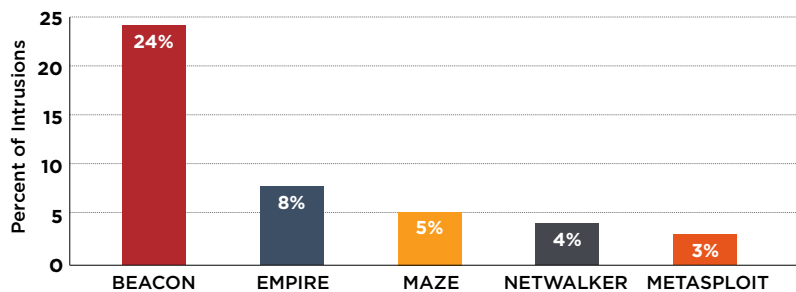


Most Frequently Seen Malware Families

The top five malware families seen most frequently during intrusions investigated by Mandiant experts were BEACON, EMPIRE, MAZE, NETWALKER, and METASPLOIT. BEACON was so prevalent in 2020 that it was observed at nearly a quarter of all the intrusions Mandiant investigated. Mandiant experts also observed a lack of cross-pollination with respect to the malware used across incidents.

Just 3.4% of malware families seen during an intrusion were observed at 10 or more intrusions, and 70% percent of malware families seen were only observed during a single intrusion.

MOST FREQUENTLY SEEN MALWARE FAMILIES, 2020



- **BEACON** is a backdoor that is commercially available as part of the Cobalt Strike software platform and commonly used for pen-testing network environments. The malware supports several capabilities, such as injecting and executing arbitrary code, uploading and downloading files and executing shell commands. Mandiant has seen BEACON used by a wide range of named threat groups including APT19, APT32, APT40, APT41, FIN6, FIN7, FIN9 and FIN11, as well as nearly 300 UNC groups.
- **EMPIRE** is a publicly available PowerShell post-exploitation framework that allows users to run PowerShell agents without the use of powershell.exe. PowerShell Empire also allows actors to run various types of post-exploitation modules and make adaptable communications while evading detection. Mandiant experts track 90 threat groups that have utilized EMPIRE including APT19, APT33, FIN10, FIN11 and 86 UNC Groups.
- **MAZE** is a ransomware family that encrypts files stored locally and on network shares. MAZE can be configured to infect remote and removable drives as well as send basic system information via HTTP. Mandiant has observed a dozen distinct financially motivated threat groups leverage MAZE ransomware.
- **NETWALKER** is a ransomware family capable of deleting volume shadow copies and encrypting files on a victim host and any mapped network drives using a combination of SALSA20 and Curve25519 encryption algorithms. Mandiant tracks eight threat groups that have used NETWALKER ransomware to further their monetary end goals.
- **METASPLOIT** is a penetration testing platform that enables users to find, exploit, and validate vulnerabilities. Mandiant has seen METASPLOIT used by APT40, APT41, FIN6, FIN7, FIN11 and 40 UNC groups with end goals ranging from espionage and financial gain to penetration testing.

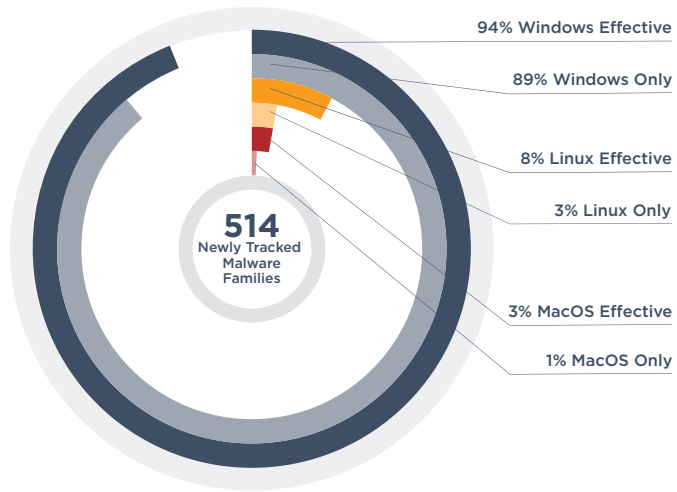


The **operating system effectiveness** of a malware family is the operating system(s) that the malware can be used against.

Operating System Effectiveness

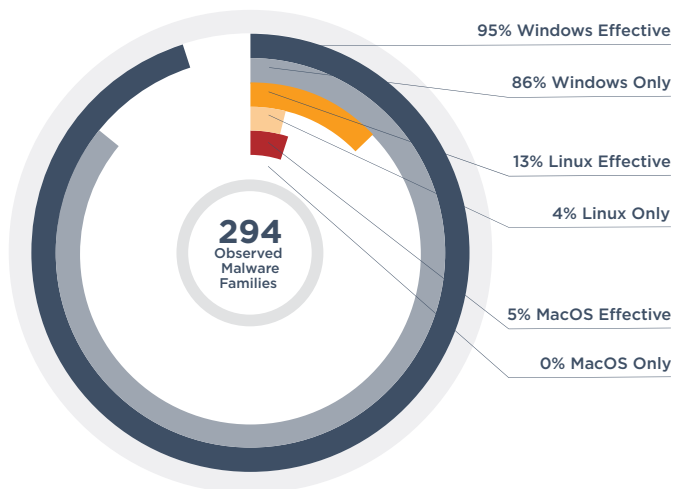
In keeping with previous trends, the majority of newly tracked malware families were effective on Windows. Only 8% and 3% of newly tracked malware families were effective on Linux and MacOS, respectively.

EFFECTIVENESS OF NEWLY TRACKED MALWARE FAMILIES BY OPERATING SYSTEM, 2020



Similar to trends seen for newly tracked malware families, the majority of malware families observed during Mandiant investigations were effective on Windows. Malware effective on Linux and MacOS was also observed but accounted for only 13% and 5% of malware families, respectively.

EFFECTIVENESS OF OBSERVED MALWARE FAMILIES BY OPERATING SYSTEM, 2020





MITRE ATT&CK® is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations. The ATT&CK knowledge base is used as a foundation for the development of specific threat models and methodologies in the private sector, government and the cyber security product and service community.

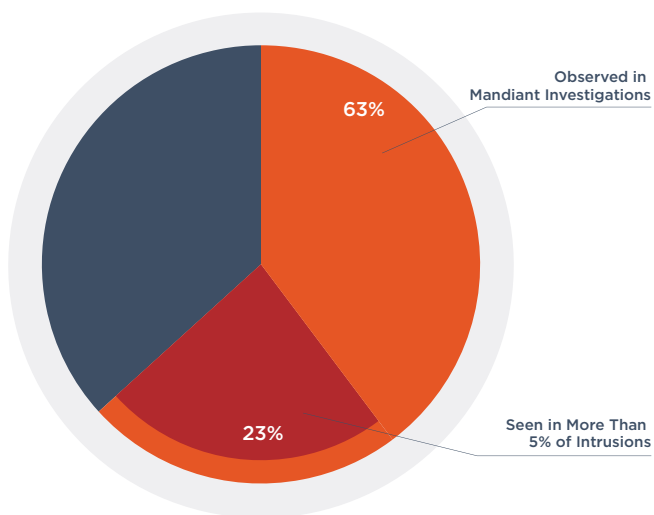
Threat Techniques

Mandiant continues to support community and industry efforts by mapping its findings to the MITRE ATT&CK framework. In 2020, significant changes were made to the MITRE ATT&CK framework with the introduction of sub-techniques and the incorporation of PRE-ATT&CK in Enterprise ATT&CK. Due in part to these changes and the continued refinement of its data model, Mandiant now has MITRE ATT&CK techniques mapped to more than 1800 Mandiant techniques and subsequent findings.

When making security decisions, organizations must consider the likelihood of specific techniques being used during an intrusion. In 2020, Mandiant experts observed attackers use 63% of MITRE ATT&CK techniques and 24% of sub-techniques. However, only 37% of the techniques observed (23% of all techniques) were seen in more than 5% of intrusions.

In more than half of the intrusions investigated in 2020, Mandiant observed that adversaries used obfuscation, such as encryption or encoding, on files or information to make detection and subsequent analysis more difficult (T1027). Adversaries regularly used a command or scripting interpreter to further intrusions (T1059) and 80% of those cases involved the use of PowerShell (T1059.001). System services (T1569) were also a popular execution method, represented in 31% of intrusions, all of which used Windows services (T1569.002). Adversaries also used Remote Services (T1021) to further intrusions, with 88% of those using the Remote Desktop Protocol (T1021.001). Adversaries often take advantage of what is available in a victim's environment; this tendency is highlighted by how frequently adversaries used PowerShell, Windows services and Remote Desktop.

MITRE ATT&CK TECHNIQUES USED MOST FREQUENTLY, 2020



FREQUENTLY TARGETED TECHNOLOGIES, 2020

88% Remote Desktop Protocol (T1021.001) for intrusions using **remote services (T1021)**

Used in 25% of all intrusions

100% Windows services (T1569.002) for intrusions using **system services (T1569)**

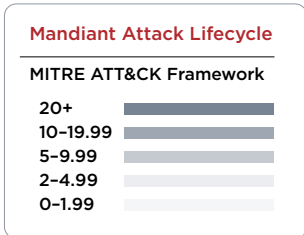
Used in 31% of all intrusions

80% PowerShell (T1059.001) for intrusions using **command or scripting interpreter (T1059)**

Used in 41% of all intrusions

MITRE ATT&CK TECHNIQUES RELATED TO ATTACK LIFECYCLE, 2020

Initial Reconnaissance				Initial Compromise			
Reconnaissance				Initial Access			
T1595: Active Scanning	0.2%			T1190: Exploit Public-Facing Application	21.0%		
Resource Development				T1566: Phishing	14.2%	T1566.001: Spearphishing Attachment	8.1%
T1588: Obtain Capabilities	21.3%	T1588.003: Code Signing Certificates	21.0%			T1566.002: Spearphishing Link	7.1%
T1583: Acquire Infrastructure	7.8%	T1583.003: Virtual Private Server	7.8%			T1566.003: Spearphishing via Service	0.5%
T1584: Compromise Infrastructure	5.1%			T1133: External Remote Services	11.5%		
T1587: Develop Capabilities	1.2%	T1587.003: Digital Certificates	1.2%	T1078: Valid Accounts	6.8%		
				T1199: Trusted Relationship	3.2%		
				T1189: Drive-by Compromise	1.5%		
				T1091: Replication Through Removable Media	0.5%		
				T1195: Supply Chain Compromise	0.5%	T1195.002: Compromise Software Supply Chain	0.5%
				T1200: Hardware Additions	0.5%		



Establish Foothold

Escalate Privileges

Persistence				Privilege Escalation			
T1053: Scheduled Task/Job	15.2%	T1053.005: Scheduled Task	6.6%	T1055: Process Injection	18.1%	T1055.003: Thread Execution Hijacking	1.0%
T1505: Server Software Component	12.2%	T1505.003: Web Shell	12.2%			T1055.012: Process Hollowing	0.5%
T1133: External Remote Services	11.5%			T1053: Scheduled Task/Job	15.2%	T1053.005: Scheduled Task	6.6%
T1098: Account Manipulation	9.0%			T1543: Create or Modify System Process	9.0%	T1543.003: Windows Service	9.0%
T1543: Create or Modify System Process	9.0%	T1543.003: Windows Service	9.0%	T1078: Valid Accounts	6.8%		
T1078: Valid Accounts	6.8%			T1134: Access Token Manipulation	5.9%	T1134.001: Token Impersonation/Theft	0.2%
T1136: Create Account	6.1%	T1136.001: Local Account	0.2%	T1547: Boot or Logon Autostart Execution	4.2%	T1547.001: Registry Run Keys / Startup Folder	4.2%
		T1136.002: Domain Account	0.2%			T1547.009: Shortcut Modification	0.2%
T1547: Boot or Logon Autostart Execution	4.2%	T1547.001: Registry Run Keys /Startup Folder	4.2%	T1546: Event Triggered Execution	3.2%	T1546.008: Accessibility Features	1.2%
		T1547.009: Shortcut Modification	0.2%			T1546.011: Application Shimming	1.2%
T1546: Event Triggered Execution	3.2%	T1546.008: Accessibility Features	1.2%			T1546.003: Windows Management Instrumentation Event Subscription	0.7%
		T1546.011: Application Shimming	1.2%	T1574: Hijack Execution Flow	3.2%	T1574.001: DLL Search Order Hijacking	2.4%
		T1546.003: Windows Management Instrumentation Event Subscription	0.7%			T1574.002: DLL Side-Loading	2.4%
T1574: Hijack Execution Flow	3.2%	T1574.001: DLL Search Order Hijacking	2.4%			T1574.008: Path Interception by Search Order Hijacking	0.2%
		T1574.002: DLL Side-Loading	2.4%	T1548: Abuse Elevation Control Mechanism	0.7%	T1548.002: Bypass User Account Control	0.5%
		T1574.008: Path Interception by Search Order Hijacking	0.2%			T1548.001: Setuid and Setgid	0.2%
T1197: BITS Jobs	0.7%			T1068: Exploitation for Privilege Escalation	0.2%		
T1542: Pre-OS Boot	0.2%	T1542.003: Bootkit	0.2%	T1484: Domain Policy Modification	0.2%	T1484.001: Group Policy Modification	0.2%

Internal Reconnaissance

Lateral Movement

Discovery

T1082: System Information Discovery	24.2%		
T1083: File and Directory Discovery	21.8%		
T1012: Query Registry	13.0%		
T1016: System Network Configuration Discovery	13.0%		
T1497: Virtualization/Sandbox Evasion	12.7%	T1497.001: System Checks	1.5%
T1057: Process Discovery	12.0%		
T1518: Software Discovery	11.5%		
T1033: System Owner/User Discovery	9.8%		
T1049: System Network Connections Discovery	5.4%		
T1007: System Service Discovery	4.9%		
T1482: Domain Trust Discovery	4.9%		
T1087: Account Discovery	4.2%	T1087.004: Cloud Account	0.2%
		T1087.002: Domain Account	0.2%
T1010: Application Window Discovery	2.4%		
T1069: Permission Groups Discovery	2.4%	T1069.003: Cloud Groups	0.2%
T1046: Network Service Scanning	1.7%		
T1124: System Time Discovery	1.0%		
T1018: Remote System Discovery	0.2%		
T1135: Network Share Discovery	0.2%		
T1217: Browser Bookmark Discovery	0.2%		
T1538: Cloud Service Dashboard	0.2%		
T1580: Cloud Infrastructure Discovery	0.2%		

Lateral Movement

T1021: Remote Services	28.4%	T1021.001: Remote Desktop Protocol	24.9%
		T1021.002: SMB/Windows Admin Shares	3.9%
		T1021.004: SSH	3.2%
		T1021.006: Windows Remote Management	1.0%
		T1021.003: Distributed Component Object Model	0.2%
T1091: Replication Through Removable Media	0.5%		
T1550: Use Alternate Authentication Material	0.5%	T1550.002: Pass the Hash	0.2%
		T1550.003: Pass the Ticket	0.2%
T1563: Remote Service Session Hijacking	0.5%	T1563.002: RDP Hijacking	0.2%
T1534: Internal Spearphishing	0.2%		

Maintain Persistence

Persistence			
T1053: Scheduled Task/Job	15.2%	T1053.005: Scheduled Task	6.6%
T1505: Server Software Component	12.2%	T1505.003: Web Shell	12.2%
T1133: External Remote Services	11.5%		
T1098: Account Manipulation	9.0%		
T1543: Create or Modify System Process	9.0%	T1543.003: Windows Service	9.0%
T1078: Valid Accounts	6.8%		
T1136: Create Account	6.1%	T1136.001: Local Account	0.2%
		T1136.002: Domain Account	0.2%
T1547: Boot or Logon Autostart Execution	4.2%	T1547.001: Registry Run Keys /Startup Folder	4.2%
		T1547.009: Shortcut Modification	0.2%
T1546: Event Triggered Execution	3.2%	T1546.008: Accessibility Features	1.2%
		T1546.011: Application Shimming	1.2%
		T1546.003: Windows Management Instrumentation Event Subscription	0.7%
T1574: Hijack Execution Flow	3.2%	T1574.001: DLL Search Order Hijacking	2.4%
		T1574.002: DLL Side-Loading	2.4%
		T1574.008: Path Interception by Search Order Hijacking	0.2%
T1197: BITS Jobs	0.7%		
T1542: Pre-OS Boot	0.2%	T1542.003: Bootkit	0.2%

Mission Completion

Collection			
T1560: Archive Collected Data	15.2%	T1560.001: Archive via Utility	3.4%
		T1560.002: Archive via Library	1.5%
T1056: Input Capture	4.9%	T1056.001: Keylogging	4.9%
T1213: Data from Information Repositories	4.2%	T1213.002: Sharepoint	0.2%
T1113: Screen Capture	3.2%		
T1114: Email Collection	3.2%	T1114.003: Email Forwarding Rule	1.5%
T1115: Clipboard Data	2.7%		
T1530: Data from Cloud Storage Object	0.5%		
T1074: Data Staged	0.2%		
T1123: Audio Capture	0.2%		
T1125: Video Capture	0.2%		
Exfiltration			
T1567: Exfiltration Over Web Service	0.2%		
Impact			
T1489: Service Stop	13.4%		
T1529: System Shutdown/Reboot	3.2%		
T1490: Inhibit System Recovery	2.7%		
T1486: Data Encrypted for Impact	2.2%		
T1496: Resource Hijacking	2.0%		
T1565: Data Manipulation	1.7%	T1565.001: Stored Data Manipulation	1.7%
T1531: Account Access Removal	1.0%		
T1491: Defacement	0.7%	T1491.002: External Defacement	0.7%

Across the Lifecycle

Credential Access				Command and Control			
T1003: OS Credential Dumping	8.8%	T1003.001: LSASS Memory	4.4%	T1105: Ingress Tool Transfer	24.2%	T1573.002: Asymmetric Cryptography	15.9%
		T1003.003: NTDS	3.4%	T1573: Encrypted Channel	15.9%		
		T1003.002: Security Account Manager	0.7%	T1095: Non-Application Layer Protocol	13.0%		
		T1003.006: DCSync	0.2%	T1071: Application Layer Protocol	9.5%	T1071.001: Web Protocols	7.6%
		T1003.008: /etc/passwd and /etc/shadow	0.2%			T1071.004: DNS	1.7%
T1110: Brute Force	6.1%	T1110.003: Password Spraying	2.0%	T1572: Protocol Tunneling	5.4%	T1071.003: Mail Protocols	0.5%
		T1110.001: Password Guessing	1.2%			T1071.002: File Transfer Protocols	0.2%
T1056: Input Capture	4.9%	T1056.001: Keylogging	4.9%			T1090: Proxy	4.9%
T1555: Credentials from Password Stores	1.7%	T1555.003: Credentials from Web Browsers	1.0%	T1090.004: Domain Fronting	0.2%		
T1552: Unsecured Credentials	1.0%	T1552.004: Private Keys	0.5%	T1102: Web Service	1.0%		
		T1552.001: Credentials In Files	0.2%				
T1111: Two-Factor Authentication Interception	0.7%			T1219: Remote Access Software	0.7%		
T1558: Steal or Forge Kerberos Tickets	0.7%	T1558.003: Kerberoasting	0.2%	T1001: Data Obfuscation	0.2%		
T1187: Forced Authentication	0.2%			T1568: Dynamic Resolution	0.2%	T1568.002: Domain Generation Algorithms	0.2%
T1539: Steal Web Session Cookie	0.2%			T1571: Non-Standard Port	0.2%		

Execution

T1059: Command and Scripting Interpreter	51.3%	T1059.001: PowerShell	40.8%
		T1059.003: Windows Command Shell	15.4%
		T1059.005: Visual Basic	5.9%
		T1059.007: JavaScript/JScript	2.7%
		T1059.006: Python	1.0%
T1569: System Services	30.6%	T1569.002: Service Execution	30.6%
T1053: Scheduled Task/Job	15.2%	T1053.005: Scheduled Task	6.6%
T1204: User Execution	11.5%	T1204.001: Malicious Link	7.3%
		T1204.002: Malicious File	4.2%
T1203: Exploitation for Client Execution	4.9%		
T1047: Windows Management Instrumentation	2.7%		
T1106: Native API	0.2%		

Across the Lifecycle

Defense Evasion

T1027: Obfuscated Files or Information	52.6%	T1027.001: Binary Padding	0.2%	T1140: Deobfuscate/Decode Files or Information	2.7%		
		T1027.004: Compile After Delivery	0.2%	T1218: Signed Binary Proxy Execution	2.4%	T1218.010: Regsvr32	1.0%
		T1027.005: Indicator Removal from Tools	1.0%			T1218.002: Control Panel	0.5%
		T1027.002: Software Packing	8.1%			T1218.005: Mshta	0.5%
		T1027.003: Steganography	0.5%			T1218.003: CMSTP	0.2%
T1070: Indicator Removal on Host	24.4%	T1070.004: File Deletion	18.1%			T1218.011: Rundll32	0.2%
		T1070.006: Timestomp	5.9%	T1564: Hide Artifacts	2.2%	T1564.003: Hidden Window	2.0%
		T1070.001: Clear Windows Event Logs	4.2%			T1564.004: NTFS File Attributes	0.2%
		T1070.005: Network Share Connection Removal	1.2%	T1036: Masquerading	1.5%	T1036.003: Rename System Utilities	0.7%
T1553: Subvert Trust Controls	21.3%	T1553.002: Code Signing	21.0%			T1036.001: Invalid Code Signature	0.5%
T1055: Process Injection	18.1%	T1055.003: Thread Execution Hijacking	1.0%			T1036.005: Match Legitimate Name or Location	0.2%
		T1055.012: Process Hollowing	0.5%	T1480: Execution Guardrails	1.5%		
T1112: Modify Registry	15.6%			T1197: BITS Jobs	0.7%		
T1497: Virtualization/Sandbox Evasion	12.7%	T1497.001: System Checks	1.5%	T1548: Abuse Elevation Control Mechanism	0.7%	T1548.002: Bypass User Account Control	0.5%
T1562: Impair Defenses	9.8%	T1562.001: Disable or Modify Tools	5.9%			T1548.001: Setuid and Setgid	0.2%
		T1562.004: Disable or Modify System Firewall	5.1%	T1578: Modify Cloud Compute Infrastructure	0.5%	T1578.002: Create Cloud Instance	0.5%
		T1562.007: Disable or Modify Cloud Firewall	0.2%			T1578.003: Delete Cloud Instance	0.2%
T1078: Valid Accounts	6.8%			T1550: Use Alternate Authentication Material	0.5%	T1550.002: Pass the Hash	0.2%
T1134: Access Token Manipulation	5.9%	T1134.001: Token Impersonation/Theft	0.2%			T1550.003: Pass the Ticket	0.2%
T1202: Indirect Command Execution	3.7%			T1127: Trusted Developer Utilities Proxy Execution	0.2%	T1127.001: MSBuild	0.2%
T1574: Hijack Execution Flow	3.2%	T1574.001: DLL Search Order Hijacking	2.4%	T1211: Exploitation for Defense Evasion	0.2%		
		T1574.002: DLL Side-Loading	2.4%	T1484: Domain Policy Modification	0.2%	T1484.001: Group Policy Modification	0.2%
		T1574.008: Path Interception by Search Order Hijacking	0.2%	T1542: Pre-OS Boot	0.2%	T1542.003: Bootkit	0.2%

RANSOMWARE



Ransomware Evolves Into Multifaceted Extortion



Our understanding of ransomware was appropriate for 2019 but the way ransomware attacks are conducted today has changed, resulting in different business consequences and different protections must be put in place. When business leaders and risk managers hear “ransomware,” they often envision scenarios of malware encrypting files, making them inaccessible to legitimate users, and ultimately resulting in some level of business disruption. They also believe that the best protection against these sorts of attacks is solid offline backups. Now, however, the problem is fundamentally different, yet we still refer to the problem as ransomware. This mischaracterization does not serve organizations well and they are unprepared when an attack’s true nature is revealed in the midst of a real incident. To better confront and mitigate these incidents, Mandiant has adopted the term “multifaceted extortion” to characterize this evolved form of ransomware.

The Facets of Multifaceted Extortion

1. Deployment of ransomware encryptors.

The target organization's files are encrypted and made unavailable. The attacker demands a payment for the decryption tool and key.

2. Theft of sensitive data.

The organizations files are stolen and the attacker demands a payment not to publish the sensitive data. This extortion is much more consequential than the first and it may give the attacker more leverage. With a multifaceted extortion, the attacker turns a service disruption into a data breach. Data breaches may have more serious business consequences than service disruptions. A data breach can result in greater reputational damage, regulatory fines, class action lawsuits and derailed digital transformation initiatives. These consequences were typically not seen with traditional ransomware before 2019. Organizations may not expect such consequences if they continue to think of modern multifaceted extortion attacks simply as ransomware.

3. Publication of stolen data on a "name-and-shame" website.

Many multifaceted extortion threat actors operate such sites on the Tor network. The actors may engage security and technology media organizations to amplify their attacks and attempt to coerce victims into paying.

With multifaceted extortion, in addition to deploying ransomware encryptors and disrupting business operations, the threat actors steal data, publish it and shame victims. Having good backups only addresses part of the problem.

4. Additional coercive tactics.

To compel victims into paying extortion demands, threat actors have applied pressure in various ways:

- Convinced news and media organizations to write stories on victim security incidents
- Called and harassed employees
- Notified business partners of data theft, creating friction in relationships and prompting breach disclosures
- Conducted distributed denial of service attacks to further disrupt operations

Disruption and Brand Damage Due to Multifaceted Extortion

Multifaceted extortion continues to be a leading concern for organizations as threat actors evolve their technology and tradecraft in response to changes in the security landscape. Mandiant has observed these actors trending away from purely opportunistic campaigns, which seek to maximize the volume of attacks, toward campaigns which require greater complexity. Attacker behavior patterns have begun to emerge which demonstrate adoption of tools, tactics and procedures (TTPs) which are more consistent with more advanced threat groups. As defense technology advances so do multifaceted extortion operators. The impetus to maintain operations and increase monetary gains has forced these operators to adopt new techniques and even more dramatic extortion demands. While attacker objectives remain consistent, their methods have evolved, leading to greater monetary gains for threat actors and a growing population of threat groups seeking to replicate these successes.


Beginning in November 2019 and increasing throughout 2020, Mandiant observed threat actors combining ransomware operations with data theft extortion. During these types of operations, malicious actors steal sensitive data before deploying ransomware on the network and then threaten to release the stolen data if the ransom is not paid. Actors use the release of sensitive data during negotiations to drive up the price of extortion and secure more timely payments from a victim organization. The exposure of data which triggers a notification requirement is often a major concern during an organization's incident response process. While some stolen sets of data would not trigger regulatory reaction, these actors would often prey upon the potential brand damage which could occur as a result in a loss of confidence and trust from partners and customers.

Historically, post-compromise ransomware actors have used data theft to obtain information that can expand their presence in a target network. However, they are now seeking additional types of data to support their extortion operations. As ransomware operators evolve into multifaceted extortion operators they seek access to sensitive information which can provide enhanced leverage during negotiations, and the opportunities to detect them increase dramatically. Mandiant has observed these operators steal a diverse set of data from target environments, including termination agreements, contracts, medical records and encryption certificates. Depending on the organization's degree of network segmentation, access to the enclaves which would house these data types would require the use of multiple credentials across disparate systems. Each system introduces further opportunities for the attacker to be detected and evicted from the network prior to any theft of sensitive data or activation of encryption tools.

Multifaceted extortion operators have used a wide range of tactics to increase pressure on data theft victims. Operators of the MAZE ransomware threatened to use stolen data to conduct targeted spam campaigns while the operators of the Ragnar Locker ransomware used Facebook ads to shame victims. The most common tactic was the creation and maintenance of name-and-shame sites where these operators would post data stolen from victims who refused to pay the extortion. Analysis of the content available on various shaming websites from October 1, 2019 to September 30, 2020, highlights several trends among victims; the data is obviously skewed toward those who chose not to acquiesce to extortion demands. There is a distinct upward trend in both the number of victims that have appeared on these sites (Fig. 1) and the number of groups using this methodology to pressure victims.

This method of combining ransomware and data theft has proven successful for several multifaceted extortion operations. Consequently, other threat actors have attempted to follow suit. MAZE operators started the shaming website trend in December 2019 and in February 2020, SODINOKIBI and DOPPELPAYEMER followed suit with their own versions of shaming websites. By March 2020, NEMTY, NEFILIM, CLOP, Sekhmet and m1x operators also started using shaming websites. Since then, there has been an average of at least one new shaming website each month through September 2020. Based on these shaming websites, media outlets and Mandiant incident response engagements, Mandiant Threat Intelligence has identified more than 800 alleged multifaceted extortion victims who likely had data stolen. This number has continued to grow steadily as more operators have started websites to support extortion demands and data publishing.

Figure 1.
 Percentage of victims from December 2019 to September 2020.



Most of the victim organizations listed on shaming websites have been based in the U.S. and spanned nearly every industry vertical, demonstrating the widespread targeting basis of these operations. Based on available data, organizations within the manufacturing sector have been impacted more than other industries. There are likely several contributing factors including the perception these organizations may be more likely to pay to prevent monetary losses due to production downtime and their overall cyber security posture relative to other sectors. Several ransomware families have also been deployed alongside scripts with process kill lists for industrial processes. Other industries experiencing frequent multifaceted extortion attacks included professional services, retail, technology, financial services, healthcare and construction and materials.

Over the last few years, traditional malware based ransomware has evolved into multifaceted extortion through repeated and deliberate operations costing organizations and governments millions of dollars. In response, many organizations took steps to limit the potential impact of broad-scale encryption by ensuring their disaster recovery plans included a similar scenario. However attacker-encrypted files is just one of many impacts victims face in a multifaceted extortion incident. New technologies and better visibility into the methodologies used by attack groups have increased opportunities for early detection. As defensive security evolves and organizations continue to invest in a broader security position, attackers are guaranteed to modify their operations to keep pace with those changes. Until multifaceted extortion operations can be reliably interrupted through legal and policy changes, the burden of their effects will continue to fall on organizations and the security teams which support them.

Steps Toward Proactive Hardening Against Ransomware in Multiple Environments



Given the surge of ransomware events observed throughout 2020 and into 2021, organizations must proactively ensure environments are hardened to mitigate the potential impact of ransomware deployment. After reviewing ransomware engagements supported throughout 2020, Mandiant experts uncovered several actions organizations should prioritize to mitigate the risk of ransomware incidents. These actions would address several common issues observed, including:

- Large numbers of highly privileged accounts in Active Directory
- Highly privileged non-computer accounts configured with service principal names (SPNs)
- Security controls not configured to minimize the exposure and usage of privileged accounts across endpoints
- Attackers modifying Group Policy Objects (GPOs) for ransomware deployment

Large Numbers of Highly Privileged Accounts in Active Directory

When conducting initial reconnaissance within target environments, Mandiant experts observed ransomware attackers “living off the land”: using native built-in tools (such as cmd.exe, PowerShell, WMI) to query information from Active Directory. Attackers also often used open-source tools, such as AdFind and BloodHound, to identify privileged accounts and accounts that presented a path to Domain Admin.

Privileged accounts within Active Directory represent more than just accounts assigned membership in the built-in domain-based privileged groups (Domain Admins, Enterprise Admins, Schema Admins, Administrators and Server Operators). Apart from the domain-based privileged groups, many organizations have delegated permissions to additional groups and accounts throughout Active Directory, which significantly increases the number of resources deemed

to be highly privileged. If an attacker is able to capture valid credentials or even impersonate access from an account assigned privileged access, this can escalate the attacker's ability to move laterally, access data and deploy ransomware to many endpoints.

Organizations should proactively review Active Directory for accounts resident in the default domain-based privileged groups as well as accounts which:

- Have been delegated permissions at the root of the domain—including permissions for DS-Replication-Get-Changes and DS-Replication-Get-Changes-All (which can be used to initiate a DCSync attack)
- Have been explicitly assigned elevated permissions on domain controllers
- Have been delegated explicit permissions for organizational units (OUs) that contain many computer and user objects
- Have local administrative permissions on many endpoints
- Are configured for unconstrained or constrained Kerberos delegation
- Are not configured to protect against delegation (not members of the Protected Users Security Group – or do not have the “Sensitive and Cannot Be Delegated” attribute configured)
- Are protected by AdminSDHolder
- Have the ability to edit, link or unlink Group Policy Objects (GPOs)
- Have the ability to change passwords for many accounts (User-Force-Change-Password permissions)
- Have user rights assignment permissions configured in Group Policy that permit remote logon capabilities to a large scope of endpoints

Highly Privileged Non-Computer Accounts Configured with Service Principal Names (SPNs)

Service principal names (SPNs) are legitimately used by Active Directory to identify unique service instances with service logon accounts for Kerberos authentication. Non-computer accounts configured with an SPN represent accounts that ransomware attackers will initially target for Kerberoasting.¹ If an account's password can be brute-forced using the Kerberoasting technique, an attacker can potentially use the account for lateral movement, privilege escalation and ransomware staging and deployment. Unfortunately, many organizations have highly privileged accounts configured with an SPN, which makes this attack-path highly successful.

Organizations can proactively review the scope of accounts configured with an SPN using PowerShell. (Fig 2.)

```
get-aduser -filter {(ServicePrincipalName -like "*")}
```

Figure 2.
PowerShell
command to
identify non-
computer accounts
assigned an SPN.

[2. The MITRE Corporation \(2015-2021\). Steal or Forge Kerberos Tickets: Kerberoasting](#)

Non-computer accounts configured with an SPN are expected to exist within Active Directory since some service accounts will likely have this configuration. Organizations should prioritize permissions assigned to these accounts and ensure enforcement of security controls to minimize privileges and lateral movement capabilities.

Security Controls Not Configured to Minimize the Exposure and Usage of Privileged Accounts Across Endpoints

When endpoint security controls were not configured to minimize the exposure of privileged accounts, ransomware attackers were able to capture valid credentials from memory to expand the scope of their access. The recent Mandiant white paper, Ransomware Protection and Containment Strategies,³ recommends several proactive hardening measures to protect privileged accounts:

- Use the Protected Users security group to house non-service privileged accounts.
- Disable methods that store clear-text credentials in memory on endpoints (such as WDigest and Windows Credential Manager). This also includes using a Group Policy configuration to automatically reapply these settings if they were to be modified on the local endpoint by an attacker.
- Enforce Credential Guard and Remote Credential Guard on Windows 10 and Windows Server 2016+ endpoints. For older endpoints, Restricted Admin Mode should be used when remote desktop connections are initiated using privileged accounts.
- Use Microsoft LAPS⁴ or other third-party tools to randomize the password for the built-in local administrator account on endpoints.
- Implement a tiered model to guide enforcement of guardrails that define where and how privileged accounts can be used. Guardrail enforcement can be defined within Group Policy Objects (GPOs) or when using authentication silos.
- Ensure privileged actions are only initiated from dedicated privileged access workstations (PAWs) or jump boxes.
- Protect privileged accounts against delegation (such as “Sensitive and Cannot Be Delegated” enforcement).
- Enforce the Windows Firewall to restrict protocols which could be used for lateral movement, remote access and ransomware deployment across endpoints.
- Restrict the scope of accounts and groups where permissions are configured to allow lateral movement across endpoints. Specifically, privileged accounts should be denied the ability to logon to Tier 1 and Tier 2 endpoints (with local or Group Policy settings):
 - Deny access to this computer from the network (SeDenyNetworkLogonRight)
 - Deny log on through Remote Desktop Services (SeDenyRemoteInteractiveLogonRight)
 - Deny log on locally (SeDenyInteractiveLogonRight)
 - Deny log on as a service (SeDenyServiceLogonRight)

³. FireEye (2020). Ransomware Protection and Containment Strategies.

⁴. Microsoft (2021). Microsoft Local Administrator Password Solution (LAPS).

- Restrict the scope of accounts and groups where permissions could be used for privilege escalation and data access, including assigned user rights that provide capabilities to:
 - Debug programs (SeDebugPrivilege)
 - Back up files and directories (SeBackupPrivilege)
 - Restore files and directories (SeRestorePrivilege)
 - Take ownership of files or other objects (SeTakeOwnershipPrivilege)

Attackers Modifying Group Policy Objects (GPOs) for Ransomware Deployment

Group policy was a common method used by attackers to deploy ransomware to many endpoints. After compromising an account with GPO edit permissions, attackers commonly targeted GPOs which were linked at the root of the domain (such as Default Domain Policy), and then added scheduled tasks, logon scripts or software installation packages to mass deploy and execute encryptors.

Organizations should proactively review Active Directory to enumerate accounts with the ability to edit existing GPOs. Additionally, organizations should identify which accounts can link and unlink GPOs within the domain.

A recently-published Mandiant white paper⁵ details specific command references for reviewing GPO permissions as well as strategies for monitoring event logs to detect GPO creations and modifications.

[5. FireEye \(2020\). Ransomware Protection and Containment Strategies.](#)

Recovery and Reconstitution Challenges in Post-Ransomware Scenarios



Following a ransomware attack, sustained business operations are predicated on an organization's ability to restore and reconstitute systems, data and access in a secure manner.

Ideally, recovery and reconstitution should run parallel to the investigation into how an attacker was able to gain access, move laterally, steal data (if applicable) and deploy ransomware. If recovery and reconstitution steps are not performed securely, an attacker will likely maintain their access, resulting in continued risk, extended downtime and future attacks.

Investigative Steps to Empower recovery and Reconstitution

Mandiant experts noted that organizations that led successful recovery and reconstitution efforts were empowered by facts and evidence identified from investigating the ransomware incident prior to taking action. To ensure a secure recovery and reconstitution plan when responding to ransomware incidents, Mandiant experts prioritize answering the following questions:

- What persistence mechanisms (backdoors) are being used by an attacker to maintain access within the environment?
- What command-and-control (CnC) channels are being used by an attacker for access that need to be blocked either at ingress and egress points, or using host-based endpoint firewalls?
- How was the ransomware deployed (manual propagation using a tool such as PsExec, group policy modification, scheduled tasks or logon script)?
- Can the ransomware propagation vector be stopped and contained?
- Which compromised accounts can be used for lateral movement and ransomware deployment?

- Which endpoints still exhibit a running encryptor?
- What was the primary vector used for initial access?
- Which accounts have privileged access within the environment?

As specific facts are uncovered and correlated, proper isolation and hardening steps must be aligned to recovery and reconstitution workstreams. Until specific investigative information about the ransomware event is understood (primary persistence methods and backdoors, CnC channels, ransomware propagation mechanisms and scope of compromised accounts), an organization may be forced to temporarily disconnect access from the Internet to begin triaging, reviewing and securely planning for hardening and recovery.

Use of Secure Enclaves

Organizations that were able to successfully recover and reestablish operations used processes that included the use of secure enclaves (such as VLANs) to safely restore and reconstitute systems. Initially, systems housed within the enclaves were restricted from communicating with most impacted systems and were designated as a clean (green) source of trust and security within the environment.

Systems within the enclaves were only permitted to communicate with trusted endpoints and security (EDR) tools and with investigative tools used to ensure the endpoints were not exhibiting signs of active encryptors or other malicious activity.

Restored or newly-built domain controllers were usually the first endpoints established within the secure enclaves and connectivity to additional endpoints was expanded once the initial ransomware outbreak had been successfully contained.

Common Contributors to a Degraded Recovery and Reconstitution Workstream

Active Directory Lock-Out

Throughout 2020, Mandiant observed attackers not only deploying ransomware within victim environments but also locking out administrators from Active Directory. Even if domain controllers were operational, after an attacker obtained domain administrative privileges, they changed the passwords for valid administrator accounts. As a result, organizations were unable to regain access and subsequent control within Active Directory to begin recovery and reconstitution actions. To regain access to Active Directory, organizations often had to boot a domain controller from physical media and replace the Utility Manager binary (utilman.exe) with cmd.exe. Consequently, when the user clicked on Utility Manager (on the login screen) they actually launched a command prompt from which a new password could be set for a known account.

Domain Controller Restoration Challenges

When ransomware is deployed within an environment, domain authentication and the ability to connect to domain-based resources are generally impacted. To begin recovery and reconstitution efforts, most organizations will prioritize restoring domain controllers to an operational state. Unfortunately, when domain controller backups are unavailable or have been corrupted, organizations may be faced with rebuilding their domain infrastructure.

The corruption of SYSVOL is often one of the main factors resulting in domain authentication and domain services being unavailable. SYSVOL refers to a set of folders and files that are present and replicated amongst each domain controller. When the contents of SYSVOL (which include group policy templates and settings as well as scripts) are either encrypted or corrupted and replication occurs, the contents cascade to all domain controllers and impact domain operations. Without timely backups, rebuilding and reconstituting SYSVOL can be a complex process. To avoid having to rebuild SYSVOL entirely, organizations should ensure backup processes include system state backups of domain controllers. If system state backups are not possible, organizations should at least back up the SYSVOL directory and data tree (%SYSTEMROOT%\SYSVOL) as well as the Active Directory database (%SYSTEMROOT%\ntds\ntds.dit). Ideally, domain controller backups should be stored and secured either offline or on storage clusters that are logically segmented and restricted from being directly accessed using domain accounts.

When restoring Active Directory from domain controller backups was the only viable option to reconstitute domain services, organizations first needed to ensure they had a working and tested backup plan and strategy to guarantee the availability and integrity of the schema and domain services which needed to be reconstituted. The following domain controller recovery and reconstitution best practices should be proactively reviewed by the organization:

- Offline backups—offline domain controller backups should be secured and stored separately from online backups.
- Encryption—backup data should be encrypted both during transit (over the wire) and when at rest or mirrored for offsite storage.
- DSRM Password validation—the Directory Services Restore Mode (DSRM) password should be set to a known value for each domain controller. This password is required when performing an authoritative or non-authoritative domain controller restoration.
- Alert configuration for backup operations—backup products and technologies should be configured to detect and provide alerting for operations critical to the availability and integrity of backup data (such as deletion of backup data, purging of backup metadata, restoration events and media errors).
- Role-based access control—Access to backup media and the applications that govern and manage data backups should use role-based access controls to restrict accounts with access to stored data and configuration parameters.
- Testing and verification—both authoritative and non-authoritative domain controller restoration processes should be documented and tested.

Insufficient Backup and Restoration Processes

Organizations that lacked effective backup and restoration processes struggled to reestablish business operations in a timely manner. This resulted in extended downtime, financial impacts, regulatory challenges and an overall impact to brand reputation. Many organizations were either unable to restore and reconstitute data to meet business operations requirements or had to fully rebuild core systems and applications, which further impacted business continuity.

Organizations that had effective backup and restoration processes were able to mobilize quickly and invoke secure restoration activities which ran in parallel to the overall investigation and environment hardening workstreams. Examples of effective backup and restoration processes include:

- Clear delineation of responsibility for managing and verifying data and application backups.
- Alignment of backup and restoration processes with business continuity and disaster recovery plans.
- Online and offline data backup retention policies, including initiation, frequency, verification and testing (for both on-premises and cloud-based data).
- Backup infrastructure that was segmented within the environment and only accessible using dedicated accounts for interfacing with and managing backup platforms and storage.
- A well-defined understanding of crown jewels data and supporting applications which align to backup, failover and restoration tasks that prioritize mission-critical business operations.
- Role-based access control that restricts the scope of accounts with access to backup media and the applications that govern and manage data backups, as well as stored data and configuration parameters.
- Training and validation exercises for failover and restoration tasks that focuses not only on data, but also crown jewel applications and services.
- Established service level agreements (SLAs) with vendors to prioritize application and infrastructure-focused support.

Summary

Ransomware resiliency defenses encompass a combination of technical and process-oriented controls. When organizations are unprepared, the impact can be time-consuming and costly. Throughout 2020, Mandiant worked with many organizations to proactively review, test, and enhance defenses to combat ransomware. While challenges remain, applying best practices can go a long way towards realizing robust recovery and reconstitution efforts.

NEWLY NAMED THREAT GROUPS



FIN11



FireEye tracks thousands of threat actors and pays attention to groups that carry out repeated intrusions across organizations. Such groups often pursue their objectives over longer periods, typically months or years. They rapidly adapt to a victim organization's attempts to remove them from the network and frequently target the same victim again if access is lost.

In 2020, FireEye promoted one attack group from a previously tracked TEMP group to a FIN group.

How a Threat Activity Cluster Becomes an “APT” or “FIN” Group

Mandiant analysts review threat activity data from a variety of sources—such as FireEye security product telemetry and Mandiant incident response engagements and research—to identify noteworthy clusters. Our team of technical and threat researchers, analysts and reverse engineers begin their work from known indicators and attempt to find related indicators, activity or other data. When only a small cluster of activity is found, we reference that activity in finished intelligence (FINTEL) without a formal name to various channels, which may include Mandiant Advantage and external blogs. Example: “Suspected Iran-based nation-state threat actors sent spear phishing emails....”

Uncategorized (UNC) groups are raw attribution analysis that were previously kept primarily in house. An UNC is a cluster of cyber intrusion activity—which includes observable artifacts such as adversary infrastructure, tools, and tradecraft. UNC’s are created based on a defining, anchoring characteristic often discovered during a single incident. As we discover new artifacts associated with other incidents and proactive collection efforts, the UNC provides a framework to join discrete pieces of evidence together.

Some clusters develop further with sufficient or consistent research that identifies their tactics, techniques, and procedures (TTPs). In these cases, the cluster is given a temporary “TEMP.<xxx>” group name. For example, APT37 was previously reported as the “TEMP.Reaper” group.

As our knowledge of a TEMP group becomes sufficiently mature, we apply a consistent, rigorous methodology to assign the actor a formal APT or FIN number. Advanced persistent threat (APT) groups are generally focused on espionage activities. Financially-motivated (FIN) groups are highly organized criminal groups engaging in crime for financial gain, such as payment card data theft, business email fraud and extortion activities.

FIN11: A Widespread Ransomware and Extortion Operation

FIN11 is a financially motivated group that has been active since at least 2016. The group uses malware such as FlawedAmmy and FRIENDSPEAK in widespread phishing campaigns that have impacted organizations across a broad range of sectors and geographic regions. To monetize its operations, FIN11 deploys CLOP ransomware and extorts victims for the non-release of stolen data shared via its public leak site. In at least one case, Mandiant Threat Intelligence observed FIN11 use point-of-sale (POS) malware. Although the group does not exhibit a high level of technical sophistication, FIN11 appears to consistently evolve its malware delivery tactics and techniques. The group has also relied on several support services to accomplish its mission (Fig. 3).

FIN11 includes a subset of activity that some security researchers call TA505. This term has been widely used in the security community to discuss large-scale spam campaigns that date to 2014 and initially distributed malware families such as the Dridex Trojan and Locky ransomware. We have not attributed TA505’s early operations to FIN11 and caution against using the names interchangeably.

Figure 3.
Resources used by
FIN11.



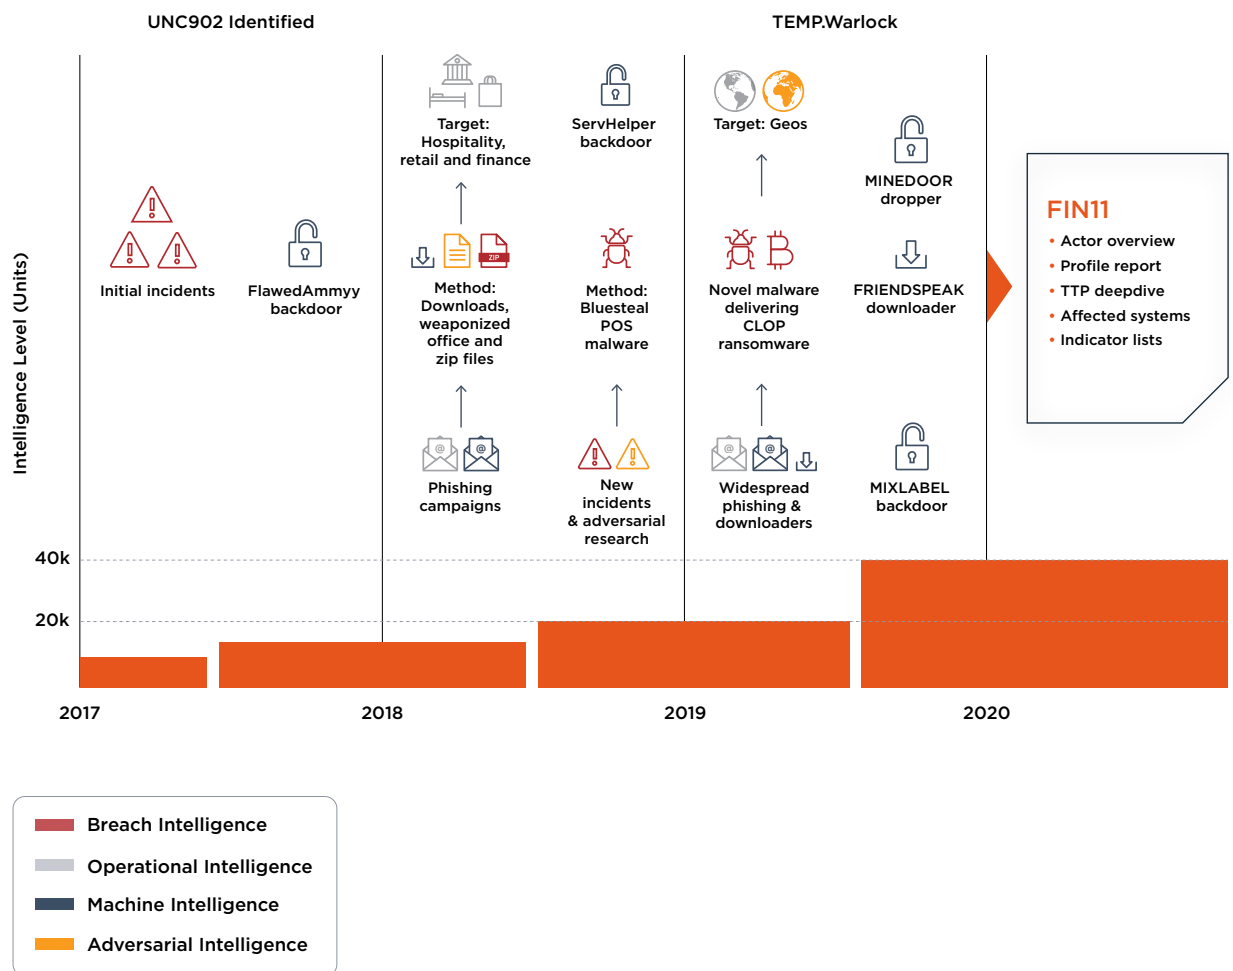
We have no evidence that indicates when FIN11 was formed. We began tracking FIN11 (as UNC902) in 2017 (Fig. 4). It is plausible that the historical TA505 activity, which dates to at least 2014, was FIN11 (or a FIN11 member); however, we can't independently verify those claims, and the purported use of services (such as Dridex and Emotet) complicates attribution analysis. Given the lack of first-hand knowledge and the mutable nature of criminal groups, any beliefs about when the group was formed would be speculative.

FIN11 was chosen for graduation due to the group's high activity level, its successful compromise of multiple organizations, our insight into the group's TTPs across the attack life cycle and our clients' interest in FIN11 and ransomware. There were several developments in Fall 2019 that supported this decision. From 2017 through mid-2019, we had limited insight into the later stages of FIN11 intrusions. In early Fall 2019, Mandiant responded to multiple FIN11 CLOP deployments that helped fill in those knowledge gaps. Further, in September 2019 the group replaced the FlawedAmmy backdoor with FRIENDSPEAK and MIXLABEL, which created additional client interest.

FIN11 campaigns have impacted a wide variety of sectors and geographical regions. The group's spam campaigns from 2017 to 2018 primarily targeted organizations in the financial, retail and restaurant sectors. In 2019 and 2020, FIN11 expanded its targeting to a larger, more indiscriminate and diverse set of industries and countries, often using generic financial lures. However, a portion of FIN11's 2019 and 2020 campaigns targeted organizations in specific industries or regions; to appear more legitimate, the group often used the target's native language coupled with manipulated email sender information, such as spoofed email display names and email sender addresses. The shift in targeting observed during the past two years may be the result of FIN11's transition from point-of-sale (POS) malware to ransomware as their main monetization method.

A hallmark of FIN11 activity since at least January 2019 has been its rapid evolution of phishing campaign TTPs. Throughout its 2019 and 2020 phishing campaigns, the group has made small changes to its initial delivery mechanisms, likely in attempts to circumvent victims' detection regimes. For example, in September 2019, the actors attached macro-laden Office files directly to phishing emails, but over the next few months the infection chain became incrementally more convoluted. By March 2020, most FIN11 phishing campaigns used HTML attachments to load a redirect from a compromised URL to a download domain that subsequently delivered a macro-laden Office file. We assess that these relatively minor and less novel modifications are not reflective of the group's sophistication.

Figure 4.
FIN11 timeline.



While FIN11 doesn't exhibit a particularly high level of technical sophistication, the group's methods are relatively efficient and effective (Fig. 5). The group rarely uses exploits as their initial infection vector and instead relies on phishing emails and user execution (opening macro-laden Microsoft Office files) to obtain an initial foothold. The actors then use common exploitation frameworks and publicly available utilities for reconnaissance, privilege escalation and lateral movement.

We suspect that several malware families are exclusive to FIN11, but it isn't clear if FIN11 develops its tools or outsources development. It is plausible that FIN11 uses gaps in phishing activities for tool development, as we have often observed new versions of the MIXLABEL backdoor being deployed shortly after the group resumes operations. Some researchers have stated that CLOP is a Cryptomix variant, which could suggest that the tool was acquired. However, while Mandiant did identify some notable overlaps between the two ransomware families, we have insufficient evidence to support the theory that CLOP is a variant of Cryptomix.

We have no reason to believe that other actors will begin using the malware families that appear exclusive to the group. FIN11 seems to be the exclusive user of FlawedAmmyy; we have not observed FlawedAmmyy activity since FIN11 ceased using the backdoor in 2019. Similarly, the group has been using BARBWIRE and MIXLABEL since 2018, and we have no evidence to suggest that other groups use these tools.

FIN11 has used several similar droppers—MINEDOOR, SPOONBEARD and FORKBEARD—to deliver malware such as FRIENDSPEAK and BARBWIRE. We suspect that these droppers are not exclusive to the group, as they have also been used to pack malware typically associated with other threat groups.

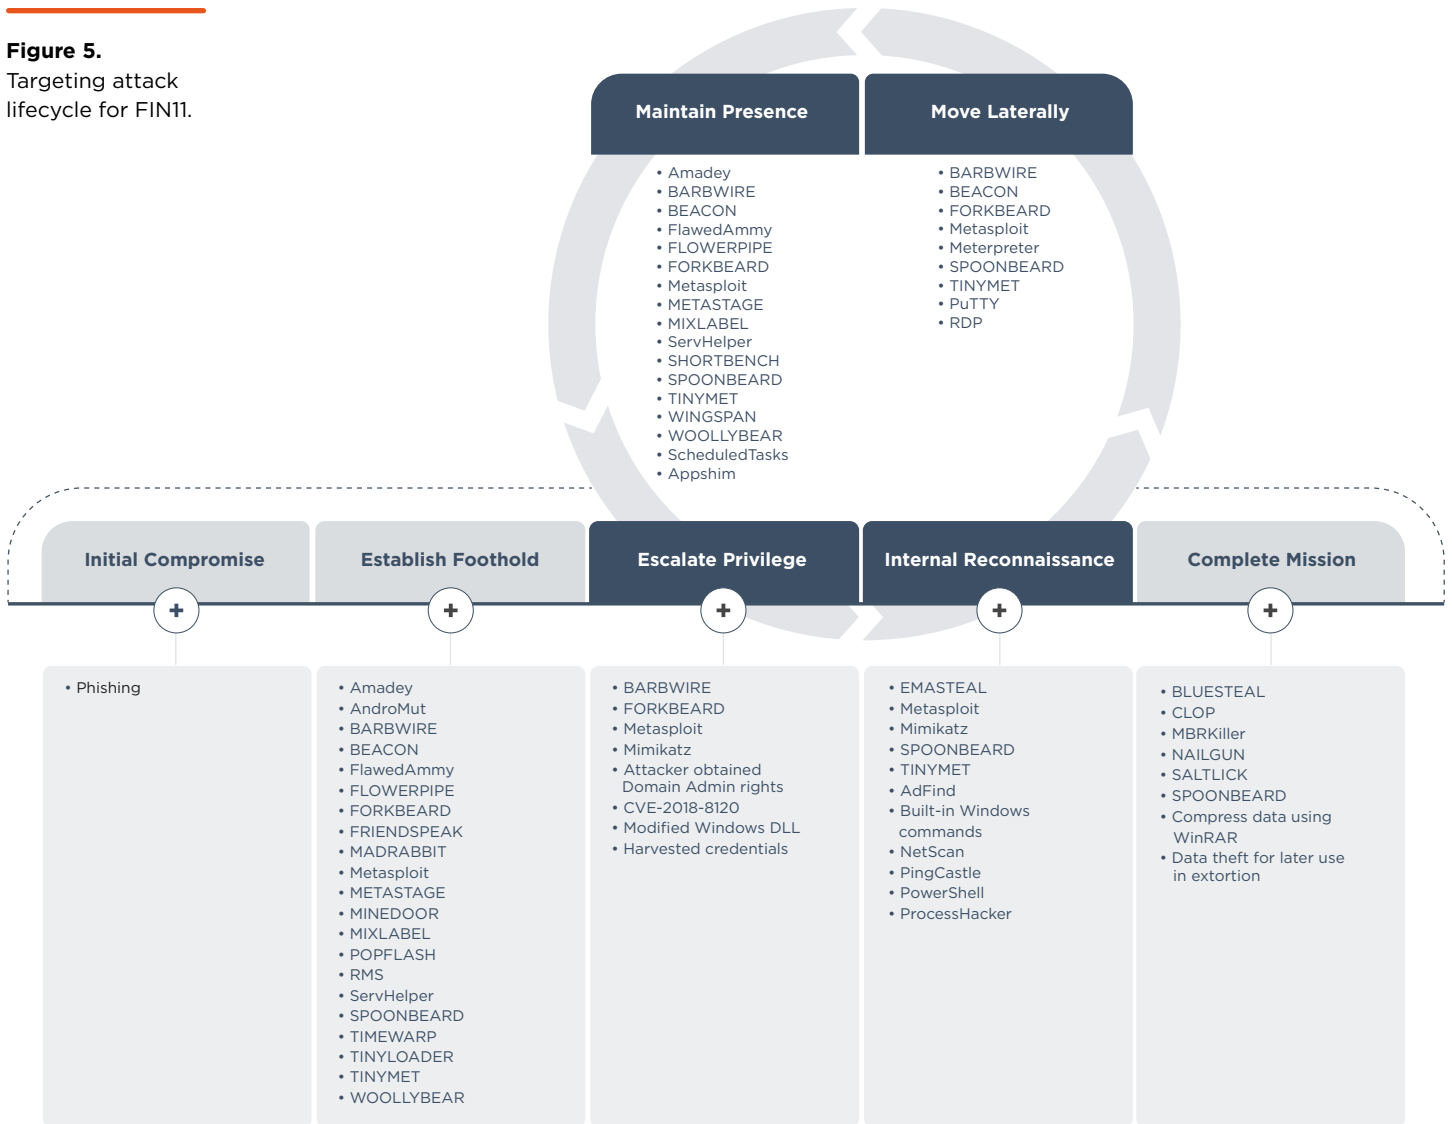
Despite the group's widespread high-volume phishing campaigns, we have only observed evidence of FIN11 successfully monetizing its operations in a handful of cases. In late 2018, Mandiant observed FIN11 attempt to monetize its operations using the point-of-sale (POS) memory scraping tool BLUESTEAL at a restaurant organization. Since then, FIN11 has deployed CLOP ransomware at a variety of organizations and incorporated data theft to increase the pressure on victims to pay extortion fees.

In 2020, FIN11 has conducted hybrid extortion attacks, combining ransomware with data theft to pressure its victims to submit to extortion demands. In cases where we observed data theft, the actors accessed several dozen systems, staged data in RAR archives, uploaded the files to MegaSync servers, deployed CLOP ransomware and then sent an email threatening to publish the data. The exfiltrated data was later posted to a dark website named CLOP^_- LEAKS. Given that we have only observed CLOP distributed by FIN11, we judge that the group also maintains this site.


We assess with moderate confidence that FIN11 is likely operating out of the Commonwealth of Independent States (CIS) based on Russian-language file metadata, avoidance of CLOP deployments in CIS countries and the observance of the Russian New Year and Orthodox Christmas holiday period.

FIN11's frequent, high-volume phishing campaigns are likely an attempt to cast a wide net rather than a reflection of the group's capabilities to monetize an expansive number of victims simultaneously. Even if the campaigns have a relatively low success rate, it is unlikely that FIN11 has the resources to monetize each intrusion prior to being detected. FIN11 may selectively choose victims to exploit further based on criteria such as their geolocation, sector or perceived security posture. Mandiant expects FIN11 spam campaigns to continue in the immediate future and, barring law enforcement action, with continued diversification in delivery tactics.


Figure 5.
Targeting attack
lifecycle for FIN11.



PANDEMIC-RELATED THREATS



Threats Against Organizations Working with COVID-19 Information and Research



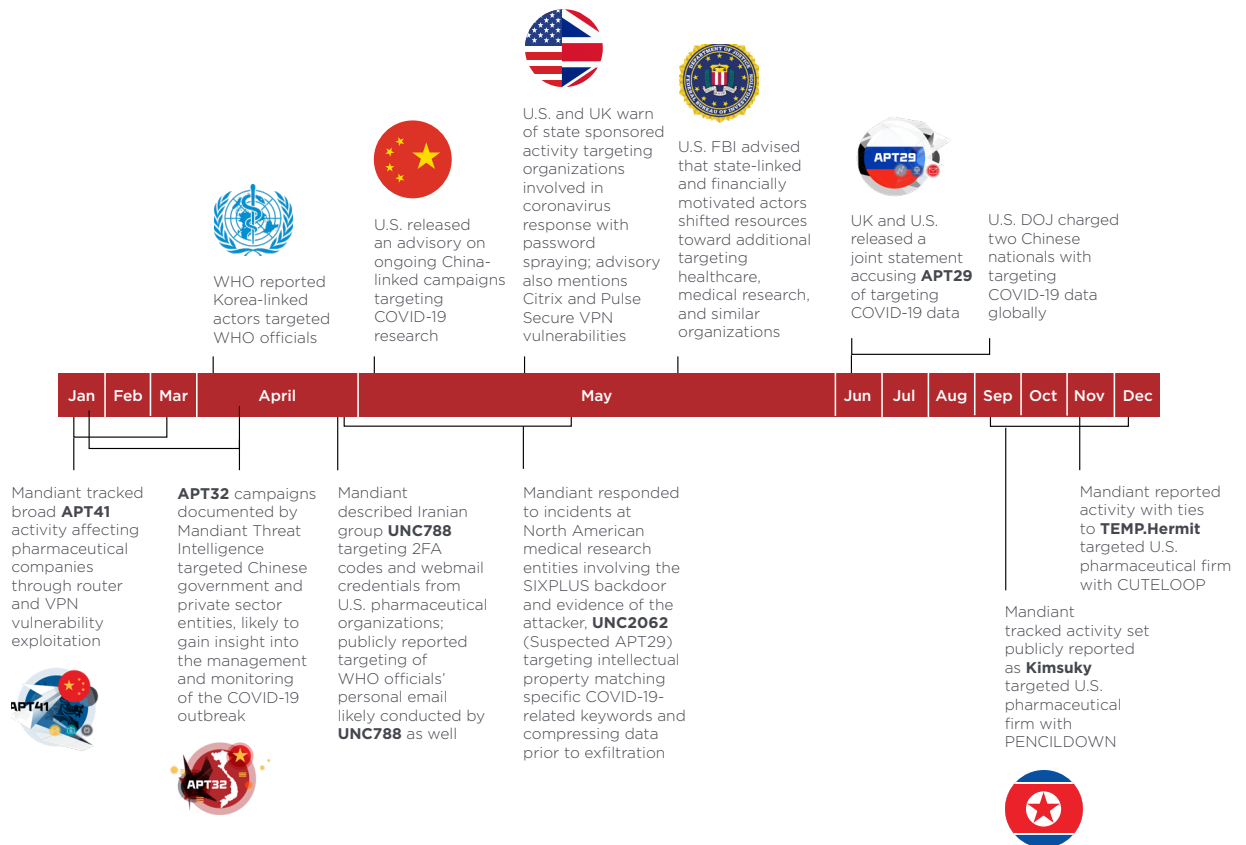
The coronavirus (COVID-19) pandemic caused systemic disruptions to nearly every institution and community, which introduced significant pressure on global leaders to find solutions to prevent or minimize the spread of the illness and economic fallout. In this climate, a successful coronavirus vaccine, as well as proven strategies for containing spread and treating symptoms, are highly prized. Mandiant experts have observed that governments deployed cyber espionage capabilities in pursuit of any information that may support their own domestic recovery and vaccine development efforts and provide a strategic advantage internationally.

Throughout 2020, Mandiant Threat Intelligence tracked many cyber espionage campaigns likely seeking COVID-19 vaccine or treatment data, including Vietnamese, Chinese, North Korean, Iranian and Russian threat groups. Many of these threat groups have been active for years and have been successful in the past.

State-sponsored activity potentially targeting COVID-19 research and response began in January when Mandiant tracked broad China-nexus APT41 activity affecting pharmaceutical companies through router and VPN vulnerability information (Fig. 6). Around the same time frame, Vietnam's APT32 conducted a campaign targeting Chinese public and private entities, likely in efforts to gain insight into the management and monitoring of the COVID-19 outbreak.

In March, the World Health Organization reported on Korea-linked actors targeting its officials, and in May, the U.S. released an advisory on ongoing China-linked campaigns targeting COVID-19 research. In April, Iran's UNC788 targeted US pharmaceutical organizations; in addition, we observed the actor tracked as UNC2062 target medical research facilities in the spring of 2020. In May through July, the U.S. government released multiple statements and advisories, some of which were released jointly with foreign partners, related to state-sponsored campaigns targeting COVID-19 data. In the fall, public reporting indicated that Russia's APT28 targeted coronavirus treatment and vaccine research. During this time and into 2021, we also observed North Korean-related activity move to target US pharmaceuticals.

Figure 6. State-sponsored activity potentially targeting COVID-19 research.



APT32

APT32 is a Vietnam-nexus cyber espionage actor that conducts foreign and domestic surveillance using commercial tools against internal targets. FireEye has tracked APT32 since 2012, and observed targeting of foreign governments, journalists, dissidents, and foreign corporations that may have a vested interest in Vietnam's manufacturing, consumer products, and hospitality sectors.

From at least January to April 2020, APT32 carried out intrusion campaigns that Mandiant believes were designed to collect intelligence on COVID-19. The campaigns targeted select victims in Beijing and Wuhan, where COVID-19 was first identified. The spear phishing emails sent to these targets contained malicious attachments and embedded tracking links; a subset of the emails included COVID-19 themes. Targeted organizations included China's Ministry of Emergency Management and the Wuhan government.

APT32 likely used coronavirus-themed malicious attachments against additional Chinese-speaking targets. While we have not uncovered the full attack chain, we identified a METALJACK loader displaying a COVID-19 decoy document with a Chinese language title while launching its payload.

APT41

APT41 is a Chinese state-sponsored espionage activity in addition to financially motivated activity potentially outside of state control. Its activity traces back to 2012 when individual members of APT41 conducted primarily financially motivated operations focused on the video game industry before expanding into likely state-sponsored activity.

Between January and March 2020, Mandiant observed APT41 attempt to exploit vulnerabilities in remote access and network appliances at more than 75 customers. APT41 targeted Citrix NetScaler/ADC, which includes VPN functionality (CVE-2019-19781), Cisco routers and Zoho ManageEngine Desktop Central (CVE-2020-10189) in attempts to gain access to organizations in the healthcare, government, education, aerospace and defense, transportation, public and non-profit sectors.

UNC788

UNC788 is a cluster of threat activity suspected of conducting cyber espionage operations on behalf of the Iranian government. Also reported on by ClearSky and CERTFA as "Charming Kitten" and by Microsoft as "Phosphorus," UNC788 is characterized by credential theft operations against corporate and personal email accounts.

In April 2020, UNC788 carried out a credential-harvesting campaign against multiple targets, including the U.S. pharmaceutical industry. The campaign was aimed at gaining personal webmail credentials and likely targeted the victim with a spear phishing email which included a link to a spoofed webmail login page. The group is also believed to have targeted global health employees using similar tactics.

UNC2062

In July 2020, a joint report from the U.K.'s National Cyber Security Centre (NCSC), Canada's Communications Security Establishment (CSE), and the U.S. Department of Homeland Security (DHS) detailed the use of three malware families used to compromise medical research companies in the U.K., U.S., and Canada. The report attributed the observed activity to the Russian sponsored threat group APT29. While plausible, we have not yet been able to corroborate this attribution. APT29 is believed to be affiliated with the Russian Federal Security Service (FSB) or the Russian Foreign Intelligence Service (SVR).

Mandiant tracks related activity as UNC2062 and observed the group using the SIXPLUS backdoor (publicly referred to as WELLMESS) to target medical research entities in the spring of 2020. We have also observed this threat group using publicly available tools and targeting data related to COVID-19.

Threats Originating From North Korean Actors

Since October 2020, Mandiant Threat Intelligence has tracked multiple distinct North Korean activity sets expand outside of established targeting patterns to pharmaceuticals and medical research.

In November 2020, Mandiant reported to intelligence subscribers that a cyber espionage campaign was distributing the CUTELOOP downloader through employment-themed lure material. Since at least April 2020, this group had marginally updated its toolset and expanded its targeting from aerospace and defense to a U.S. pharmaceutical company. We believe this activity set to be North Korean in origin and noted potential ties to another North Korean threat actor we call TEMP.Hermit. A separate North Korean activity cluster we typically see seeking intelligence on nuclear or international relations issues was also observed targeting the same U.S. pharmaceutical company with the PENCILDOWN downloader. We noted indications of this activity set also targeting organizations involved in pharmaceutical research, including universities in Germany and South Korea, a South Korean pharmaceutical company and a U.S. vaccine development company. We also identified associated domains spoofing pharmaceutical, biotechnology, research and other health institutions in September and October 2020, including the World Health Organization (WHO).

The actors using PENCILDOWN almost certainly continued to register domains spoofing pharmaceutical and healthcare companies in October 2020. They also deployed the VENOMBITE loader against COVID-19 researchers in South Korea from December 2020 to January 2021. These actors typically carry out strategic intelligence collection surrounding ongoing Korean peninsula geopolitical issues and most closely align with publicly reported "Kimsuky" operations. Mandiant has detected multiple incidents of activity attributed to this activity cluster, and this group is partially responsible for an increase in the volume of observed North Korean espionage targeting multiple industry verticals in 2020. This actor set possibly shares resources with both TEMP.Hermit and APT38; however, the lines of demarcation between North Korea-linked groups remains disputed. Infrastructure related to VENOMBITE also appears to indicate targeting of research universities and pharmaceuticals as well. The pace of the group's focus and targeted efforts on medical and health entities underscore North Korea's immediate needs related to COVID-19 infections and treatments.

Outlook and Implications

Mandiant assesses with moderate confidence that COVID-19-related targeting of the healthcare, pharmaceutical, medical research and closely related industries will continue to be prominent for the foreseeable future while the pandemic continues. Although targeting volume may not increase, the types of information targeted and the depth of intrusions may be more impactful than many incidents before the pandemic. Most of this COVID-19-related targeting will very likely continue to be from espionage actors, with a much smaller portion likely from financially motivated actors—particularly ransomware operators.

Additionally, contact tracing systems and applications deployed by governments, and often developed and/or operated by third parties, will likely provide additional targets of interest to both espionage and financially motivated actors, given the value of large-scale databases for intelligence gathering, sale in underground markets and development of phishing campaigns.

UNC2452



Mapping UNC2452 Activities to the Targeted Attack Lifecycle Framework

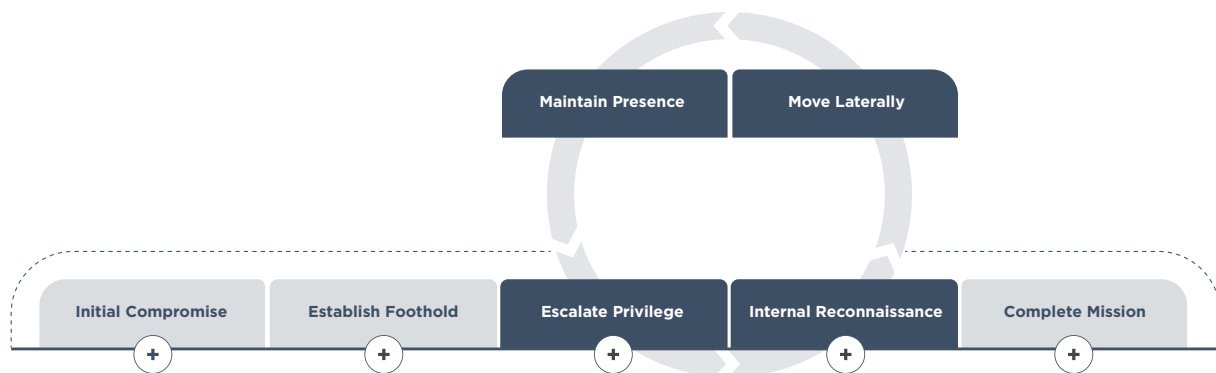


On December 13, 2020, FireEye published a report which detailed a supply chain attack called SUNBURST—an implant in the SolarWinds Orion platform being used to compromise target environments. Mandiant tracks this activity as UNC2452 and suspects the group is state-sponsored. While supply chain attacks are not unheard of in the information security space, this attack used significantly complex, methodical and deliberate tradecraft.


UNC2452 has demonstrated a facility with network intrusions which stand well outside the standard set of activities attackers commonly employ during a breach. Mandiant has observed UNC2452 take advantage of areas in an environment that may be monitored less intensely than others, and remain within those areas as long as possible to reduce opportunities for detection.

Mandiant analysts use the Targeted Attack Lifecycle framework to model attacker behaviors. Over thousands of engagements, the framework has provided a common nomenclature to describe both individual incidents and attacker methodologies. During an investigation, Mandiant consultants will often collect forensic artifacts associated with specific stages of the Targeted Attack Lifecycle. Mandiant consultants also use the framework to develop a structured narrative for reporting investigative findings.

Figure 7.
Targeted attack
lifecycle framework.



In 2020, Mandiant consultants were tasked with scoping the initial stages of SUNBURST activity across many complex and disparate environments. The ability to tie specific indicators and methodologies to the Targeted Attack Lifecycle allowed them to apply a tiered approach to triage potential incidents. Modelling UNC2452's activity not only allowed consultants to respond appropriately but, more importantly, helped highlight how organizations can prepare for and potentially detect UNC2452 and the kinds of activity advanced actors can bring to bear during a complex intrusion. FireEye Mandiant recently published guidance on investigating and remediating UNC2452 activity within Azure cloud environments.⁶



Initial Compromise

In the Initial Compromise phase, the attacker successfully executes malicious code on one or more systems within an environment. UNC2452 targeted SolarWinds and implanted SUNBURST into the build cycle of the SolarWinds Orion product. UNC2452 managed to subvert the build process for SolarWinds Orion from March 2020 through June 2020; the trojanized binary maintained the normal trappings of a legitimate software package. SolarWinds reported that 18,000 of their approximately 33,000 Orion customers downloaded the trojanized dynamically-loaded library (DLL) which constitutes SUNBURST. UNC2452 appeared to target environments selectively based on the profiling data submitted to the Command and Control (CnC) nodes.⁷ It is unclear whether UNC2452 would target victims based on industry vertical or had specific organizations in mind, but the ability to scope this phase of the lifecycle accurately is a critical step for security teams.

6. FireEye (January 19, 2021). [Remediation and Hardening Strategies for Microsoft 365 to Defend Against UNC2452.](#)

7. SolarWinds (January 29, 2021). [SolarWinds Security Advisory.](#)

After initial installation, SUNBURST potentially sleeps for up to 14 days, attempts to interrupt security instrumentation prior to launch and then starts beaconing via DNS requests for avsvmcloud.com. These requests include encoded information about the environment from which SUNBURST is beaconing. Any organization that used SolarWinds Orion and updated to the trojanized version was practically guaranteed to have at least attempted DNS requests (Stage 1).

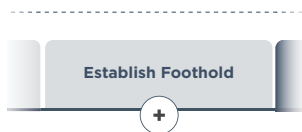
Scoping the Initial Compromise phase from the perspective of the endpoint focuses heavily on the SolarWinds Orion endpoints themselves. Security teams should work to identify evidence that the malicious hotfix served by the SolarWinds update process was downloaded and activated on their SolarWinds Orion endpoints. Ideally, security teams would have imaged their SolarWinds Orion systems and analyzed the image or investigated the endpoint using forensically sound endpoint detection and response (EDR) capabilities. Fortunately, host based indicators (HBIs) of SUNBURST are fairly straightforward.

During the update process the hotfix package is downloaded to the %PROGRAMDATA%\SolarWinds\Installers\ directory along with a corresponding SolarWinds-Core-v2019.4.5220-Hotfix5.msp.log file in the %PROGRAMDATA%\SolarWinds\Logs\Installer\<<YYYY-MM-DD_HH-MM_SS> directory. After installation, the trojanized DLL associated with SUNBURST is written to the %PROGRAMFILES(x86)%\Solarwinds\Orion\ directory.

The identification of the SUNBURST installer provides two valuable data points for a follow-on investigation. First, confirming the existence of SUNBURST is critical, but defining the earliest evidence of compromise can help drive the pace of the investigation. While it was reported SUNBURST may sleep for 14 days after installation, the check is performed against the NTFS Standard Information Modified timestamp of the BusinessLayer DLL. In some cases, this timestamp can persist from the time the archive was created, making this date less reliable. Therefore, the NTFS Create timestamp for the BusinessLayer DLL is used as the start of the timeline for a compromise in which UNC2452 used SUNBURST as the initial vector.

Passive DNS logging makes Scoping Stage 1 SUNBURST requests a trivial exercise. While it is possible to log requests on a local DNS server within the organization, a more stable solution includes passive network collection of DNS requests. These requests usually include the source and destination of the request, the time of the request and relevant DNS-specific details such as the query, the answer and various status indicators of both. In this configuration the sensor's placement and its effect on the quality of the data being captured are critical.

Depending on the placement of a passive sensor, the data captured may lack visibility, reducing the accuracy of its logs. If a sensor is placed north of a local server that provides DNS and recursive lookups, DNS requests emanating from the server may not indicate the true source of the domain lookup. This limitation may be overcome in relatively small or localized SolarWinds deployments, but can generate more questions in a large-scale deployment. It is more desirable if the sensor sits between the individual endpoint and any DNS resolution service.



Establish Foothold

During the Establish Foothold phase, the attacker seeks to strengthen their position in the environment by installing a persistent backdoor on the just-compromised endpoint. SUNBURST acts as a fully featured backdoor with administrative access to the endpoint and runs as SYSTEM on the SolarWinds Orion endpoint. Consequently, UNC2452 can maintain access to an environment as long as the SolarWinds Orion systems remain active.

During startup SUNBURST checks the ReportWatcherRetry value in the SolarWinds.Orion.Core.BusinessLayer.dll.config file. If it is set to any value other than 3, SUNBURST will launch; a value of 3 in ReportWatcherRetry acts as a SUNBURST kill switch, which may be automated or manual. Mandiant experts analyzed SUNBURST and identified a hardcoded list of IP address blocks that control the malware's behavior. DNS records returning values within these blocks would terminate the malware and subsequently update the ReportWatcherRetry value to prevent further execution. However, Mandiant Consulting engagements have identified instances during which UNC2452 set the ReportWatcherRetry value to 3 after successful acquisition of a secondary means of access to the environment. When security teams identify an inactive ReportWatcherRetry value they may need to determine why and how it was set.

```
<add key="ReportWatcherRetry" value="3" />
```

Figure 8.
Sample setting for
ReportWatcherRetry.

To further interact with the environment, UNC2452 can switch the DNS resolution requests for avsvmcloud.com to a CNAME record response containing a domain to which SUNBURST will subsequently communicate (Stage 2). While the full breadth of the UNC2452 campaign remains opaque, it is safe to assume that even an organization with state sponsorship would not be able to manage the workload of 18,000 compromised networks. As of March 2021, the number of environments that entered Stage 2 is suspected to be in the hundreds instead of the thousands. Security teams can focus on the search for CNAME record responses. Depending on how DNS traffic is monitored, available data may differ, but the best monitoring provides the query, response and response type. If security teams identify CNAME responses associated with SUNBURST Stage 2, their investigation is likely to advance dramatically.

Ideally, Stage 2 CNAME DNS requests would be captured by passive DNS logging within the environment. Secondary sources such as firewall logs can identify potential Stage 2 activity in the environment but they rely on publicly available threat intelligence. While searching firewall logs for the IP addresses of known Stage 2 domains may help identify Stage 2 activity in an environment, it is unlikely that every Stage 2 domain will be publicly disclosed. In fact, Mandiant has identified the use of multiple domains during Stage 2 SUNBURST activity in individual environments. Therefore, security teams should hunt for other indications of attacker activity in their environment consistent with UNC2452 across the Targeted Attack Lifecycle.

In a few engagements, Mandiant has observed UNC2452 use the memory-only dropper TEARDROP, which launches Cobalt Strike BEACON. If Cobalt Strike BEACON was used by UNC2452, the CnC infrastructure was separate from the infrastructure used by SUNBURST. UNC2452 has demonstrated significant caution

regarding actions which would have exposed SUNBURST; this pattern persisted in the instances where TEARDROP was deployed. While Cobalt Strike has a reputation for stealth, UNC2452 appeared to take pains to ensure that if Cobalt Strike was detected, the SUNBURST backdoor would not be burned in the process.

UNC2452 clearly understands both security operations and incident response. An incident's initial attack vector often remains a mystery. UNC2452 was likely counting on this knowledge gap to protect SUNBURST as long as possible. If the Cobalt Strike implant were detected and an incident response process initiated around the event, UNC2452 may have hoped that by avoiding a direct line between SUNBURST and BEACON they could protect the former, a more costly and advantageous tool.



Escalate Privileges

In the Escalate Privileges phase, the attacker seeks to obtain further access to systems and data in the target environment. This can include credential harvesting, keystroke logging or the compromise of authentication systems. Success with credential harvesting depends on the attacker's skillset. UNC2452 is quite able to surmount obstacles, research vulnerabilities and move towards their objective. UNC2452 tries to progress to lateral movement from the SYSTEM privileges provided by a SolarWinds Orion endpoint. In some cases, Mandiant has observed a modified version of Mimikatz used to harvest password hashes from memory. In other cases, Mandiant has observed UNC2452 use DCSync, a technique that simulates the Windows domain controller (DC) replication process to an unauthorized endpoint.


As an investigation moves through the cyclical portion of the Targeted Attack Lifecycle, tracking malware-adverse adversaries relies on an in-depth and well-maintained monitoring strategy. Groups such as UNC2452, which understand both security operation center and incident response procedures, are likely to plan their actions around activities which leave behind the least forensic evidence. For example, the actions performed during a DCSync attack are legitimate actions within an Active Directory (AD) environment. The ability to replicate AD objects is a convenience that can add efficiency to an AD environment. When attempting to identify a DCSync attack, security teams will often, by necessity, focus on events that fall outside of common communications and actions between a DC and another endpoint. Where enabled, the Windows Event ID 4662, which details operations performed on directory service objects, can be used to identify DCSync activity. It can be overwhelming to track all Event ID 4662 events. To better manage the task, security teams should only enumerate those Event ID 4662 events that include values in their Operation > Properties field, such as Replicating Directory Changes All or a combination of the globally unique identifiers (GUIDs) listed in Figure 9.

Figure 9. Control access rights associated with DCSync.

GUID	Control Access Right Symbol
1131f6ad-9c07-11d1-f79f-00c04fc2dcd2	DS-Replication-Get-Changes-All
9923a32a-3607-11d2-b9be-0000f87a36b2	DS-Install-Replica
1131f6ac-9c07-11d1-f79f-00c04fc2dcd2	DS-Replication-Manage-Topology

Unfortunately, Event ID 4662 does not provide the source of any operation request. Security teams will need to carefully scrutinize suspicious events identified through Event ID 4662 analysis and identify any correlated events that point to the source of the DCSync. Searching for instances of access from SolarWinds systems to domain controllers may help reduce the dataset of potential connections. However, UNC2452 has been observed to run DCSync from systems affected by SUNBURST as well as other endpoints.

On the network, DCSync is accomplished through DCE/RPC sessions between the unauthorized system and the DRSUAPI RPC interface on a domain controller. DCE/RPC calls such as DSBind, which creates the context necessary to call other functions, and DSGetNCChanges, which requests AD object updates from the domain controller, are both key indicators for identifying DCSync. Unfortunately, not all customer environments conduct the degree of internal monitoring necessary to map RDP/DCE to actions taken across the network. While egress monitoring has become more common in the last decade, monitoring internal traffic exacerbates issues that have slowed the adoption of comprehensive network monitoring solutions. The UNC2452 attack highlights network monitoring gaps which have historically been acceptable risk areas in an organization's monitoring stance.



Internal Reconnaissance

In the internal reconnaissance phase, the attacker explores the organization's environment to better understand its infrastructure, how and where it stores information of interest, and the roles of critical users. In engagements where attackers are less worried about stealth, Mandiant has often observed threat actors use reconnaissance utilities that generate large volumes of data and create detection opportunities. For UNC2452, the converse has been observed across multiple organizations. UNC2452 was often found to use common Microsoft Windows-related tools. For example, one of the few "noisy" actions conducted by UNC2452 was the enumeration of file system shares. UNC2452 would regularly peruse filesystem contents on remote systems, an activity that generates a considerable amount of noise because of how verbose the SMB protocol can be. In fact, SMB monitoring is rarely implemented for the same reason.

In organizations which did have SMB monitoring, reconnaissance performed by UNC2452 would stand out because regular and in-depth enumeration of network shares sourced from a SolarWinds system are atypical events. While this allowed Mandiant consultants to identify the types of files UNC2452 would target, it also provided insight into the data sources the attacker would avoid on the way to their objective. As of March 2021, Mandiant has not observed UNC2452 target access to customer data, personally identifiable information or financial data.

UNC2452 targeted data stores that acted as central stores of knowledge. Mandiant has observed UNC2452 targeting access to online documentation stores, code repositories, and IT and Infosec file shares. Logging resources associated with such data stores are valuable hunting grounds for identifying reconnaissance activities. Searching for access from uncommon sources or sources known to have been affected by SUNBURST can provide substantial indicators to advance hunting efforts. For example, SMB sessions between SolarWinds endpoints affected by SUNBURST and a file server or administrative share may highlight times when the attacker was particularly active on an endpoint, allowing security teams to identify user accounts accessing data during those times. Similarly, quantifying serial access to documentation stores by individual IP addresses within a set period may lead back to a system compromised by UNC2452.



Lateral Movement

In the Lateral Movement phase, the attacker uses the accounts and knowledge of the network gathered in previous phases to move to additional systems in the environment. Tracking UNC2452's lateral movement is quite challenging, given the large amount of network traffic originating from the SolarWinds Orion platform. Once again, the quantity and quality of logging performed in the environment will determine the likelihood of accurately identifying attacker activity in this phase.

Creating timelines by identifying windows of attacker activity in the environment can help reduce the amount of data to be analyzed. Reviewing common Windows logon event types in aggregated windows event logs for the identified timelines can help pinpoint instances of lateral movement. Every time UNC2452 progresses further into an environment, the process of generating timelines, identifying suspicious activities and hunting must be repeated. Every indicator the security team can identify, no matter how small requires the re-examination of existing datasets through a constantly changing lens of attacker motivations. However, teams can quickly identify patterns undertaken by UNC2452. Every identification of lateral movement becomes an opportunity to further quantify UNC2452 behavior and build a comprehensive understanding of their operations and objectives.



Maintain Presence

During the Maintain Presence phase, the attacker ensures continued access to the environment. Mandiant observed UNC2452 move away from SUNBURST as a persistence mechanism when possible. UNC2452 commonly maintained access to an environment through legitimate VPN sessions. After VPN access was stabilized, UNC2452 sometimes activated the kill switch on SUNBURST and relied solely on VPN access.

In some investigations where Mandiant observed UNC2452 shift to VPN, a distinct pattern emerged. UNC2452 applied the same level of tradecraft to source VPN sessions as they did to other phases of the attack lifecycle. The threat actor not only maintained an individual IP address from which an account would authenticate over the lifetime of the account's use, it also ensured the GeoIP location of the IP address showed an address local to the target environment. In some cases, Mandiant observed the threat actor changing the hostname of the Windows-based virtual machines they were using to that of a legitimate hostname in the target environment. While the actions taken to mask the true origin of the VPN sessions were not costly or difficult, the effect on investigations could be dramatic. Traditionally, incident response engagements use Windows Event Logs as a trusted primary source and efforts to subvert those logs often rely on bulk deletion. While this common tactic limits the ability of security professionals to quantify attacker activity during a specific timespan, it still allows the qualification of attacker activity. Windows endpoints log the action of Event Log deletion to the event log which can become an actionable indicator of compromise. In fact, most incident response firms look for quick wins by enumerating these events early in the investigation. UNC2452 attacked this limitation from the other direction.


Instead of deleting logs, UNC2452 activities reduced the confidence security professionals could maintain in a critical incident response dataset. Instead of a one-to-one mapping of hostname to system, analysts would instead have to consider a one-to-many mapping of hostname to systems—with at least one being the legitimate host. These actions did not alter the intrinsic value or

accuracy of the logs themselves; instead, they made it dramatically more difficult to compensate for the different mappings as the number of scenarios producing this effect increased. A counterfeit hostname for an attacker-controlled system is quite rare. As investigations progress, new findings spawn follow-on actions to pursue deeper analyses. When analysts are not aware that a hostname might be falsified, they may initiate follow-on workstreams which assume the legitimate host is compromised. As analysts review endpoints which have been erroneously identified as compromised, more time is required to disambiguate attacker activity and legitimate activity when in truth attacker activity never originated from the legitimate system. These unnecessary workflows draw valuable resources away from the investigation—a worthwhile effect in the eyes of an attacker. While UNC2452 was able to place roadblocks in the path of investigations, the roadblocks sometimes provided opportunities for detection.

UNC2452 relied on virtual private server (VPS) infrastructure to limit ingress IP addresses to geolocations which matched the target organizations geographical spread, which meant VPN logs could be used for broad scale analysis. The source IP address from which a VPN connection is established can be cross-referenced with public data, including the autonomous system number (ASN) to which the IP address belongs. Mandiant analyzed correlations between IP addresses extracted from VPN logs and corresponding ASNs. From this dramatically reduced dataset, analysts could begin to rule out VPN sessions based on an ASN's owner and quickly target ASNs associated with VPS providers. While some advanced users may connect to a corporate VPN through a VPS, it should be a relatively rare practice. Security teams should ensure they review each VPN session sourced from a VPS provider because Mandiant has observed UNC2452 use multiple VPS providers in individual environments.

By using legitimate VPN credentials, UNC2452 allowed Mandiant to successfully perform “Impossible Travel” analysis, which correlates each VPN logon session's originating IP Address to the geolocation of the IP. Individual accounts with successive logons from locations which are too geographically disparate for a person to travel through using conventional means dramatically reduces the dataset of potentially suspicious VPN sessions. The UNC2452 campaign was conducted during the global lockdown resulting from COVID-19 policies, which meant organizations were likely heavily invested in remote work. This enabled the curation of a more accurate baseline of normal user activity. While UNC2452 used legitimate credentials and logged on from localities which broadly matched the organization's, this technique was not infallible. Mandiant observed instances where accounts for users who were on leave were used by UNC2452 as well as cases where international users would log on from US-based locales.

While the impact of using hostnames taken from targeted environments as VPS hostnames had an appreciable effect on the pace of investigations, the identification of this technique led to novel analysis methodologies. Mandiant analysts were able to use services that perform data collection and historical analysis of changes across the Internet to identify potential systems of interest. One field commonly captured by these services is the RDP SSL certificate which can leak the configured hostname of an endpoint. By querying for endpoints which presented leaked hostnames matching those of legitimate hosts within customer environments, Mandiant consultants were able to identify suspicious VPN sessions and then review VPN log data to generate timelines of attacker activity.



Complete Mission

In the Complete Mission phase, the attacker accomplishes the objectives of the intrusion. In some intrusions this can be intellectual property theft or disruption of business operations. Investigation during this phase of the Targeted Attack Lifecycle depends entirely on the accuracy and consistency of the timelines the analysts previously built based on available data. In environments where UNC2452 was able to acquire VPN access, they could freely access data through common protocols such as HTTP(S), SMB and SSH. The cross-referencing of timeline data such as the VPN IP address with access logs for data stores can help security teams identify individual accesses which result in data loss.

Conclusion

The UNC2452 campaign was extremely challenging to uncover and address. The attacker was sophisticated enough to implant SUNBURST in the widespread and broadly respected SolarWinds Orion platform. While UNC2452's knowledge of operational security was higher than most incident responders are likely to witness first-hand, it doesn't change the mission: security professionals must work to guard against similar attacks from copycat threat actors. Well-designed environment monitoring practices and procedures, along with rigorous investigation methodologies serve as consistent ways to shine a bright light on the actions of advanced attackers.

CASE STUDIES



Insider Threat Risks to Flat Environments



In mid-2020, a technology provider engaged Mandiant to investigate a suspected intrusion in their environment. The client notified Mandiant that an unauthorized user had accessed their development/test and Amazon Web Services (AWS) environments remotely. The client had also identified evidence to indicate the unauthorized user deleted Relational Database Service (RDS) backups within the AWS environment. For this investigation, the insider threat team was engaged to:

- Identify the attacker: Collect and deliver evidence that could identify the unauthorized user.
- Determine any loss of client data: Review all forensic evidence for data theft.

The Mandiant insider threat team investigates client environments using a Follow the Data model that echoes methodologies from the National Institute of Standards and Technology (NIST) and the National Insider Threat Taskforce (NITTF). Their focused data collection allows clients to take necessary legal actions, as coordinated with their third-party legal counsels.

Analysis

The Mandiant insider threat team analyzed AWS CloudTrail logs and Linux authentication logs and identified a timeline of events which detailed the logon activities of the suspected insider threat actor (Fig. 10).

While the threat actor took steps to provide a level of anonymity to their actions, Mandiant experts were able to correlate discrete artifacts across sessions to tie the activity to a single user.

Fifteen days before the employee was laid off, the Linux server logged a remote SSH session from the IP address x.x.x.121 which authenticated using an SSH key associated with the ubuntu account—a default account in Ubuntu-based Amazon Machine Images (AMI). At the time of analysis, the IP address x.x.x.121 was associated with a residential Internet provider.

Three days before the employee was laid off, AWS CloudTrail logged the employee authenticating to cloud infrastructure from the IP address x.x.x.113 using their legitimate credentials. At the time of analysis, the IP address x.x.x.113 was associated with an anonymized VPN provider. During this session, AWS CloudTrail logged the creation of users and as well as an EC2 instance by the employee.

Twelve hours after the employee had been laid off, the Linux server logged an SSH session from the same x.x.x.113 IP address used to access AWS cloud infrastructure; it was authenticated using the private key for the ubuntu account. By correlating individual artifacts across multiple logging sources Mandiant consultants were able to develop an evidence-backed narrative which associated the recently discharged employee with both the anonymous VPN IP and the private key used for certificate-based authentication.

AWS CloudTrail logs and host-based forensic artifacts allowed Mandiant consultants to detail the actions undertaken by the recently laid off employee after the employee had authenticated to the development server as the ubuntu account. Command line history and sudo logs extant on the Ubuntu-based EC2 instance provided key artifacts to drive further investigation into the AWS CloudTrail logs. Mandiant experts identified logs that indicated the installation of the awscli toolset which facilitates the submission of API requests to the AWS API.

Thirteen hours after the employee had been laid off, the ubuntu account used the awscli utility to provision a new AWS IAM user account in the client’s AWS tenant.

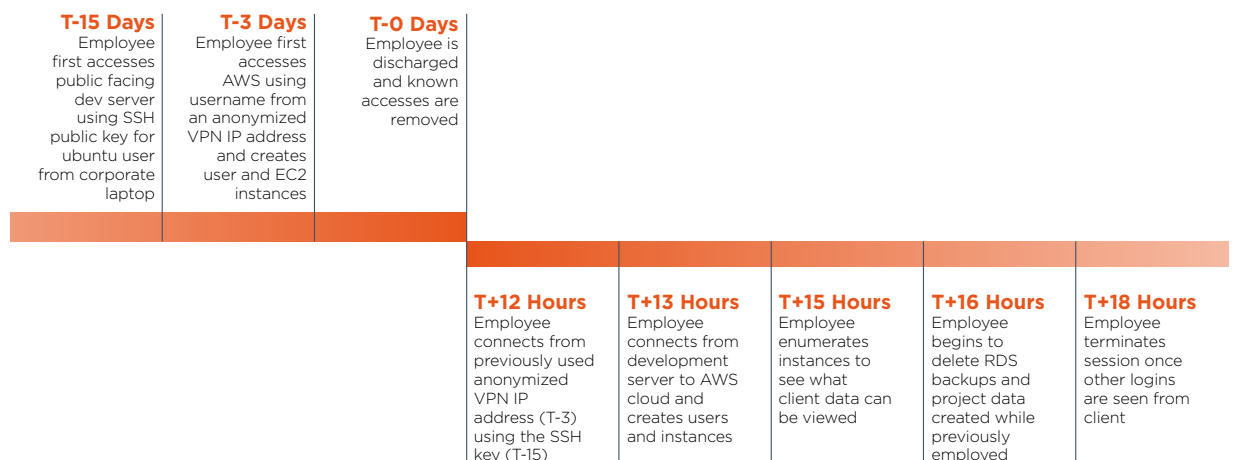


Figure 10. Timeline of events constructed by the insider threat case study.

Three hours later, after failing to enumerate the EC2 instances to which the new AWS IAM account would have access, the ubuntu account issued API calls which deleted RDS backups, removed the accounts it had created and terminated the session. We assumed the enumeration failure was the result of the IAM principal lacking required permissions.

Outcome

The Mandiant insider threat team was able to build a timeline of unauthorized activity, assure the customer there was no evidence of client data loss and subsequently remediate the environment to prevent similar future insider and external threats. At the end of the engagement, the client's environment was more secure.

Recommendations

Credential Revocation

The client believed the insider was not aware of the upcoming layoffs. Mandiant often recommends that clients discharging employees should coordinate access removal with notification to protect both the organization and the employee. This would hold true even if an employee voluntarily resigns or retires. All SSH keys, PEM files, service passwords and application passwords to which the individual had access should be rotated for all environments. Similarly, provisions within multi-factor authentication (MFA) should be immediately unenrolled when an employee or contractor leaves the organization.

Network Segmentation

Historically, Mandiant has found that most client networks enable at least a basic level of network segmentation. However, there often is no separation between the development and production networks. Access controls between development and production networks can limit opportunities for an unauthorized user to move laterally between zones. Mandiant recommends organizations thoroughly review their network segmentation and limit unnecessary traffic between sensitive and less-trusted environments. Also, segmentation at the AWS account level would prevent access to customer data regardless of the principal's IAM permissions. These measures will help prevent insider threats from moving laterally or connecting from an internal network segment to a cloud environment. In general, systems that do not need to be public facing should be segmented from public access and restricted as much as possible.

Logging

In this case, the client had enabled logging and event aggregation to a security information and event management (SIEM) system. This ensured the authenticity of the logs used during the investigation. Linux based operating systems log activity in plain text which provides an opportunity for an attacker to manually edit the log entries generated by their activity. By streaming logs to a secondary source, the client not only followed best-practice logging guidance but maintained critical information necessary for an in-depth investigation.

Least Privilege

This client had many administrative controls for most of its production systems. However, the development network had users with extensive permissions that extended to the creation of accounts in both the client network and the cloud environment. Mandiant recommends organizations implement comprehensive user access controls across all environments on their networks to ensure that user, developer and administrator accounts have only the necessary level of permissions to complete their tasks and maintain business continuity.

Red Team Makes the Most of Social Engineering and System Misconfigurations



Mandiant red team consultants perform targeted, objective-based assessments using a combination of publicly available and internally developed tools. Once the red team is able to access an environment, observing how legitimate employees interact with workstations and applications allows the red team to blend in by adopting behaviors common to the environment.

A financial services firm engaged Mandiant to evaluate the effectiveness of its information security team's detection, prevention and response capabilities.

The Mandiant red team's objectives were to avoid detection while accomplishing the following tasks:

- Access Executive Email
- Access Sensitive Production Network

During this engagement, the red team consultants emulated advanced nation-state attackers regularly seen on the frontlines of Mandiant incident response engagements.

Initial Compromise

While performing open-source intelligence (OSINT) reconnaissance to identify, catalogue and analyze the client's Internet-facing infrastructure, the red team discovered a job application portal through which interview candidates could submit documents. According to Mandiant threat intelligence, social engineering is one of the most common initial attack vectors used by advanced attackers. For this engagement, the red team used the job application portal to deliver malicious documents which might be opened by employees during regular business operations. This allowed the red team to use a trusted resource without having to bypass email-based security controls. To increase the chance the resume would be opened, the red team crafted a suitable candidate for employment based on common job openings published by the customer and submitted the job application and malicious resume through the application portal.

The resume contained a malicious macro and a request that attempted to convince the user to enable macro content to allow the execution of the embedded code on their endpoint. Historically, the red team has found that applying personal touches to attempts at social engineering increases the chance of success. In this case, the red team sent a follow-up email to the HR department to demonstrate enthusiasm from the applicant and check on the status of the resume.

Establishing a Foothold and Maintaining Access

Shortly after sending the follow up email to HR, a recruiter working for the customer opened the malicious resume and enabled content, which provided a Cobalt Strike command and control connection. The successful pre-text and macro execution enabled the red team to gain a foothold into the environment.

Once a foothold is established APT groups often prioritize creating a means through which consistent access to the environment can be maintained. After establishing a successful Cobalt Strike session Mandiant's red team sought to do the same. To maintain access, the red team installed a Startup Folder persistence mechanism within the profile of the compromised user, which would launch the backdoor each time the user logged on to the compromised system.

Escalating Privileges and Lateral Movement

The ability to quickly reconnoiter an environment while remaining undetected is a capability which often separates opportunistic attackers from advanced attackers. The Mandiant red team uses tradecraft common to APT groups to profile the behaviors of legitimate users and reduce the chance of detection. The red team performed host-based reconnaissance on the HR user's workstation to gain situational awareness regarding the user's privileges and systems to which the user had access. This helped the red team identify an unquoted service path on the system which could be used to escalate privileges. An unquoted service path allows attackers to intercept execution flow by placing a binary in a directory higher than that of the intended executable. The red team exploited this misconfiguration to gain elevated access to the system. Using the elevated privileges, the red team used Mimikatz to access the local administrator password hash on the system.

While the privileges acquired through the unquoted service path vulnerability provided a greater level of access to the local host, it did not directly facilitate moving deeper into the environment. The development and adoption of the Microsoft Local Administrator Password Solution (LAPS) has dramatically reduced the attack surface associated with the reuse of Local Administrator passwords. However, deployment of LAPS can be inconsistent across environments. Looking through the client's Active Directory environment, the Mandiant red team enumerated systems that were not managed by LAPS by querying for the absence of the "ms-Mcs-AdmPwdExpirationTime" computer object attribute. This produced a list of endpoints to which the red team could potentially authenticate using the recently acquired local administrator password hash. The red team tested the validity of the compromised Local Administrator account and identified several servers on which the credentials were valid. One of those servers was configured with the unconstrained delegation property which can allow attackers with administrative access to impersonate any account which has successfully authenticated to that system. To force the client's domain controller to authenticate to the just-compromised unconstrained delegation server, the red team used the 'printer bug' attack which sends an

'RpcRemoteFindFirstPrinterChangeNotification' request to a domain controller running the Print Spooler service. As a result of this request, the domain controller tests the connection, which provides the Kerberos Ticket-Granting Ticket (TGT) for the domain controller machine account.

The red team used the captured Kerberos TGT to impersonate the client's domain controller machine account and perform a DCSync attack. A successful DCSync attack provides the NTLM password hashes for any account in the target domain, including domain administrator users. The password hashes acquired through DCSync allowed the red team to move through the client's environment freely and focus on their access objectives.

Accessing High-Value Objectives

During the initial scoping for the engagement, the client tasked the Mandiant red team to gain access to the contents of an executive's email account and move laterally into a sensitive production network.

Executive Email Access

Due to newly adopted work from home practices paired with heavy adoption of cloud services, the client's employees were not centrally located as would be expected with a traditional office model. Their Office 365 email access was protected by multi-factor authentication (MFA), so the red team chose to target executive workstations for email access. The NTLM hashes gathered during the DCSync attack were processed and fed through an offline password-cracking server internal to Mandiant, which returned plain text credentials for an executive account. The red team then monitored the environment to see when executive workstations were connected to the internal network either directly or through a Virtual Private Network (VPN) concentrator. Once an executive's workstation was connected to the internal network, the red team tested the validity of credentials by mapping the C\$ share and gained interactive access to the endpoint—and the user's email—with the executive's compromised credentials.

Accessing Sensitive Production Network

Access to the production network also presented substantial obstacles for the red team. The customer required all connections between the corporate and production networks to traverse a jumphost protected by MFA. To identify the jumphosts that bridged the corporate and production environments, the red team performed reconnaissance within Active Directory and identified a jumphost that allowed inbound Remote Desktop Protocol (RDP) and Windows Management Instrumentation (WMI) connections. While RDP would require the use of MFA, access via WMI was not similarly constrained. Using the Domain Administrator account, the red team queried WMI on the jumphost to identify a list of users who had recently authenticated to the server. The red team then conducted DCSync attacks against the list of recently authenticated users to obtain the NTLM hashes for each account, which were submitted to a Mandiant-internal password cracking service. The cracking service successfully identified the plain text credentials of multiple users from the list, giving the red team the first element needed to authenticate to RDP. While MFA provides a high degree of security for environments, convenience features that help improve adoption rates are often built into the products. One such feature is the ability to generate a push notification on the user's device and request authorization. Advanced attackers and red teams alike can use this feature and hope that a busy employee may assume a request is legitimate and authorize the request.

In this case, the red team initiated a Duo Push request that was accepted by the end user, which gave the red team interactive access to the jumphost through an RDP connection. To perform follow-on actions efficiently and move further into the sensitive production network, the red team needed to place a Cobalt Strike beacon on the jumphost but firewall policy restricted the transfer of files to the jumphost. Using a compromised user workstation, the red team temporarily stored a DLL sideloading payload on a corporate-wide file share which could then be accessed from the jumphost while the red team was connected through RDP. With a Cobalt Strike beacon in place, the red team could maintain access to the jumphost and complete the client's requested objective.

Outcomes

In this case study, at the client's request, the Mandiant red team gained a foothold in the client's environment, obtained full administrative control of the company domain, accessed executive email and connected to the sensitive production network without any software or operating system exploits. Instead, the red team focused on identifying system misconfigurations, conducting social engineering attacks and using the client's internal tools. They were able to achieve their objectives despite the obstacles presented by the client's MFA, network segmentation and employee social engineering awareness.

CONCLUSION



More Security Awareness to Build Best Practices



2020 was a year where we were reminded how events in the physical world and cyber security are intertwined. In past M-Trends we reported on how geopolitical events often have repercussions on cyber security. In 2020 we saw how a global pandemic changed business operations and, as a result, the attack surface and risk profile of most businesses. Organizations around the world struggled with adapting to the new norm and maintaining their defenses as attackers took advantage of these unprecedented times.

Over the past year we were also reminded of the complexity and impact of supply chain attacks, a trend we first mentioned in M-Trends 2013, where we discussed how attackers used third-party service providers to compromise their victims. While many of the trends we have observed in 2020 are not new, we have seen these issues reach new and memorable levels of sophistication and proportion.

We also witnessed how ransomware evolved to multi-faceted extortion and continues to escalate. In 2020, we saw the rise of “name and shame” websites in addition to encryptor deployment. Threat actors capitalized on infrastructure deployed to support a remote workforce by exploiting new and old vulnerabilities for initial access. These trends underscore the importance of sound fundamentals such as vulnerability and patch management, least privilege and hardening.

Security organizations need to continue to be prepared for ongoing escalations with threat actors and deal with changes to their own environment and attack surface. While much has stayed the same, we are seeing a continued evolution of past trends that requires security teams to remain vigilant, adapt and evolve.

To learn more about FireEye, visit: www.FireEye.com

To learn more about Mandiant Solutions, visit: www.FireEye.com/mandiant

FireEye, Inc.

601 McCarthy Blvd. Milpitas, CA 95035
408.321.6300/877.FIREEYE (347.3393)
info@FireEye.com

©2021 FireEye, Inc. All rights reserved. FireEye and Mandiant are registered trademarks of FireEye, Inc. All other brands, products, or service names are or may be trademarks or service marks or service marks of their respective owners.
M-EXT-RT-US-EN-000372-03

About FireEye

At FireEye, our mission is to relentlessly protect organizations with innovative technology, intelligence and expertise gained on the frontlines of cyber attacks. Learn how at www.FireEye.com.

About Mandiant Solutions

Mandiant Solutions brings together the world's leading threat intelligence and frontline expertise with continuous security validation to arm organizations with the tools needed to increase security effectiveness and reduce business risk.